Agricultural Development and Structural Change, Within and Across Countries*

JACOB MOSCONA†

May 5, 2019

Abstract

This study exploits rapid technological development during the Green Revolution (1960-1990) to estimate the causal effect of agricultural productivity growth on structural change both within and across countries. I use variation in ecological characteristics that determined the maximum potential impact of new crop-specific technologies on agricultural productivity to construct an instrument for agricultural productivity growth. Across districts in India, agricultural productivity growth spurred income growth, employment, and land use in the agricultural sector; it also reduced urban development and manufacturing employment. Across countries, agricultural productivity increased specialization in agricultural production and reduced urbanization. I find no evidence that agricultural productivity growth increased national income on average. Estimated effects are most pronounced for districts and countries that were more open to trade in 1960 and had a negative impact on income in countries that were most open.

JEL Codes: J43, O13, O14, O33, O47, Q10, Q16.

Keywords: Agricultural productivity, structural change, economic development, Green Revolution.

*I especially thank Nathan Nunn for his support and guidance from the project's outset. I am grateful to Daron Acemoglu, Abhijit Banerjee, Esther Duflo, Claudia Goldin, Eduardo Montero, Ben Olken, Jonas Poulsen, Carl Riskin, James A. Robinson, Karthik Sastry, Andrei Shleifer, and members of the Harvard Economic History Workshop and M.I.T. Political Economy Lunch for extremely helpful feedback. First Draft: May, 2016.

†Department of Economics, M.I.T., 50 Memorial Drive, Cambridge, M.A. 02142, U.S.A. (email: `moscona@mit.edu`; website: `http://economics.mit.edu/grad/moscona`)

1 Introduction

How economies respond to productivity growth in the agricultural sector has been and remains a central question in economic development. In *The Stages of Economic Growth*, Walt Rostow averred, "[R]evolutionary changes in agricultural productivity are an essential condition for successful take-off [since] modernization of a society increases rapidly its bill for agricultural products" (Rostow 1960 p. 8). Numerous scholars—in large part based on the experience of the British industrial revolution—have suggested that agricultural productivity growth is a necessary pre-requisite for the expansion of manufacturing and overall economic growth.[1] The relevance of this conventional wisdom for the 20th and 21st centuries, when industrial development may be less reliant on local demand, is unclear.

Recent empirical work has focused on within-country variation in agricultural productivity and found a limited or negative relationship between agricultural productivity and non-farm growth.[2] The consequences of these regional dynamics for national development, however—the focus of Rostow and his contemporaries—are ambiguous. Canonical models of structural change suggest that the sub-national and national implications of agricultural productivity growth could be drastically different. Matsuyama (1992), for example, argues that agricultural productivity growth causes growth in closed but not in open economies. If modern countries, like 18th century England, more closely resemble closed economies than regions within countries, analyses of within-country and country-level agricultural productivity growth could come to opposite conclusions. Alternatively, modern countries' trade integration and access to global markets might lead the country-level relationship between agricultural productivity and structural change to depart from received wisdom.[3]

This study makes two central contributions. First, I provide a rigorous empirical analysis of the consequences of modern agricultural productivity growth caused by the Green Revolution (1960-2000), a transformative episode in the global agricultural sector. Over the course of just a few decades, global annual food grain production doubled and a widely predicted Malthusian famine was warded off (Khush, 2001, Pingali, 2012). Yet the Green Revolution remains highly controversial in both academic and policy debates.[4] This paper estimates the causal effect of the

[1]See Nurkse (1953); Schultz (1953); Rostow (1960); A.H. John (1965); and more recently Irz et al. (2001) and Kogel & Prskawetz (2001).

[2]For example, Foster & Rosenzweig (2004) and Hornbeck & Keskin (2015). Bustos et al. (2016) find the opposite effect studying the release of genetically engineered soy in Brazil after 2003, but a similar effect for genetically engineered maize. In their context, however, genetically engineered soy was "strongly labor saving" whereas most recent technological development in agriculture has complemented labor. This is especially true of the technological advances that drove the Green Revolution—this is discussed in more depth below, in Section 2.1.

[3]Relatedly, in Murphy, Shleifer & Vishny (1989), agricultural wealth spurs industrialization by generating demand for manufactures. It is unlikely, however, that there is a tight relationship between supply and demand of manufactures within local areas or countries that are very open to trade. 20th century economic growth in countries like South Korea and Taiwan was not driven by local demand for manufactured goods, but predominately by foreign demand and export-led growth.

[4]According to political journalist Alexander Cockburn, "Aside from Kissinger, probably the biggest killer of all to have got the peace prize was Norman Borlaug, whose 'green revolution' wheat strains led to the death of peasants by the million" (Quoted here: https://www.theguardian.com/global-development/poverty-matters/2014/apr/01/norman-borlaug-humanitarian-hero-menace-society?fbclid=IwAR3pBkdNRVTn_Vc7x2naQM9TsQwryDnL8OoiPzHiA0c0gR7ZsLxyA12DXSQ.). Prior

Green Revolution on rural development and structural change.

Second, using the same identification strategy, I estimate the impact of agricultural productivity growth on structural transformation at multiple levels of aggregation: first, across districts in India, and second, across countries. India, the focus of the within-country component of the analysis, was the epicenter of twentieth century agricultural productivity growth and variation in productivity growth within India was substantial.[5] This makes it possible both to better understand the equilibrium effects of modern agricultural productivity growth, and to investigate whether different units of analysis—sub-national regions and countries, or even countries that are more or less open to trade—respond differently to a productivity shock in the agricultural sector.

In order to identify the causal link between changes in agricultural productivity and development, I exploit rapid increases in crop-level potential productivity that resulted from the staggered release of new crop varieties during the Green Revolution. Output growth was driven almost entirely by technological development—in particular, the release of new high-yielding crop varieties (HYVs) (Ball et al., 1997). Due to differences in ecological and geographic characteristics, different regions were able to adopt and reap the benefits of Green Revolution technologies with very different levels of success. These differences generate exogenous variation in agricultural productivity growth brought about by the Green Revolution.

I compute the maximum *potential* impact of HYV releases on regional or country-level agricultural productivity using theoretical models of crop yield from the Food and Agriculture Organization (FAO), and use this measure of ”predicted productivity” as an instrument for actual agricultural productivity growth. The FAO model is constructed from time-invariant geographic, ecological, and climatic conditions; crucially, it can be computed assuming either the use of traditional agricultural technologies or the systematic use of modern varieties. The instrument is constructed by aggregating grid-cell-level measures of maximum potential crop yield for a set of staple crops within each region or country, assuming the use of only traditional technology before the Green Revolution and the use of modern varieties after the Green Revolution. Regions with high predicted productivity growth are those for which maximum potential yield using modern technology is substantially higher than maximum potential yield using traditional technology. Thus, variation in the instrument is driven by time-invariant characteristics allowing regions to benefit differentially from the release of modern agricultural varieties.[6]

Using the predicted productivity instrument, I first investigate the impact of the Green Revolution across districts in India. I find that agricultural productivity improvements spurred growth in the agricultural sector and reduced urbanization and manufacturing growth. Agricultural productivity growth increased several proxies for rural public goods provision in Indian districts, including access to education, healthcare, and quality roads, as well as district-level agricultural wages;

academic and policy research on the consequences of the Green Revolution are discussed Section 2.1 and later in the Introduction.

[5]In the state of Punjab, for example, overall agricultural productivity increased by 138% between 1960 and 1980 whereas in some other areas, changes were minimal (based on my calculation from Sahghi et al. (1998)—described below).

[6]Section 3 of this paper is devoted entirely to developing and testing the instrument, as well as verifying its validity.

at the same time, it increased inequality in land ownership.[7] It also led to greater employment in farming and a larger portion of district land devoted to agriculture. Moreover, agricultural productivity growth led to reduction in district-level urbanization and manufacturing employment. Thus, the Green Revolution impeded structural transformation across Indian districts; it led to agricultural development at the expense of urban and manufacturing growth.

The extent to which districts could specialize in agricultural production depended on their openness to trade. To investigate whether openness to trade mediated the impact of agricultural productivity growth on structural change, I test whether the estimated effect of the Green Revolution was larger for more districts with greater trade involvement at the start of the sample period. As a proxy for trade involvement, I use the share of the labor force involved in trade and commerce in 1961.[8] I find that the negative impact of agricultural productivity growth on structural change larger in magnitude for districts that were more involved in trade. This is consistent with a mechanism in which the Green Revolution increased the relative size of the districts' agricultural sectors, particularly in open districts that could specialize more completely, by increasing their comparative advantage in agricultural production.

Next, I turn to the cross-country analysis, which in theory and as suggested by the heterogeneous district-level effects, could yield quite different results from the within-India analysis. However, I find estimates that are qualitatively very similar to the results from within India, albeit intuitively smaller in magnitude.[9] Agricultural productivity growth from the Green Revolution led to greater fractions of land and labor devoted to agriculture and lower levels of urbanization at the national level. I find no evidence that the Green Revolution had a positive impact on national total or per-capita income on average: the estimated relationship is in fact negative, and economically meaningful positive values occupy only a small fraction of the estimated confidence intervals.

Consistent with more affected countries developing a comparative advantage in agricultural production, agricultural productivity growth reduced agricultural imports and increased (albeit insignificantly) agricultural exports. Moreover, all baseline effects are more pronounced for countries that were more open to trade in 1960 (i.e. countries for which trade was a larger share of GDP). Among the countries for which trade data in 1960 are available, agricultural productivity growth had a positive effect on income in the 25% least open countries, while it had a significant and neg-

[7] Due to the absence of data on income and consumption during the early part of the sample period period, these are commonly used proxies for district-level rural wealth (see Banerjee & Iyer, 2005, p. 1199). Similar measures for the urban sector do not exist during the early part of the sample period (see Burgess et al., 2017 p. 19), which further motivates studying the effect of agricultural productivity on income at the national level in the second part of the analysis. Even if manufacturing wage data were to exist, cross-country analysis is still important in order to determine the impact of the Green Revolution on income. A Ricardo-Viner (specific factors) model with labor movement across sectors would predict a manufacturing wage increase in tandem with the agricultural wage and with the relative growth of the agricultural sector. The effect on overall income, however, could still be negative.

[8] While this is an imperfect measure of openness to trade, it does not simply stand in for the pre-period share of the labor force outside of agriculture. Indeed, I do not find that districts with differing pre-period employment in non-trade/commerce, non-agricultural occupations responded differently to the agricultural productivity shock.

[9] The only difference in the identification strategy is that I use a broader set of crops to construct the instrument in order to accommodate the broader geographic diversity globally. This is discussed in detail in Section 5.

ative impact on per-capita and total GDP in countries that were more open to trade.[10] While at first glance this finding may seem counterintuitive, agricultural productivity growth can have a negative impact on income if there are "learning-by-doing" or other externalities in urban sectors, a focus of existing theoretical work (e.g. Matsuyama, 1992, Young, 1991). Recent empirical analysis has found substantial evidence of these externalities, making the theory's proposed mechanism a plausible explanation for my findings.[11] Intuitively, I estimate a negative impact of productivity growth on income in the set of open countries where productivity growth also reduced urbanization and increased specialization in agriculture; for countries that only experienced the positive direct effect of agricultural productivity growth and no countervailing decline in urban development, I find productivity growth increased national income.

Taken together, these results demonstrate that the Green Revolution has, on average, led both regions in India and countries to specialize in agricultural production. Specialization in agriculture came at the cost of urban development. The findings also highlight the role of trade in mediating the impact of agricultural productivity growth, both across districts and countries; openness to trade may be a key feature that distinguishes subnational analyses from country-level analyses of structural change. It may also be part of what sets the modern context apart from agricultural productivity growth that led to the English Industrial Revolution, which has guided much existing thought about the role of agriculture in structural change.[12] Moreover, the results indicate that, perhaps because of dynamic externalities in manufacturing, by entrenching countries in agricultural production the Green Revolution had a negative impact on income in many countries. National policy designed to increase agricultural productivity may not be an efficient route to long run growth.

1.1 Related Literature

This study extends existing knowledge about the role of agricultural productivity in growth and structural change. The sub-national component of this analysis builds especially on the seminal work by Foster & Rosenzweig (1996, 2004), who find that the Green Revolution increased returns to schooling and slowed non-farm growth across rural Indian villages. I depart from their work

[10]These results are robust to controlling flexibly for a series of potentially confounding factors that might be correlated with openness, including trends in initial country size and income.

[11]Most recently, Bartelme et al. (2018) and Lashkaripour and Lugovskyy (2018) estimate a positive relationship between the size and productivity of manufacturing sectors across countries, suggesting that manufacturing externalities could indeed generate a negative relationship between agricultural productivity growth and income in the long run. Focusing on India, Patibandla & Petersen (2002), Lall & Chakravorty (2004), and Lall & Mengistae (2005) find evidence of external economies of scale in Indian cities. A range of papers find evidence of plant-level learning by doing and productivity growth over time; for example, Levitt, List & Syverson (2013) and Thompson (2010) for a survey. Rosenberg (1982) cites substantial qualitative evidence. Finally, see Acemoglu, (2009), p. 719 for a broad discussion.

[12]Matsuyama (1991, 1992) also noted the potential issues with focusing on the experience of England. Matsuyama (1991 p. 643), for example, notes that, "[A] takeoff is possible in an economy with less productive agriculture while an economy with productive agriculture will be trapped into the state of preindustrialization. This result, once stated, is intuitive. A low productivity in agriculture implies an abundant supply of 'cheap labor' that the manufacturing sector can rely on." The "conventional law" that there are "positive links between agricultural productivity and industrialization" is based almost entirely on the history of England during the Industrial Revolution (Matsuyama, 1992).

first by developing an instrument for district-level agricultural productivity growth, which means that the estimates remain valid when the adoption of new seed varieties and resulting change in productivity are endogenous.[13] Second, rather than focus on villages, I analyze labor movement between the rural and urban sectors and the district-level economy as a whole. By doing so, I highlight the equilibrium effects of agricultural productivity change, its effect on broader sectors of the economy, and the role of migration and district-level trade exposure.

A range of additional studies investigates the impact of agricultural productivity growth. Subnational analyses from other contexts argue that there are limited or negative effects of agricultural productivity growth on local non-farm development (e.g. Hornbeck & Keskin 2015; Bustos et al. 2016, for maize but not soy). Previous cross-country analyses—which rely primarily on correlational and observational evidence—tend to find positive effects of agricultural productivity on income growth.[14] My baseline results are a departure from these claims, at least in the context of the Green Revolution.

This paper also contributes to a deeper understanding of the impact of the Green Revolution (Evenson & Gollin, 2003 [2]; Dethier & Effenberger, 2012; and Pingali, 2012 provide helpful reviews of recent work; section 2.1 describes prior work in more depth). Most related to my baseline country-level analysis, Gollin, Hansen, and Wingender (2018) have a contemporaneous study, using a very different identification strategy from this paper, that argues that country-level HYV adoption had a positive impact on income growth. Among other differences, their empirical design relies on the assumption that aggregate crop-level HYV adoption and the timing of HYV adoption across countries are exogenous. Data on realized HYV adoption are used to construct the instrument. In this study, I construct the instrument using time-invariant data on local agricultural potential at various input levels; the instrument would remain valid if the timing of HYV adoption and aggregate variation across crops were endogenous. This difference in empirical approach could lead to the differences between our findings if they bias the GHW (2018) IV estimates toward OLS correlations. Indeed, my OLS results closely mirror the OLS and IV results in GHW (2018) (see Appendix Table A9).[15]

[13] The HYV index used in Foster & Rosenzweig (2004) is calculated from information on yield potential from "informed sources in the village" based on past and existing crop yields, whereas the variation in my measure of predicted agricultural productivity is from exclusively baseline geographic and ecological characteristics. Last, Foster & Rosenzweig (2004) focus on a slightly different time period. Their analysis begins in 1971, which is after the release of several HYVs—including HYV rice, wheat, and maize—that play a major role in my empirical results (see Figure 1 of this paper).

[14] Early works by Nurkse (1953), Schultz (1953), Rostow (1960), and A.H. John (1965), and more recently Irz et al. (2001) and Kogel & Prskawetz (2001), argue that agricultural productivity growth is highly associated with—and perhaps a precondition for—industrial development. An exception is McArthur & McCord (2017) who investigate the impact of changes in input use for cereal production globally.

[15] In particular, the time variation in GHW (2018)'s equation (2), and hence variation in the instrument, comes from realized variation in HYV adoption (p. 17). The cross sectional variation in their instrument is pinned down by differences in which crops countries produced during the pre period, which itself may be different for countries on different income trends (p. 18, equation (3)). Both features may bias their IV estimates toward OLS. Another reason their results may differ from mine is related to the heterogeneous effects that I uncover. The "compliers" in their instrumental variables strategy are countries that drive global variation in crop-specific HYV adoption—likely to be larger and generally more closed economies. These are precisely the countries where I find some evidence that the Green Revolution had a positive effect on income.

This paper is linked to a large body of work that studies the impact of resource abundance and resource based specialization since a central finding is that productivity growth in agriculture increased specialization in the agricultural sector (e.g. Cordon & Neary, 1982; Auty, 1993; Sachs & Warner, 2005; and Michaels, 2007).[16] These findings also add to a large literature on the determinants of migration—rural-urban migration in particular—and variation in wages across sectors (e.g. Banerjee & Newman, 1998; Young, 2013; Bryan, Chowdhury & Mobarack, 2014; Munshi & Rosenzweig, 2016; Imbert & Papp, 2016).[17]

My identification strategy builds on a growing body of work in economic development and economic history that uses variation in land characteristics as a source of plausibly exogenous variation impacting technology adoption (e.g. Nunn & Qian, 2011; Alesina et al., 2013; Bustos et al., 2016). Bustos et al. (2016), in particular, use the difference in local soybean potential yield at high and low input levels to estimate the impact of the introduction of genetically engineered soybeans in Brazil during the 1990s.

The paper proceeds as follows. In the next section, I describe the main source of recent variation in agricultural productivity, the Green Revolution, in more depth, along with the data used in the subsequent analysis. Section 3 details the construction of the instrument and presents the first stage results for the sub-national analysis, along with a series of robustness checks. Since the identification strategy is very similar for both the sub-national and cross-country analysis, it is discussed in depth in Section 3 and only briefly in the context of the cross-country analysis. Section 4 presents the main empirical results on the sub-national impact of the Green Revolution in India. Section 5 presents the cross-country first stage and main empirical results. Section 6 concludes.

2 Background and Data

2.1 The Green Revolution

The 20th century's Green Revolution was a period of steep increases in agricultural productivity that resulted from coordinated global research into the development of high yield varieties (HYVs) of a set of staple crops. Before 1960, few major advances in crop productivity had taken place and the high-yield technologies that did exist were largely limited to conditions in the developed world (Evenson & Gollin 2003). While individuals began experimenting with improved agricultural technologies long before the mid-20th century (e.g. van Zanden 1991, Sonnenfeld 1992), during the 1960s a process of rapid and coordinated discovery led to drastic increases in potential crop yield

[16] Michaels (2007) is of particular relevance—studying the impact of oil abundance in U.S. counties during the 20th century, he finds that oil abundance increased county-level per-capita income and in the long run reduced the manufacturing employment share. He shows, however, that oil abundance did not reduce the absolute size of the manufacturing sector and argues that this is because it led to population in-flows. This is distinct from the experience of the Green Revolution in India, where I show that that agricultural productivity growth reduced net inward migration in Indian districts, contributing to the contraction of the non-farm sector.

[17] Interestingly, work in developing countries, including the present study, on the relationship between agricultural employment and structural change stands in contrast to Eckert & Peters (2018), who find that in the US, most spatial reallocation occurred within labor markets.

(e.g. Evenson & Gollin 2003, Foster & Rosenzweig 1996, Sasaki et al. 2002). Khush (2001 p. 815) notes, "It took almost 10,000 years for food grain production to reach 1 billion tons, in 1960, and only 40 years to reach 2 billion tons, in 2000."

Several scholars cite 1960—especially the release of much higher yielding rice and wheat varieties during the mid-1960s—as the start of this period of rapid change. A series of institutional changes occurred during the 1960s that encouraged investment in agricultural research and that may have resulted in the rapid production of higher yielding crop varieties. Before the 1960s, there was "no effective intellectual property of crop varieties...[but] in the 1960s, Plant Breeders' Rights were developed in order to provide incentives for private breeding programs" (Evenson & Gollin 2003 p. 2). This institutional adjustment increased potential profits from private sector investment. At the same time, international organizations, backed by an international set of donors, formed to coordinate the development of high yielding crops. This resulted in the establishment of several international agricultural research centers (IARCs) which ultimately coalesced to form the Consultative Group for International Agricultural Research (CGIAR). This combination of private and public sector institutional changes increased crop breeding research and boosted the release novel crop varieites.

The canonical example of targeted development of high yield technology has become the release of genetically engineered HYV rice in 1966. Genetic engineering allowed wet and dryland rice plants to produce significantly greater quantities than had previously been possible. The first, and arguably most game-changing breakthrough in rice technology was the the IR8 variety—so called "miracle rice"– developed during the early 1960s by Peter Jennings at the International Rice Research Institute (IRRI) in the Philippines. IRRI, one of the IARCs, was established by the Ford and Rockefeller foundations in 1960 to promote crop variety research. The striking difference between IR8 and pre-existing rice varieties was first noted in 1966, when S. K. De Datta documented that while pre-existing varieties produced on average one ton of rice per hectare, the IR8 variety could produce 5 tons per hectare even without the use of fertilizer. Moreover, IR8 took 30-40 fewer days to mature than other rice varieties.

Distinct from earlier periods of agricultural development, agricultural technologies developed during the Green Revolution had several complementarities with labor. The "labor intensive" work of modern variety crops has been noted by several scholars (Ladejinsky 1970, Hossain 1998). More recently, researchers have sought to address several shortcomings of the Green Revolution and, in particular, the fact that Green Revolution technologies required high amounts of labor, preventing the movement of labor towards non-farm industries. For example, Lal et al. (2016) propose a new "Evergreen Revolution that would introduce technologies that would be...labor saving" and that would "tend to free labor for nonfarm uses" (p. 438). While the Green Revolution indeed made agriculture more productive, significant amounts of labor were required to implement new techniques and utilize new technologies.

Polarizing debate on the impact of the Green Revolution remains far from conclusive—few development projects have been simultaneously so praised and decried. The Green Revolution has been credited with drastically reducing hunger and poverty in the developing world and with countries' ability to support rapidly growing populations (e.g. Davis et al. 2010, Pingali 2012, Rao 1985, Thirtle et al. 2003, Evenson & Gollin 2003, Frankel, 1971). At the same time, it has been associated with higher poverty levels within India (e.g. Griffin 1974, Harriss 1977), with lower local industrial development (e.g. Foster & Rosenzweig 2004), with growing social inequality in several countries (e.g. Freebairn 1995 for a review of early work; Junankar 1975 for India, Sonnenfeld 1992 for Mexico, Niazi 2004 for Pakistan) and with higher levels of political discord and violent conflict (e.g. Ladejinsky 1970, Frankel 1971, Shiva 1991). However, previous work has focused almost entirely on local effects of technological change in agriculture, and provides predominately correlational and anecdotal evidence.

2.2 Data: Constructing the Instrument

To estimate the impact of changes in agricultural productivity, I exploit variation in exposure to the impact of the Green Revolution to calculate a measure of "predicted productivity." I then use predicted productivity as an instrument for actual agricultural productivity at both the district and country-level. The predicted productivity measure relies on theoretical models of maximum potential crop yield from the Food and Agriculture Organization Global Agro-Ecological Zones (FAO GAEZ).[18] These data:

> "...reflect yield potentials with regard to temperature, radiation and moisture regimes prevailing in the respective grid-cells. The model requires the following crop characteristics: Length of growth cycle (days from emergence to full maturity); length of yield formation period; maximum rate of photosynthesis at prevailing temperatures, leaf area index at maximum growth rate; harvest index; crop adaptability group; sensitivity of crop growth cycle length to heat provision; development stage specific crop water requirements, and coefficients of crop yield response to water stress" (FAO GAEZ).

Crucially, the FAO potential yield model is constructed using parameters derived from controlled experiments, and not from data on actual agricultural inputs and output (see Costinot & Donaldson 2016, p. 18).[19] The data are reported by the FAO as a $9.25km \times 9.25km$ raster grid, with each grid cell containing the maximum attainable yield for a given crop in that area.

The data are constructed at both "low" and "high" input levels. The low-input level version assumes the general use of "traditional cultivars" in agriculture while the high-input level assumes the use of high-yield variety crops (FAO GAEZ). While the low input level data account for the fact

[18] As noted in the Introduction, FAO data have been used in several recent works to estimate the variation in suitability for agricultural technologies or individual crops, including Qian (2008), Nunn & Qian (2011), Alesina et al. (2013), Fiszbein (2015), Bustos et al. (2016), Costinot & Donaldson (2016).

[19] The FAO produces two basic measures of potential crop growth: maximum potential yield (described in the quote above) and a crop suitability index. I calculate instrument values using the potential yield measure because it seems less likely to be influenced by endogenous district-level characteristics or growth patterns. The suitability index is calculated using "edaphic rating[s] for each soil/slope combination" along with soil type. It seems likely that soil quality and local terrain may have been influenced by human behavior and technology. Soil quality does not factor into the FAO calculation of potential yield.

that there might be some unsystematic use of early high-yield technologies—as indeed there was before 1960—it makes the assumption that they are "treated in the same way as local cultivars" and that the use of these technologies was not coordinated (FAO GAEZ). I also collect information about timing of high yield variety release, primarily from a series of crop reports compiled in Evenson & Gollin (2003) and from Dalrymple (1986).

Since this paper's main results focus on long-difference specifications, they do not hinge on the exact HYV release years that were identified; however, HYVs seem to have been adopted rapidly and broadly following their release. The fraction of India's rice, wheat, and maize cropland devoted to HYVs between 1957-1987, averaged across the 271 districts in that comprise my sample, is presented in Figure 1 (these data are from Sanghi et al., 1998). The vertical dotted line in each graph indicates the HYV release year that I identified. In all cases, the fraction of cropland devoted to the HYV began to increase in the indicated year. By 1981, on average, the fraction of a district's land devoted to rice cultivation in which HYVs were used was over 40%. The proportion was even greater for wheat and maize. In India and other developing regions, HYV adoption was widespread and quickly followed technological innovation.[20]

2.3 Data: Actual Productivity and Outcome Variables

For the analysis of Indian districts, actual agricultural productivity is constructed from the India Agriculture and Climate Dataset, compiled by Sanghi et. al. (1998), which calculates yield measures for 5 major and 15 minor crops at the district level for 271 districts from 1957-1987. These data have been used extensively in the economics and political science literature on agriculture in India.[21] Crop-level agricultural output in Sanghi et. al. (1998) is measured in hectograms per hectare. In order to calculate a consistent weighted measure of agricultural productivity that incorporates nutritional content, I matched each crop in the data set to United States Department of Agriculture (USDA) data on calorie content to estimate the caloric load of a hectogram of each agricultural product. The district-level productivity measure is the sum of the output of all crops divided by the area under cultivation for all crops.[22] Calories are a convenient unit to aggregate productivity across crops, particularly since data on producer prices are scarce during the early part of the sample period, especially for the cross-country part of the analysis; moreover, crop-specific prices

[20]This pattern of rapid and widespread adoption was not confined to India. Evenson & Gollin (2003) find that globally the adoption of modern varieties—aggregated across all crops—had reached 46% by 1990. The figure was 63% by 1998. The adoption rate was greater in many parts of East and South East Asia—for example, by 1981, the proportion of rice growing area devoted to high-yielding varieties was 74.5% in Indonesia (p. 50). The proportion in 1981 was 45.4% in then-Burma (Dalrymple 1986, p. 41), 43.5% in Pakistan (p. 56), 78.3% in the Philippines (p. 60), and 70.8% in Sri Lanka (p. 62).

[21]For example, Duflo & Pande (2007), Edmonds, Pavcnik & Topalova (2010), Iyer (2011), Dasgupta (2014), Taraz (2017).

[22]Crops included are: Bajra (pearl millet), jowar (sorghum), maize, rice, wheat, barley, groundnut, gram, "other pulses," potato, ragi, tur, and soybean. Since the focus of this paper—and of the Green Revolution—is the impact of productivity growth in consumption crops, I exclude cash crops from my calculation of district level productivity. In the main results, I do not factor tobacco, sunflower, cotton, sesame, sugar, jute, or rapeseed production into the calculation of productivity. However, all results are robust to the inclusion of cash crops in the calculation of productivity and, if anything, results using consumption-crop productivity that I present in the body of the paper are more conservative (see Table A3).

are themselves directly affected by agricultural productivity growth whereas calories are (largely) constant across space and time.[23] Country-level agricultural productivity data are from the FAOSTAT database of the FAO and were similarly matched to USDA data on calorie content in order to calculate a calorie weighted measure of agricultural productivity.

Outcome variables derive from a variety of sources. Data on population composition, employment, migration, and literacy rates in Indian districts are from India's 1961, 1971, and 1981 censuses.[24] For data on public goods provision (schools, education, and roads), I use standardized measures that were computed from India's censuses in Iyer (2010). Data on agricultural wages are from Sanghi et. al. (1998).[25] Country-level outcome variables are from the World Bank Development Indicators, with the exception of data on trade in agricultural goods, which are from FAOSTAT, and data on the share of employment in agriculture, which are from Wingender (2014).

3 Predicted Productivity & Sub-National First Stage Results

3.1 Calculating Predicted Productivity

Construction of the instrument and the logic of the first stage are very similar in the sub-national and cross-country analyses. I describe the instrument, its construction, and the necessary exclusion restriction in detail here—in the context of the sub-national analysis—allowing for a shorter exposition in the cross-country analysis below. Using the data described in section 2.2, I construct the predicted productivity instrument as:

$$P_{it} = \frac{\sum_c [(1 - I_{ct}) P^L_{ci} + I_{ct} P^H_{ci}]}{N_{it}} \tag{1}$$

where i indexes districts, t indexes time, and c indexes crops. P_{it} is predicted agricultural productivity in district i at time t. I_{ct} is an indicator variable that takes on the value 1 in the year that high-yield variety of crop c is released and in all subsequent years. P^L_{ci} and P^H_{ci} are the theoretical maximum potential yield measures from the FAO for crop c in district i at low and high input levels respectively, converted to calorie units. That is, for each crop I use the low-input level version of the FAO's theoretical data for the years before the release of its HYV and switch to the high-input level version for all years after the release of the HYV. N_{it} is the number of crops used in the construction of the instrument that it is possible to grow, according to the FAO GAEZ data, in district i at time t.

For the analysis of India, I calculate P_{it} use the FAO models of rice, wheat, and maize cultiva-

[23]This implies that computing productivity from contemporaneous prices would substantially bias the estimate; moreover, using pre-period prices would add significant measurement error because of the large changes in price resulting from the productivity shock. Nevertheless, the results are similar if I take this second approach, however, intuitively, both the first and second stages estimates are biased toward zero.

[24]These were compiled for analysis by Reeve Vanneman at the India District Database

[25]In order to approximate the real wage, I normalized the nominal wage using a district-level price index constructed from price data in Sanghi et. al. (1998) (discussed below).

tion since in India, a large portion of HYV-induced productivity growth is attributable to these three crops (Evenson & Gollin 2003, Pingali 2012, Hazell & Ramasamy 1991).[26] To incorporate larger geographic and ecological variation, I use a broader set of crops to calculate the country-level version of the instrument (discussed below).

Figure 2a shows the change in the instrument value between 1961 and 1981 for all 271 districts in the sample; Figure 2b shows the corresponding change in actual agricultural productivity. The first stage is presented graphically as a partial correlation plot in Figure 3a with the change in the instrument on the x-axis and the change in actual productivity on the y-axis—the t-statistic is 7.51, a first indication of a strong first stage relationship.[27] Visually, this relationship appears general and does not seem to be driven by a small number of influential observations.

The validity of the instrument relies on the assumption that $Cov(P_{it}\epsilon_{it}) = 0$, where ϵ_{it} is the error term in the second stage regression (equation (3), below). As (1) shows, variation in the instrument is determined entirely by baseline district-level (or country-level) characteristics that influence potential responsiveness to the HYVs of a set of important crops, combined with the timing of global technological developments. Since the indicator variable I_{ct} changes when the HYV for crop c was released globally, in the same year for all districts, potentially endogenous rates of HYV adoption do not bias the instrument. Moreover, a majority of the analysis relies on long-difference specifications, which do not incorporate variation in the timing of HYV releases and rely exclusively on baseline characteristics. It therefore seems unlikely that this instrument is correlated with any regional changes in economic, institutional, or geographic conditions. Nevertheless, I show that each set of results is robust to the inclusion of a broad range controls capturing differences in baseline district-level geographic and economic characteristics.

3.2 "Zeroth" Stage

Part of the logic behind the first stage is that HYVs were more broadly adopted in districts with a larger potential productivity gain from adoption. Table A1 documents a strong, positive correlation between predicted productivity—the instrument—and district-level HYV adoption. This relationship is robust to the inclusion of a broad set of controls and remains strong when rice, wheat, or maize is dropped from the construction of predicted productivity.

3.3 First Stage Estimates

The first stage relationship between predicted and actual productivity is modeled as:

$$X_{it} = \beta P_{it} + \eta_i + \xi_t + \mathbf{Z}'\gamma + u_{it} \tag{2}$$

[26]Clearly, however, predicted productivity can be constructed with any set of crops for which FAO potential yield models exist. In particular, I show in Section 3.3 that the first stage relationship does not hinge on choosing exclusively rice, wheat, and maize by constructing a version of the instrument that also includes sorghum and barley.

[27]Latitude, longitude, initial log agricultural wages, initial population density, and the initial male literacy rate are partialled out.

X_{it} is actual agricultural productivity in district i at time t. All specifications include both time and district fixed effects. The coefficient of interest is β, which measures correlation between the predicted productivity measure and actual district-level output per hectare. $\mathbf{Z'}$ is a vector of controls that changes depending on the specification and u_{it} is an error term.

Table 1 presents the baseline first stage results. All regressions are "long difference" specifications, including just the years 1961 and 1981. These years were chosen since data in Sanghi et al. (1998) covers 1957-1987 and, within that period, 1961 and 1981 are the years in which the census was conducted that are farthest apart. Data from the census are used extensively in the second stage so, in order for the first stage regressions to correspond appropriately to the second stage, in most of the analysis only 1961 and 1981 are considered. Both the instrument and actual productivity are measured in kilocalories of output.

The first column includes exclusively the instrument and the fixed effects on the right hand side and suggests a strong first stage relationship between predicted and actual productivity. The second column adds a set of district-level geographic characteristics—latitude, longitude, and a coastline indicator—interacted with an indicator variable that equals 1 in 1981 ("postyear indicator" interaction). These controls allow for differential trends based on initial geographic characteristics. Despite the inclusion of these controls, the point estimate is slightly larger and the first stage relationship increases in statistical significance (F-statistic = 27.66). The third column adds to these the postyear indicator interacted with the 1961 adult male literacy rate, population density, and (log of) 1961 agricultural wages, potential proxies for development at the start of the period. The first stage coefficient magnitude and statistical significance remain similar. Despite the inclusion of these controls, designed to capture the role of initial geographic and economic characteristics on trends in agricultural productivity, the first stage relationship remains robust.

A remaining potential concern is that variation in the instrument is part determined by agricultural productivity at the start of the period. If initial agricultural productivity impacted subsequent development, it could bias the IV estimates. First, it is important to note that both pre-period and post-period values of the instrument are based on theoretical models of potential yield, calculated based on topological and climatic characteristics. As a result, values of the predicted agricultural productivity measure are not determined by potentially endogenous pre-intervention agricultural productivity or any "real world data." Second, I control directly for initial agricultural productivity in the first-stage regression model to show that the strength of the instrument does not hinge on district-level differences in initial agricultural productivity. Column 4 of Table 1 adds to the existing controls an interaction term between actual agricultural productivity in 1961 and a post-year indicator. The magnitude and significance of β remain virtually unchanged despite the inclusion of this control.

Reassuringly also, the magnitudes of the coefficient β presented in Table 1 are logical. Whereas the instrument is a measure of maximum attainable yield based on geographic and ecological characteristics, the outcome X_{it} is the yield that was actually attained in district i at time t. Therefore, it makes sense that $\beta \in (0, 1)$ in all first stage specifications.

3.4 Robustness & Falsification Tests

If the instrument is valid, it should not be correlated with changes in agricultural productivity before the Green Revolution. Figure 3a highlights the strong relationship between the change in actual agricultural productivity from 1961-1981 and the change in predicted productivity from 1964-1981 ($\beta = 0.111$, t-stat=5.82).[28] Figure 3b plots the change in actual agricultural productivity from 1957-1964—prior to the release of HYVs—against the change in predicted productivity from 1964-1981.[29] To my knowledge, district-level data on agricultural output in India do not exist prior to 1957 so the y-axis in Figure 3b represents the change in agricultural productivity over just seven years.[30] In Figure 3b, the relationship between predicted and actual productivity is far from significant and very small in magnitude ($\beta = 0.006$, t-stat = 0.49). This suggests that it is unlikely that the first stage relationship is driven by pre-existing trends in agricultural productivity.

The first stage results in Table 1 do not allow for the inclusion of lead or lagged variables on the right hand side. Moreover, they assume that contemporaneous predicted productivity impacts actual productivity. If, however, future predicted productivity is correlated with current actual productivity, this would cast doubt on the validity of the first stage. While Figure 3 suggests this is not the case, it is useful to check in regression form. Table A2 presents a series of results using a panel that includes a single observation every five years in order to allow for the inclusion of lagged and lead variables. All specifications include district and year fixed effects along with a full set of time indicator interactions with latitude, longitude, a coastline indicator, log agricultural wages in 1961, population density in 1961, and the adult male literacy rate in 1961.

For reference, the first column includes just the instrument on the right hand side (along with the controls) to show the baseline first stage when a panel, rather than long difference, is used. Column 2 of Table A2 includes a five-year lag of the district literacy rate and agricultural wages in order to control for recent "wealth" dynamics. The coefficient of interest remains virtually unchanged. Column 3 adds a five year lead of the predicted productivity instrument to the right hand side. If the future value were significant, it would cast doubt on the assumption that changes in HYV suitability that are used to construct the instrument lead to changes in actual agricultural productivity. While the current value remains significant, the lead value is insignificant and low in magnitude (coef.=-0.007. This indicates that a pre-existing trend in agricultural productivity does not drive the first stage relationship.

Column 4 uses an instrument constructed using two additional crops whose modern varieties Sanghi et al. (1998) suggest were in broad use in India (although their yield did not change as dramatically as those of rice, wheat and maize): sorghum and barley.[31] The first stage relationship

[28] All graphs in Figure 3 are partial correlation plots that conditions on latitude, longitude, log of district wages in 1961, population density in 1961 and adult male literacy in 1961.

[29] Among the rice, wheat, and maize, the earliest high yield variety release was for maize in 1965. As a result, the instrument value in 1961 is equivalent to the instrument value in 1964.

[30] A version of 3(a) where the outcome is the actual productivity change from the seven years after the first introduction of HYVs, 1964-1971, looks similar to 3(a) suggesting that the difference between 3(a) and 3(b) is not driven by the difference in time scale.

[31] Sorghum and barley are also the only two additional crops in the Sanghi et al. (1998) data that also appear on an

using this five-crop instrument is highly significant, suggesting that strength of the first stage does not rely on the instrument's restriction to rice, wheat and maize. The remaining columns repeat the same five specifications measuring predicted and actual productivity in kilocalories per hectare—the results are qualitatively identical.

4 District-Level Results

The primary estimating equation for this section is:

$$Y_{it} = \pi X_{it} + \eta_i + \xi_t + Z'_{it}\gamma + \epsilon_{it} \tag{3}$$

X_{it} remains agricultural productivity (kcal/ha) and it is instrumented using predicted productivity, P_{it}. π is the coefficient of interest. The outcome variable, Y_{it}, changes across specifications. District and year fixed effects (η_i & ξ_t) are included in all regressions and Z'_{it} is a vector of controls, noted at the bottom of each table column. ϵ_{it} is an error term. Because of the likely bias of OLS results in this context, I discuss exclusively 2SLS specifications in the main results. However, OLS estimates are reported for reference in Table A3. Reduced form estimates, which tell the same story as 2SLS, are presented in Table A5.

4.1 Rural Public Goods and Wages

While district-level measures of individual income or consumption for the early part of the period under investigation are not available, there are several proxies that have been utilized in earlier work (e.g. Banerjee & Iyer 2005, Iyer 2010). These measures include village-level access to roads, healthcare, and schooling, as well as agricultural wages. Importantly for interpretation, these measures are predominately proxies for rural income. Similar district level measures that better capture income and welfare in the urban sector do not exist for the period under investigation (see Burgess et al., 2017 p. 19).[32]

Table 2 documents a generally positive impact of agricultural productivity on agricultural wages and rural public goods provision. All regressions include the baseline geographic controls interacted with a postyear indicator on the right hand side. Data on each outcome variable are missing for some districts, so the number of observations changes slightly. Each regression uses a balanced panel of districts for which the data are available.

In Column 1, the fraction of a district's villages with access to a road is the outcome variable. The coefficient of interest, π, is positive and significant ($p<0.01$). In the second column, the fraction of villages with a medical center is the outcome variable. The outcome variables in Columns 3-5 are

expanded list of crops whose yield drastically changed during the Green Revolution (Evenson & Gollin 2003).

[32] Even if manufacturing wage data were to exist they would not substitute for the cross-country analysis below. A Ricardo-Viner (specific factors) model with labor movement across sectors would predict a manufacturing wage increase in tandem with the agricultural wage and with the relative growth of the agricultural sector. The effect on national income, however, could still be negative for reasons outlined in the Introduction.

the fraction of villages with access to a primary school, middle school, and high school respectively. Most villages included a primary school at the start of the period so the coefficient in Column 3 is mechanically low; it is included for completeness. While π is positive in all cases, it is only statistically significant (p<0.01) when access to middle school is the outcome variable. In column 6, adult male literacy is the outcome variable. Rather than measuring access to schooling and school infrastructure, literacy measures education more directly. In both panels, the π is positive and significant.

Last, I investigate the impact of agricultural productivity growth on agricultural wages. Following Burgess et al. (2017), I approximate real agricultural wages by dividing the nominal daily wage measure by an agricultural price index.[33] Column 7 of Table 2 shows a positive and significant relationship between agricultural productivity and real agricultural wages. The results from this section thus suggest that agricultural productivity growth had a direct positive effect on rural wages and public goods.

4.2 Structural Change

4.2.1 Agricultural Sector

Table 3 explores the impact of the Green Revolution on the size of the farm sector. All regressions are estimates of (3). The first three outcome variables are (log of) individuals employed in agriculture (Columns 1-2), the fraction of the workforce employed in agriculture (Columns 3-4), and (log of) the total amount of land devoted to agriculture (Columns 5-6).[34] Even numbered columns include a postyear indicator interacted with: latitude, longitude, a coastline indicator, (log of) initial agricultural wages, initial population density, and the initial adult male literacy rate.

All columns point toward a positive and significant relationship between agricultural productivity growth and the size of the agricultural sector. Agricultural productivity growth associated with the Green Revolution increased district level employment in agriculture and land used for growing crops. The estimated magnitudes are also large: column 2 implies that a one standard deviation increase in agricultural productivity led to a 0.25 standard deviation increase in agricultural employment. I interpret these results as strong evidence that agricultural productivity growth had a significant and quantitatively meaningful positive impact on the size of districts' farm sectors.

Columns 7-8 ask a slightly different question and investigate the relationship between agricultural productivity growth and land ownership. The observation that the Green Revolution may have led to growing land inequality, especially in India, has been cited extensively to critique policies that promote technology driven agricultural development (Freebairn 1995). Cleaver Jr. (1972), for example, wrote in 1972 that "there is a growing effort by landlords to acquire more land and to convert their tenants into hired laborers in order to reduce their costs" (p. 182). Initially wealth-

[33]The agricultural price index is calculated at the district level is the weighted average price of the five major crops from the Sanghi et al. (1998) data—rice, wheat, maize, pearl millet, and sorghum—where the averaging weights are the district-level share of revenue for each crop.

[34]All regressions are robust to the inclusion of (log of) district population on the right hand side.

ier landowners may have also benefitted disproportionately by investing more heavily in Green Revolution technologies, exacerbating pre-existing inequality. The outcome variable in columns 7-8 is the fraction of agricultural laborers with land ownership rights.[35] π is negative and significant in both columns. While agricultural productivity growth increased overall resources in the agricultural sector, it led to a less even distribution of land ownership.

4.2.2 Non-Farm and Urban Sectors

Next, I turn to the impact of the Green Revolution's productivity growth on the non-farm sector. Table 4 estimates (3) with measures of employment in a series of non-farm occupations as the outcome variables.[36] In Column 1, the share of the total workforce employed in jobs that are not part of the farming sector is the outcome variable.[37] The coefficient of interest is negative and significant at below the 1% level. In the second column, the outcome variable is the (log of) the district-level urban population. Again the coefficient estimate is negative and significant.

The remaining columns focus on individual occupations: manufacturing, trade and commerce, service sector, transport and communication, and household industry.[38] The coefficient of interest is negative in all specifications. Intuitively, the effect is weakest for non-tradable sectors and strongest for manufacturing. Column 3 implies that a one standard deviation increase in agricultural productivity led to a 0.20 standard deviation relative decline in manufacturing employment. When the outcome variable is employment in either transportation and communication or "other" services, the coefficient of interest is not statistically different from zero. Overall, agricultural productivity growth had a discernible and negative effect on district-level urbanization and employment in manufacturing.

Taken together, these results demonstrate that in Indian districts most exposed to the Green Revolution, the agricultural sector expanded and the non-farm sector contracted, particularly in tradable industries.

[35] India's census distinguishes between "cultivators" and "agricultural laborers." While agricultural laborers worked either for wages or "in kind," cultivators had land ownership rights. The outcome variable is calculated as the number of cultivators divided by the total number of individuals employed in the farm sector. According to the census definition, "A person was considered as cultivator if he or she was engaged either as employer, single worker or family worker in cultivation of land owned or held from government or held from private persons or institutions." This is quoted from the India District's Database, http://vanneman.umd.edu/districts/codebook/defcult.html.

[36] Again here, all results are robust to the inclusion of (log of) district population is included as a control.

[37] This includes the sum of individuals employed in mining, manufacturing, construction, trade and commerce, transportation & communication, and other services.

[38] Here is an excerpt from the definition of Household Industry in the Indian Census: "Household industry was defined as an industry conducted by the head of the household himself/herself and/or by the members of the household at home or within the village in rural areas, and only within the precincts of the house where the household lived in urban areas ...A household industry is one that is engaged in production, processing, servicing, repairing or making and selling (but not merely selling) of goods." Crucially, Household Industry often encompassed manufacturing or service jobs that were run within a household or from the home. This description was accessed at the India Districts Database, http://vanneman.umd.edu/districts/codebook/defhhind.html.

4.3 Reduced Form Results

A methodological alternative to 2SLS estimation is estimation of the reduced form effect of predicted agricultural productivity. Indeed, in their analysis of structural change in Brazil, Bustos et al. (2016) interpret measures akin to single crop versions of the predicted productivity instrument as changes in potential yield and use these measures as their primary independent variables of interest. Reduced form estimates for the primary outcome variables in the preceding analysis are reported in Table A5 and are qualitatively very similar to results from 2SLS.

4.4 Mechanisms

4.4.1 Specialization Through Trade

If trade allowed districts that were most affected by the Green Revolution to specialize in agricultural production, the baseline effects should be most pronounced in districts that were most open to trade. These districts most completely specialized in agriculture following the shift in within-country comparative advantage. To my knowledge, a direct measure of district-level openness to trade is not available at the start of the sample period. To overcome this, I construct a proxy from the census data, which is the share of the workforce employed in trade and commerce. In Table 5, I present reduced form estimates where, along with the instrument P_{it}, I include an interaction term between the instrument and the share of the workforce employed in trade and commerce in 1961. I interpret the interaction term as the differential effect of agricultural productivity growth for districts that were more exposed to trade.

The baseline impacts on employment and sector specialization increase in magnitude with the trade and commerce employment share. One possible concern with the trade and commerce employment share measure is that is might be simply a proxy for overall non-farm employment at the start of the period. To address this, in odd numbered columns I also include an interaction term between the instrument and the non-farm, non-trade share of employment in 1961. The interaction with the trade share remains similar across specifications and the interaction with the non-trade share is consistently small in magnitude. This suggests that the baseline heterogeneous effects are not driven by differences in total non-farm employment.

The results in Table 5 suggest that, as predicted by a Matsuyama (1992)-style model, the negative relationship between agricultural productivity growth and structural change is stronger in more open districts, which more completely specialize in agricultural production. The impact on farm employment (columns 1-2), urbanization (columns 5-6) and manufacturing employment (columns 7-8) are all substantially larger in more open districts; I do not find a differential effect when agricultural land is the outcome. When log of the urban population is the outcome variable, the the direct effect is positive and significant; column 6 implies that in the 25% of districts least involved in trade, agricultural productivity growth had a positive effect on urbanization. Thus, even across districts, exposure to trade mediates relationship between agricultural productivity growth and structural change.

4.4.2 Migration

The impact of changes in agricultural productivity on the composition of the labor force may be driven by either the movement of individual between sectors within a district or by movement of individuals between districts (or larger regions). Table A6 partially distinguishes between these possibilities by controlling directly for net migration from other districts in the regressions presented in Tables 3 and 4. In the first column, (log of) farm labor is the outcome variable and the coefficient π remains highly significant ($p < 0.01$) and the point estimate is very similar to Table 5. The impact on farm labor is not affected by migration. In column 2, the outcome variable is log of manufacturing labor, and the estimated effect of agricultural productivity growth is attenuated compared to Table 4. In columns 3-7, the outcome variables are employment in the remaining non-farm occupations, and the estimated effect of agricultural productivity is not statistically different from zero. In all cases, the coefficient on (log of) inward migration is positive and highly significant.

Column 9 explicitly documents a negative relationship between agricultural productivity growth and inward migration from other Indian districts. The contraction of non-farm sectors in places where agricultural productivity growth was the greatest seems to have been driven in part by the movement of labor across districts; in particular, there was lower net migration into districts that experienced a large agricultural productivity growth. the increase in agricultural labor, on the other hand, was not driven by migration. Unfortunately to my knowledge it is not possible to determine whether this is driven by fewer individuals migrating into districts with greater increases in agricultural productivity or more individuals migrating out of these districts. Nevertheless, the fact that local productivity growth affects migration across districts is further motivation for investigating the impact of agricultural productivity growth at the country-level, where the reallocation of labor across sub-national regions is taken into account.

5 Cross-Country Results

It is unclear ex ante how subnational trends presented in Section 4 might aggregate to the country-level. In this section, I use an analogous identification strategy to the previous section to study the impact of the Green Revolution across countries. Here, it is possible to test broader implications of the Green Revolution at the national level, including its effect on national income. OLS results are presented in Table A9; however, for standard reasons, these estimates are unreliable. For example, the fact that there is a highly significant positive correlation between agricultural productivity and per-capita GDP (Column 7) may result from the fact that wealthy countries are able to invest more in agricultural production. The remainder of the section therefore focuses exclusively on causal estimates from the instrumental variables approach.

5.1 First Stage Estimates

5.1.1 Baseline Estimates and Robustness

The first stage at the country level along with a series of robustness checks—analogous to Tables 1 and A2 for the sub-national analysis—are reported in Tables 6 and A7. The country-level first stage is also calculated using (1) and the first stage is an estimate of (2). The only difference is that (1) is calculated using a broader set of crops that were impacted during the Green Revolution due to the greater variation ecological and climatic characteristics in the global analysis compared to the subnational analysis.[39] In column 4 of Table A7 I show that the first stage relationship is similar, ableit weaker, if I construct the instrument using potential yield data on only rice, wheat, and maize, as in the sub-national analysis. I find a strong first stage relationship that does not appear to be driven by a pre-existing trend in agricultural productivity. The first stage coefficient estimate is also very similar to the corresponding first stage coefficient estimates from the sub-national analysis (Table 1, columns 5-6).

5.1.2 Placebo Test

If the first stage estimates are capturing the adoption of HYVs in places that stood to benefit from adoption, there should not be a first stage relationship between actual agricultural productivity and placebo instruments constructed using crops that were not actually affected by the Green Revolution. This idea motivates the placebo exercise presented in Table A8. I construct a series of placebo instruments using pairs of crops for which major HYVs were not released during the Green Revolution and a series of real instruments using crops for which major HYVs were released. First stage estimates using both real and placebo instruments are reported in Panels A and B respectively of Table A8. None of the coefficients on the placebo instruments attains statistical significance while all instruments constructed from pairs of Green Revolution crops are significantly correlated with actual productivity. In some cases, the coefficient on the placebo instrument is even negative (columns 2, 5, and 7).[40] These results suggest that the instrument is indeed capturing variation in productivity due to variation in HYV suitability and adoption.

5.2 Structural Change at the Country Level

Table 7 presents the main results from the country-level analysis. Each column reports an estimate of (3) in which i is a country and the coefficient of interest if π. All regressions use a seven-year panel, including one observation every five years for the years 1960-1990, and include country and year fixed effects; results using the full panel are similar.

[39] The full list of crops is: Rice (wetland and drylands), cassava, potato, phaseolus bean, maize, sorghum, wheat, and barley. These are all crops for consumption in Evenson & Gollin (2003) whose HYVs were released between 1960-1990, the time period for which the FAO data were calibrated.

[40] Even an instrument constructed using oats and rye (columns 1-3)—perhaps the crops that are most similar to rice and wheat but which were not nearly as extensively researched during the Green Revolution—is not significantly correlated with actual productivity, and the sign of the correlation is negative controlling for initial income.

In the first column, the fraction of land devoted to agriculture is the outcome variable. π is positive and significant, closely mirroring the subnational results. In the second column, I run the same regression excluding the wealthiest 10% of countries from the sample. This is motivated by the fact that the Green Revolution represented an explicit effort to transform agricultural production in lower income countries; countries at the top of the income distribution had less to gain from the Green Revolution's HYVs and overall growth dynamics were also likely less susceptible to changes in the agricultural sector. π increases slightly in magnitude. The same specifications are repeated in columns 3-4 with the agricultural employment share as the outcome variable, and π is positive and weakly significant in both cases. Columns 5-6 report estimates of the same two specifications with urbanization as the outcome variable—these results also mirror the results from within India. π is negative and (weakly) significant in Column 5 and both higher in magnitude and statistical significance in Column 6.[41] Across countries, agricultural productivity growth increased the relative size of the agricultural sector, in terms of both land and labor; urbanization was negatively affected.

Next, I turn to the impact of agricultural productivity growth on measures of national income. Column 7 documents a negative and (weakly) significant relationship between agricultural productivity growth and total GDP. In Column 8, when per capita GDP is the outcome variable π is negative but insignificant. Economically meaningful positive values occupy a small fraction of two standard error or Anderson-Rubin confidence intervals around the 2SLS estimates. I find no evidence of a positive baseline relationship between agricultural productivity growth and income. Reduced form estimates of all baseline cross-country results, which tell a similar story as the IV estimates, are reported in Table A10.

5.2.1 Mechanisms and Heterogeneity

The clearest explanation for these results is that countries that were more exposed to the Green Revolution developed or strengthened a comparative advantage in agricultural production. Combined with "learning by doing" or other dynamic externalities in urban and manufacturing sectors, this could generate a negative relationship between agricultural productivity growth and GDP. To investigate this mechanism, I use data on cross-country trade in agricultural goods; due to data availability, the sample size is slightly reduced. I find that agricultural productivity growth led

[41] In some cases, reducing urbanization might not be an entirely negative outcome. In a model of population movement between the rural and urban sectors, Todaro (1969) finds that "the net benefit of bringing 'city lights' to the countryside might greatly exceed whatever net benefit might be derived from luring more peasants to the city by increasing the attractiveness of urban living conditions...[or else] the lure of relatively higher permanent incomes will continue to attract a steady stream of rural migrants into the ever more congested urban slums. The potential social, political, and economic ramifications of this growing mass of urban unemployed should not be taken lightly" (p. 147). Bates (2005) also argues that in Africa, politics that favor industrial development over rural livelihood and promote policy designed to reduce prices in cities can have detrimental economic ramifications. However in the long run, the analysis of India showed that agricultural productivity growth did not seem to improve quality of life in the rural sector—indeed, rural wages were relatively lower in Indian districts that experienced greater agricultural productivity growth. Yet Bates's argument against the staunchly pro-urban policies of several African governments during the 20th century and the fact that that Africa is also the region that seems to have been least successful at adopting Green Revolution technologies is an interesting connection that deserves further exploration.

to a decline in agricultural imports ($p<0.01$) and a (statistically insignificant) increase in agricultural exports (columns 9-10). This is consistent with agricultural productivity growth leading to international specialization through trade and the Green Revolution increasing certain countries' comparative advantage in agriculture.

To further understand the mechanism, I group countries based on their openness to trade in 1960, using trade as a share of GDP as a proxy for openness. If the negative impact of agricultural productivity growth on structural transformation is driven by specialization in agriculture and foregone externalities from urbanization, the baseline effects should be more pronounced for more open countries that more completely specialized in agricultural production. From a theoretical perspective, these countries—perhaps like districts in India—more closely resemble the small open economies in Matsuyama (1992). I present reduced form estimates where both the instrument and an interaction term between the instrument and the country's pre-period trade volume as a share of GDP are included on the right hand side of the regression. Results from this heterogeneity analysis are presented in Table 8. In even columns, openness to trade is captured with an indicator that equals one if a country's 1960 trade as a share of GDP was above the median; in odd columns, I separate countries by their openness quartile. In order to isolate the impact of trade, I also control for pre-period (log of) population and GDP interacted with year indicators in order to control flexibly for country size and wealth, which may be correlated with openness.

While the lower sample size due to the limitation of the trade data (72 countries) reduces the statistical precision of the estimates, Table 8 tells a consistent story that the estimates from Table 6 are more pronounced in initially more open countries. For example, column 1 implies a that there is, if anything, a small positive relationship between agricultural productivity and urbanization for countries with below-median trade volume in 1960, but that there is a large negative effect for countries with above-median trade volume in 1960. Similarly, the positive effect of agricultural productivity growth on agricultural employment is increasing with countries' pre-period trade volume (columns 3-4). Column 6 implies that for countries in the top quartile of pre-period openness, there is a large and significant positive effect of agricultural productivity growth on agricultural exports.

Columns 5-8 turn to the heterogeneous impact on total and per-capita GDP. Again, I find that the impact is larger for initially more open countries. Intuitively, I find that agricultural productivity growth had a positive effect on income in countries in the bottom openness quartile. In these countries, the Green Revolution had a positive direct effect on agricultural productivity, and no countervailing negative impact on urban development (columns 1-2); if anything, the opposite. I also find that agricultural productivity growth had a negative and significant impact on income in countries that were above median openness in 1960. This is consistent with the large negative effect on urbanization and growth in size of the agricultural sector for these countries.

Finally, in columns 9-12, I estimate the heterogeneous impact of agricultural productivity growth on agricultural trade. Intuitively, productivity growth had the largest negative effect on agricultural imports and positive effect on agricultural exports in countries that were more open to trade.

This provides validation that the pre-period openness measure indeed captures the countries that were more likely to specialize through trade in agricultural production.

6 Conclusion

This paper develops an identification strategy to estimate the causal effect of agricultural productivity growth associated with the Green Revolution, a recent period of unprecedented agricultural productivity growth. I calculate a measure of predicted agricultural productivity using time-invariant geographic and ecological characteristics, along with the timing of global HYV releases. Using predicted productivity as an instrument, this study examines the impact of the Green Revolution on structural change, both across districts in India and across countries.

The effect of agricultural productivity growth at the national and sub-national level could in theory be very different (Matsuyama, 1992, Murphy, Shleifer & Vishny, 1989). However, both at the sub-national and national level, I find that agricultural productivity growth led to growth of the agricultural sector at the expense of non-farm development. Within India, it led to higher agricultural wages and greater employment and land use in agriculture; at the same time, it reduced urbanization and manufacturing employment, particularly in districts more involved in trade. Furthermore, the Green Revolution lowered country-level urbanization and increased the size of the agricultural sector. It had, if anything, a negative effect on income, driven by countries that were more open to trade in 1960.

More broadly, these findings show that equilibrium effects play a major role in shaping the impact of agricultural productivity gains. Especially as some policy makers call for a Second Green Revolution in agriculture, one that develops more ecologically sustainable processes to mass produce food, a deep knowledge recent changes to the agricultural sector and their impacts is critical.

7 Bibliography

Acemoglu, Daron, *Introduction to Modern Economic Growth* (Princeton: Princeton University Press, 2009).

Alesina, Alberto, Paola Giuliano, and Nathan Nunn, "On the Origins of Gender Roles: Women and the Plough," The Quarterly Journal of Economics 128, no. 2 (2013), 469-530.

Auty, Richard, *Sustaining Development in Mineral Economies: The Resource Curse Thesis* (London/New York: Routledge, 1993).

Ball, V. Eldon, et al., "Agricultural Productivity Revisited," American Journal of Agricultural Economics 79, no. 4 (1997), 1045-63.

Banerjee, Abhijit, and Lakshmi Iyer, "History, Institutions, and Economic Performance: The Legacy of Colonial Land Tenure Systems in India," The American Economic Review 95, no. 4 (2005), 1190-1213.

Banerjee, Abhijit V., and Andrew F. Newman, "Information, the Dual Economy, and Development," The Review of Economic Studies 65, no. 4 (1998), 631-53.

Bartelme, Dominick, Arnaud Costinot, Andres Rodriguez-Clare, and Dave Donaldson. "Economies of Scale and Industrial Policy: A View from Trade," Working Paper, 2018.

Bates, Robert H., *Markets and States in Tropical Africa: The Political Basis of Agricultural Policies: With a New Preface.* 2 edition (Berkeley: University of California Press, 2005).

Brass, Paul R., "The Politicization of the Peasantry in a North Indian State: I,' The Journal of Peasant Studies 7, no. 4 (1980), 395-426.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak, "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh," Econometrica 82, no. 5 (2014), 1671-1748.

Burgess, Robin, Olivier Deschenes, Dave Donaldson, and Michael Greenstone, "Weather, Climate Change and Death in India," https://www.michaelgreenstone.com/s/WD_master_170420-1.pdf, 2017.

Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli, "Agricultural Productivity and Structural Transformation: Evidence from Brazil," American Economic Review 106, no. 6 (2016), 1320-65.

Cleaver, Harry M., "The Contradictions of the Green Revolution," The American Economic Review 62, no. 1 (1972), 177-86.

Dalrymple, Dana G., "Development and Spread of High-Yielding Wheat Varieties in Developing Countries," Bureau for Science and Technology, Agency for International Development, 1982.

Dasgupta, Aditya, "The Curse of Strong Roots in Dominant Party Collapse: Evidence from India's Green Revolution," SSRN Scholarly Paper. https://papers.ssrn.com/abstract=2303336, 2014.

Duflo, Esther, and Rohini Pande, "Dams," The Quarterly Journal of Economics 122, no. 2 (2007), 601-646.

Eckert, Fabian, and Michael Peters. "Spatial Structural Change." 2018 Meeting Papers. Society for Economic Dynamics, 2018.

Evenson, Robert E. and Douglas Gollin, "Assessing the Impact of the Green Revolution, 1960 to

2000," Science 300, no. 5620 (2003), 758-62.

Evenson, Robert E. and Douglas Gollin [2], *Crop Variety Improvement and Its Effect on Productivity: The Impact of International Agricultural Research*, CABI, 2003.

Field, Alexander, "Sectoral Shift in Antebellum Massachusetts: A Reconsideration," Explorations in Economic History 15, no. 2 (1978), 146-71.

Fiszbein, Martin, "Agricultural Diversity, Structural Change and Long-Run Development: Evidence from US Counties," Unpublished Manuscript, http://www.economics.illinois.edu/ seminars/documents/ Fiszbein.pdf. 2015.

Foster, Andrew D. and Mark R. Rosenzweig, "Agricultural Productivity Growth, Rural Economic Diversity, and Economic Reforms: India, 1970-2000," Economic Development and Cultural Change 52, no. 3 (2004), 509-542.

Foster, Andrew D., and Mark R. Rosenzweig, "Technical Change and Human-Capital Returns and Investments: Evidence from the Green Revolution," The American Economic Review 86, no. 4 (1996), 931-953.

Frankel, Francine R., *India's Green Revolution: Economic Gains and Political Costs* (Princeton: Princeton University Press, 1971).

Goldberg, Jessica, "Kwacha Gonna Do? Experimental Evidence about Labor Supply in Rural Malawi," American Economic Journal: Applied Economics 8, no. 1 (2016), 129-49.

Gollin, Douglas, "Agricultural Productivity and Economic Growth," In Handbook of Agricultural Economics, edited by Prabhu Pingali and Robert Evenson, 4 (2010), 3825-66.

Gollin, Douglas, Casper Worm Hansen, and Asger Wingender. "Two Blades of Grass: The Impact of the Green Revolution." Working Paper. National Bureau of Economic Research, June 2018.

Gollin, Douglas, David Lagakos, and Michael E. Waugh, "The Agricultural Productivity Gap," The Quarterly Journal of Economics 129, no. 2 (2014), 939-993.

Griffin, Keith, *The Political Economy of Agrarian Change: An Essay on the Green Revolution* (London: Macmillan Press, 1979).

Harriss, John, "The Limitations of HYV Technology in North Arcot District: The View from a Village," In Green Revolution, edited by B. H. Farmer. Palgrave Macmillan (1980), 124-142.

Hazell, Peter and Lawrence James Haddad, *Agricultural Research and Poverty Reduction*, International Food Policy Research Institute, 2001.

Hornbeck, Richard and Pinar Keskin, "Does Agriculture Generate Local Economic Spillovers? Short-Run and Long-Run Evidence from the Ogallala Aquifer," American Economic Journal: Economic Policy 7, no. 2 (2015), 192-213.

Hossain, Mahabub, *Nature and Impact of the Green Revolution in Bangladesh*, International Food Policy Research Institute, 1988.

Hsieh, Chang-Tai, and Benjamin A. Olken, "The Missing ?Missing Middle," The Journal of Economic Perspectives 28, no. 3 (2014), 89-108.

Imbert, Clment, and John Papp, "Short-Term Migration Rural Workfare Programs and Urban Labor Markets-Evidence from India," http://wrap.warwick.ac.uk/id/eprint/82307, 2016.

Irz, Xavier, Lin Lin, Colin Thirtle, and Steve Wiggins, "Agricultural Productivity Growth and

Poverty Alleviation," Development Policy Review 19, no. 4 (2001), 449-66.

Iyer, Lakshmi, "Direct versus Indirect Colonial Rule in India: Long-Term Consequences," The Review of Economics and Statistics 92, no. 4 (2010), 693-713.

Jayachandran, Seema, "Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries," Journal of Political Economy 114, no. 3 (2006), 538-75.

John, A. H., "Agricultural Productivity and Economic Growth in England, 1700-1760," The Journal of Economic History 25, no. 1 (1965), 19-34.

Khush, Gurdev S., "Green Revolution: The Way Forward," Nature Reviews Genetics 2, no. 10 (2001), 815-22.

Kogel, Tomas, and Alexia Prskawetz, "Agricultural Productivity Growth and Escape from the Malthusian Trap," Journal of Economic Growth 6, no. 4 (2001), 337-57.

Ladejinsky, Wolf, "Ironies of India's Green Revolution," *Foreign Affairs*, January 28, 1970, https://www.foreignaffairs.com/articles/asia/1970-07-01/ironies-indias-green-revolution.

Lal, Rattan, David O. Hansen, and Norman Uphoff, *Food Security and Environmental Quality in the Developing World*, CRC Press, 2016.

Lall, Somik Vinay, and Sanjoy Chakravorty, "Industrial Location and Spatial Inequality: Theory and Evidence from India," Review of Development Economics 9, no. 1 (2005), 47-68.

Lall, Somik V., and Taye Mengistae, "The Impact of Business Environment and Economic Geography on Plant-Level Productivity?: An Analysis of Indian Industry," Policy Research Working Paper Series, The World Bank, 2005.

Levitt, Steven D., John A. List, and Chad Syverson. "Toward an Understanding of Learning by Doing: Evidence from an Automobile Assembly Plant." Journal of Political Economy 121, no. 4 (2013): 643-81.

Lashkaripour, Ahmad, and Volodymyr Lugovskyy. "Scale Economies Andthe Structure of Trade and Industrial Policy," Working Paper, 2019.

Lewis, W. Arthur, *Theory of Economic Growth* (Crows Nest: Allen & Unwin, 1955).

Matsuyama, Kiminori, "Increasing Returns, Industrialization, and Indeterminacy of Equilibrium," The Quarterly Journal of Economics 106, no. 2 (1991), 617-650.

Matsuyama, Kiminori, "Agricultural Productivity, Comparative Advantage, and Economic Growth," Journal of Economic Theory 58, no. 2 (1992), 317-34.

McArthur, John W., and Gordon C. McCord. "Fertilizing Growth: Agricultural Inputs and Their Effects in Economic Development." Journal of Development Economics 127 (July 1, 2017): 133-52.

Munshi, Kaivan, and Mark Rosenzweig, "Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap," American Economic Review 106, no. 1 (2016), 46-98.

Murphy, Kevin M., Andrei Shleifer, and Robert Vishny, "Income Distribution, Market Size, and Industrialization," The Quarterly Journal of Economics 104, no. 3, 1989: 537-64.

Niazi, Tarique, "Rural Poverty and the Green Revolution: The Lessons from Pakistan," The Journal of Peasant Studies 31 no. 2 (2004), 242-60.

Nunn, Nathan and Nancy Qian, "The Potato's Contribution to Population and Urbanization: Evidence From A Historical Experiment," The Quarterly Journal of Economics 126, no. 2 (2011), 593-650.

Patibandla, Murali, and Bent Petersen, "Role of Transnational Corporations in the Evolution of a High-Tech Industry: The Case of India's Software Industry," World Development 30, no. 9 (2002), 1561-77.

Pingali, Prabhu L., "Green Revolution: Impacts, Limits, and the Path Ahead," Proceedings of the National Academy of Sciences 109, no. 31 (2012), 12302-12308.

Qian, Nancy, "Missing Women and the Price of Tea in China: The Effect of Sex-Specific Earnings on Sex Imbalance," The Quarterly Journal of Economics 123, no. 3 (2008), 1251-85.

Rosenberg, Nathan. *Inside the Black Box: Technology and Economics*. Cambridge: Cambridge University Press, 1983.

Rostow, Walt Whitman, *The Stages of Economic Growth: A Non-Communist Manifesto* (Cambridge: Cambridge University Press, 1990).

Rudolph, Susanne Hoeber, and Lloyd I. Rudolph, "The Centrist Future of Indian Politics," Asian Survey 20, no. 6 (1980), 575-94.

Sachs, Jeffrey D., and Andrew M. Warner, "Natural Resource Abundance and Economic Growth," Working Paper. National Bureau of Economic Research, 2005.

Sanghi, Apurva, K. S. Kavi Kumar, and James W. McKinsey Jr., "India Agriculture and Climate Data Set," https://ipl.econ.duke.edu/dthomas/dev_data/datafiles/india_agric_climate.htm, 1998.

Sasaki, A., M. Ashikari, et al., "Green Revolution: A Mutant Gibberellin-Synthesis Gene in Rice," Nature 416, no. 6882 (2002), 701-2.

Shiva, Vandana, *The Violence of the Green Revolution: Third World Agriculture, Ecology, and Politics* (London: Zed Books, 1991).

Sonnenfeld, David A., "Mexico's 'Green Revolution,' 1940-1980: Towards an Environmental History," Environmental History Review 16, no. 4 (1992), 29-52.

Taraz, Vis, "Adaptation to Climate Change: Historical Evidence from the Indian Monsoon,"Environment and Development Economics, 22(5), 517-545.

Thirtle, Colin, Lin Lin, and Jenifer Piesse, "The Impact of Research-Led Agricultural Productivity Growth on Poverty Reduction in Africa, Asia and Latin America," World Development 31, no. 12 (2003), 1959-75.

Thompson, Peter. "Learning by Doing." Handbook of the Economics of Innovation. Elsevier, 2010.

Todaro, Michael P., "A Model of Labor Migration and Urban Unemployment in Less Developed Countries," The American Economic Review 59, no. 1 (1969), 138-48.

Topalova, Petia. "Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India." American Economic Journal: Applied Economics 2, no. 4 (October 2010): 1-41.

Wingender, Asger Moll, "Structural Transformation in the 20th Century: A New Database on Agricultural Employment Around the World," https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2533360, 2014.

Young, Alwyn. "Learning by Doing and the Dynamic Effects of International Trade." The Quarterly Journal of Economics 106, no. 2 (May 1, 1991): 369-405.

Young, Alwyn, “Inequality, the Urban-Rural Gap, and Migration,” The Quarterly Journal of Economics 128, no. 4 (2013), 1727-85.

Zanden, J. L. Van, “The First Green Revolution: The Growth of Production and Productivity in European Agriculture, 1870-1914.” The Economic History Review 44, no. 2 (1991), 215-239.

8 Figures & Tables

Figure 1: HYV Adoption in Indian Districts: Wheat, Rice & Maize

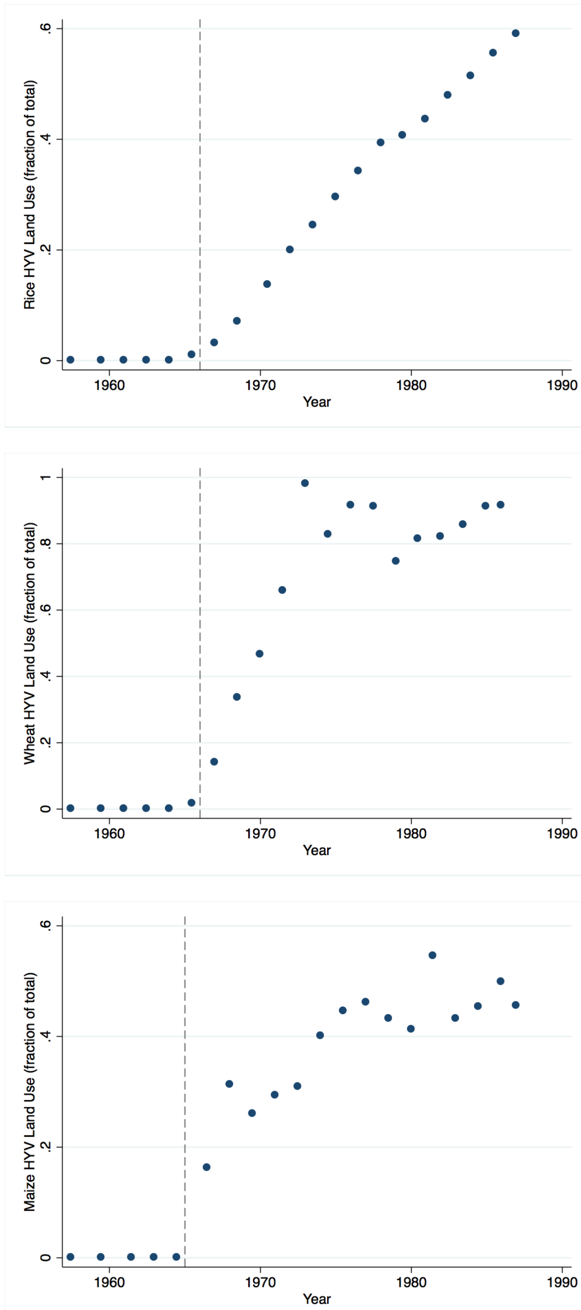

Notes: The year is on the X-axis and the fraction of crop land devoted to cultivating indicated crop in which high yield varieties were used is on the Y-axis. The dotted vertical line is the year in which the high yield variety for the graph's crop was released—that is, the year that I chose to construct the instrument—and values are averaged over the 271 Indian districts in the Sanghi et al. (1998) data set.

Figure 2: Change in District-Level Predicted and Actual Agricultural Productivity

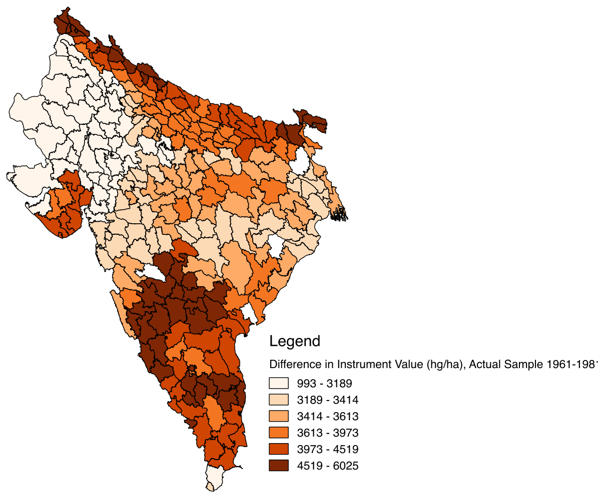

(a) Δ Predicted Agricultural Productivity, 1961-1981

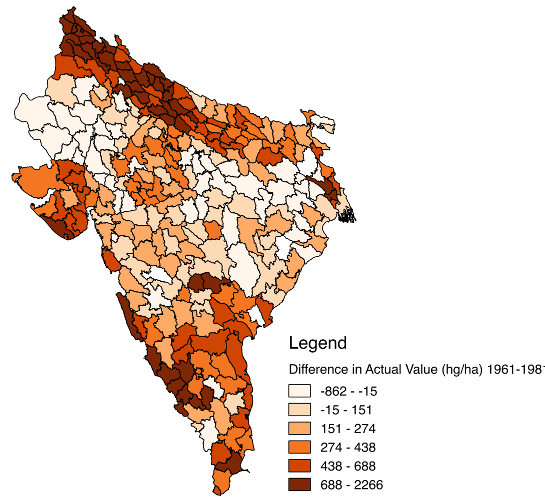

(b) Δ Actual Agricultural Productivity, 1961-1981

Notes: Change in the instrument (predicted agricultural productivity) and actual agricultural productivity for the 271 Indian districts in the sample.

Figure 3: First Stage & Pre-Trend Falsification

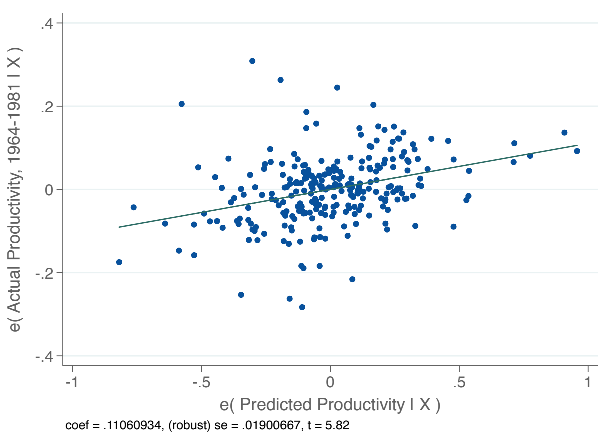

(a) Δ Predicted vs. Δ Actual Productivity, 1964-1981

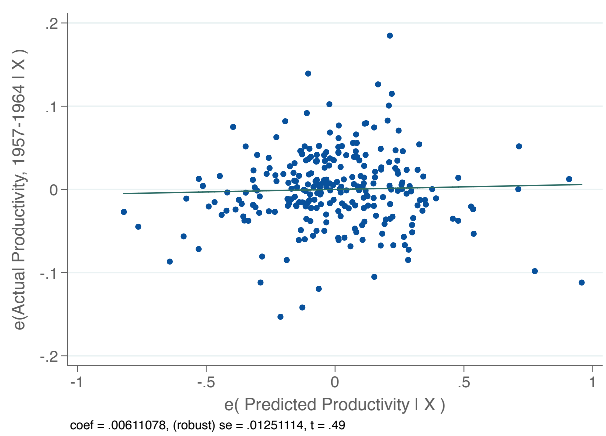

(b) Δ Predicted vs. Δ Actual Productivity, 1957-1964

Notes: Partial correlation plots illustrating the relationship between predicted and actual agricultural productivity. Estimates are conditional on latitude, longitude, initial log agricultural wages, initial population density, and the initial male literacy rate.

Table 1: Baseline District Level First Stage

	(1)	(2)	(3)	(4)
Outcome Variables:	Agricultural Productivity (kcal/ha)			
Predicted Agricultural Productivity	0.141***	0.162***	0.138***	0.121***
	(0.0350)	(0.0307)	(0.0318)	(0.0342)
Latitude x I_{1981}		Yes	Yes	Yes
Longitude x I_{1981}		Yes	Yes	Yes
Coastline Indicator x I_{1981}		Yes	Yes	Yes
1961 log Ag. Wages x I_{1981}		Yes	Yes	Yes
1961 Population Density x I_{1981}			Yes	Yes
1961 Literacy Rate x I_{1981}			Yes	Yes
1961 Agricultural Productivity x I_{1981}				Yes
District & Year Fixed Effects	Yes	Yes	Yes	Yes
F-Statistic	16.33	27.66	18.93	12.47
Districts	271	271	271	271
Observations	542	542	542	542
R-squared	0.851	0.881	0.893	0.895

Notes: Each regression uses two-year panel data of 271 Indian districts and the years 1961 and 1981. The independent variable of interest is predicted agricultural productivity. District and year fixed effects are included in all specifications. The F-statistic for the instrument, predicted agricultural productivity, is reported at the bottom of the table. Robust standard errors clustered at the district level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table 2: Rural Public Goods and Wages

Outcome Variable:	(1) Road 2SLS	(2) Medical Center 2SLS	(3) Primary School 2SLS	(4) Middle School 2SLS	(5) High School 2SLS	(6) Literacy 2SLS	(7) log Ag. Wage 2SLS
Ag. Productivity (kcal/ha) x 10^{-6}	2.140***	0.815**	0.337	0.414***	0.0175	0.185**	3.996***
	(0.787)	(0.356)	(0.363)	(0.118)	(0.0609)	(0.0808)	(0.575)
F-Statlstic	9.34	10.53	23.21	25.89	22.39	27.66	27.66
R-squared	0.792	0.762	0.809	0.886	0.884	0.936	0.774
Latitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Longitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coastline Indicator x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Districts	196	132	202	169	203	271	271
Observations	392	264	404	338	406	542	542

Notes: Each regression is a 2SLS model uses a two-year panel data of 271 Indian districts and the years 1961 and 1981. The independent variable is agricultural productivity measured in kilocalories per hectare. The outcome variables are the fraction of villages in each district with access to roads (Column 1), a medical center (Column 2), a primary school (Column 3), a middle school (Column 4), or a high school (Column 5), the adult male literacy rate (Column 6), and (log of) daily agricultural wages normalized by an agricultural price index. All of these measures are from India's national census. The F-statistic for the instrument, predicted agricultural productivity, is reported at the bottom of the table. District and year fixed effects are included in all specifications. Robust standard errors clustered at the district level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table 3: Farm Labor and Land

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Outcome Variable:	log Farm Labor		Farm Share of Employment		log Farm Land		Frac. Farmers Own Land	
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
Ag. Productivity (kcal/ha) x 10^{-6}	1.045***	1.168***	0.305***	0.371***	0.378**	0.401**	-0.315***	-0.243**
	(0.256)	(0.315)	(0.0915)	(0.120)	(0.150)	(0.188)	(0.0964)	(0.115)
F-Statlstic	28.49	19.35	27.66	18.93	28.49	19.35	28.49	19.35
R-squared	0.981	0.981	0.949	0.948	0.992	0.992	0.961	0.964
District & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Latitude x I_{1981} & Longitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coastline Indicator x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
1961 log Ag. Wages x I_{1981}		Yes		Yes		Yes		Yes
1961 Population Density x I_{1981}		Yes		Yes		Yes		Yes
1961 Literacy Rate x I_{1981}		Yes		Yes		Yes		Yes
Districts	271	271	271	271	271	271	271	271
Observations	542	542	542	542	542	542	542	542

Notes: Each regression is a 2SLS model that uses two-year panel data of 271 Indian districts and the years 1961 and 1981. The independent variable is agricultural productivity measured in kilocalories per hectare.The outcome variables are the (log of) the number of farm laborers in each district (Columns 1-2), the fraction of the total workforce that is devoted to farm labor (Columns 3-4), and the (log of) district land devoted to agriculture. In Columns 1-2 & 3-4, log population is included on the right hand side. The F-statistic for the instrument, predicted agricultural productivity, is reported at the bottom of the table. District and year fixed effects are included in all specifications. Robust standard errors clustered at the district level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table 4: Non-Farm Labor

Outcome Variables:	(1) Urban Population 2SLS	(2) Non-Farm Share of 2SLS	(3) Manufac. 2SLS	(4) Commerce 2SLS	(5) "Other" Service 2SLS	(6) Transport. & Commun. 2SLS	(7) Household Industry 2SLS
Ag. Productivity (kcal/ha) x 10^{-6}	-0.653**	-0.215***	-1.915***	-0.678**	-0.477	-0.402	-0.302*
	(0.293)	(0.0802)	(0.497)	(0.312)	(0.294)	(0.418)	(0.170)
F-Statlstic	27.41	27.65	27.65	27.65	27.65	27.65	27.65
R-squared	0.985	0.961	0.974	0.986	0.981	0.977	0.991
Latitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Longitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coastline Indicator x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Districts	270	271	271	271	271	271	271
Observations	540	542	542	542	542	542	542

Notes: Each regression is a 2SLS model that uses two-year panel data of 271 Indian districts and the years 1961 and 1981. The independent variable is agricultural productivity measured in kilocalories per hectare. The outcome variables are the share of non-farm employment, and the (log of) the urban population, number of people employed in manufacturing, commerce, "other" service sectors, transport & communication, and household industry (Columns 2-7). In Columns 2-7, log of district population population is included on the right hand side. The F-statistic for the instrument, predicted agricultural productivity, is reported at the bottom of the table. District and year fixed effects are included in all specifications. Robust standard errors clustered at the district level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table 5: District-Level Heterogeneity: Involvement in Trade

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	log Farm Labor		log Farm Land		log Urban Population		log Manuf. Employment	
Predicted Productivity	-0.0328	-0.0292	0.0818*	0.0795*	0.138**	0.133**	0.0982	0.0902
	(0.0475)	(0.0476)	(0.0439)	(0.0444)	(0.0601)	(0.0606)	(0.0968)	(0.0968)
Predicted Productivity x 1961 Commerce Share	0.0277***	0.0417***	-0.00536	-0.0119	-0.0340***	-0.0501***	-0.0567***	-0.0797***
	(0.00434)	(0.0120)	(0.00441)	(0.00724)	(0.00562)	(0.00829)	(0.00948)	(0.0170)
Predicted Productivity x 1961 Non-Farm Non-Comm. Share		-0.00319		0.00150		0.00370**		0.00528
		(0.00292)		(0.00131)		(0.00145)		(0.00321)
District & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Latitude x I_{1981} & Longitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coastline Indicator x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
1961 log Ag. Wages x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
1961 Population Density x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
1961 Literacy Rate x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	542	542	542	542	540	540	542	542
R-squared	0.983	0.983	0.982	0.982	0.986	0.986	0.976	0.976

Notes: Each column reports a reduced form estimate from a two-year panel of 271 Indian districts and the years 1961 and 1981. The independent variable is predicted agricultural productivity measured in kilocalories per hectare, and an interaction term between predicted productivity and the share of district-level labor employed in commerce in 1961. The outcome variables are listed at the top of each column. All columns include an interaction between predicted productivity. District and year fixed effects are included in all specifications, along with the full set of geographic controls.. Robust standard errors clustered at the district level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table 6: Cross-Country Analysis: First Stage

	(1)	(2)	(3)	(4)	(5)	(6)
		Outcome is Agricultural Productivity (kcal/ha)				
Sample:	Full Panel	Reduced Panel (every 5 years)			Excluding top income	Excluding bottom income
Predicted Agricultural Productivity	0.210***	0.159***	0.159***	0.158***	0.167***	0.186***
	(0.0637)	(0.0499)	(0.0501)	(0.0509)	(0.0553)	(0.0579)
Country & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
1960 log p.c. GDP x $\sum_t I_t$		Yes	Yes	Yes	Yes	Yes
1960 urbanization x $\sum_t I_t$			Yes	Yes		
1960 agricultural productivity x ∑tIt				Yes		
F-statistic	10.87	10.16	10.14	9.63	9.08	10.32
Observations	4,561	798	791	791	602	602
R-squared	0.959	0.967	0.968	0.977	0.968	0.967

Notes: Column 1 uses panel of 115 countries and all years 1961-1990. Columns 2-6 use a 7-year panel (years included are 1961, 1965, 1970, 1975, 1980, 1985, and 1990). Agricultural productivity measured in kilocalories per hectare is the outcome variable in all columns. In Column 5 countries that were in the highest income quartile in 1960 are removed from the sample and in column 6 countries that were in the lowest income quartile in 1960 are removed from the sample. The independent variable of interest is the predicted agricultural productivity instrument. Country and year fixed effects are included in all specifications, along with the characteristics listed at the bottom of each panel (log of per capita GDP in 1960, urbanization rate in 1960, and actual agricultural productivity in 1961) interacted with a full set of time dummies. Robust standard errors clustered at the country level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table 7: Cross-Country Analysis: Main Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Outcome Variable:	Frac. Land Devoted to Ag.	Land Devoted to Ag., Excl. top 10% p.c. GDP	Agricultural Employment Share (AES)	AES, Excl. top 10% p.c. GDP	Urbanization	Urbanization, Excluding top 10% p.c. GDP	log GDP	log p.c. GDP	Ag. Imports Share of 1961 GDP	Ag. Exports Share of 1961 GDP
Estimation:	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
Ag. Productivity (kcal/ha) x 10^{-6}	0.983**	1.174**	0.0197*	0.0272*	-1.467*	-1.978**	-0.0748*	-0.0252	-0.0428***	0.00329
	(0.414)	(0.503)	(0.0115)	(0.0156)	(0.761)	(0.955)	(0.0389)	(0.0321)	(0.0149)	(0.00340)
1960 log p.c. GDP x $\sum_t I_t$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F-statistic	10.15	8.13	8.03	6.08	10.11	8.08	10.15	10.15	9.6	9.6
Countries	115	102	109	97	114	101	115	115	110	110
Observations	798	714	763	679	791	707	798	798	770	770

Notes: All columns are 2SLS estimates that use a 7-year panel (years included are 1961, 1965, 1970, 1975, 1980, 1985, and 1990) and the unit of observation is the country-year. The outcome variable varies in each regression is listed at the top of the column. The F-statistic for the instrument, Predicted Agricultural Productivity, is reported at the bottom of the table along with the controls included in each specification. The first seven columns include a full set of time indicator interactions with log of per capita GDP in 1960 on the right hand side. Country and year fixed effects are included in all specifications. Robust standard errors clustered at the country level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table 8: Country-Level Heterogeneity: Involvement in Trade

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Urbanization		Ag. Employment Share		ln GDP		ln Per Capita GDP		Agricultural Exports		Agricultural Imports	
Predicted Productivity	0.0772	0.184	0.0700	-0.0775	0.0151**	0.0187**	0.0146**	0.0178**	0.0857	0.0987*	-0.0456	-0.000277
	(0.135)	(0.173)	(0.165)	(0.172)	(0.00754)	(0.00746)	(0.00725)	(0.00783)	(0.0569)	(0.0549)	(0.0724)	(0.0814)
Predicted Productivity x Above Median Openness	-0.409*		0.218		-0.0266**		-0.0352***		0.0882		-0.367*	
	(0.206)		(0.306)		(0.0124)		(0.0126)		(0.167)		(0.207)	
Predicted Productivity x 2nd Quartile Openness		-0.192*		0.251**		-0.00660		-0.00592		-0.0323		-0.0845*
		(0.110)		(0.120)		(0.00503)		(0.00512)		(0.0450)		(0.0448)
Predicted Productivity x 3rd Quartile Openness		-0.523**		0.361		-0.0308**		-0.0389***		0.0524		-0.422**
		(0.235)		(0.304)		(0.0121)		(0.0127)		(0.149)		(0.206)
Predicted Productivity x 4th Quartile Openness		-0.349		0.368		-0.0202		-0.0310**		0.412*		-0.224
		(0.267)		(0.340)		(0.0137)		(0.0142)		(0.247)		(0.266)
ln 1961 Population x Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ln 1961 GDP x Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country and Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Countries	72	72	69	69	72	72	72	72	70	70	70	70
Observations	2,880	2,880	2,760	2,760	2,880	2,880	2,880	2,880	2,100	2,100	2,100	2,100
R-squared	0.978	0.979	0.975	0.976	0.988	0.988	0.966	0.966	0.709	0.725	0.701	0.707

Notes: All columns report reduced form estimates estimates where the unit of observation is the country-year. The outcome variable varies in each regression is listed at the top of the column. Country and year fixed effects are included in all specifications; all specifications in Panel B also include 1961 log of GDP and population interacted with year indicators on the right hand side. Robust standard errors clustered at the country level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

9 Appendix Tables

Table A1: Effect of HYV Potential Benefit on HYV Adoption: Zeroth Stage

	(1)	(2)	(3)	(4)	(5)	(6)
Outcome Variable:		Fraction of Cultivated Land Using HYVs (Main Crops)				
		Baseline		Excl. Rice	Excl. Maize	Excl. Wheat
Predicted Ag. Productivity (kcal/ha) x 10^{-6}	0.324***	0.367***	0.307***	0.394***	0.341***	0.162**
	(0.0650)	(0.0623)	(0.0665)	(0.0759)	(0.0654)	(0.0782)
R-squared	0.855	0.870	0.886	0.853	0.860	0.752
Latitude x I_{1981} & Longitude x I_{1981}		Yes	Yes			
Coastline Indicator x I_{1981}		Yes	Yes			
1961 log Ag. Wages x I_{1981}			Yes			
1961 Population Density x I_{1981}			Yes			
1961 Literacy Rate x I_{1981}			Yes			
District & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Districts	271	271	271	271	271	271
Observations	542	542	542	542	542	542
R-squared	0.853	0.866	0.883	0.851	0.855	0.763

Notes: Each regression is an OLS model that uses two-year panel data of 271 Indian districts and the years 1961 and 1981. The outcome variable is the fraction of district cropland used for rice, wheat, maize, sorghum, and millet on which HYVs were used. The dependent variable is predicted productivity, measured in hg/ha in Panel A and kcal/ha in Panel B. District and year fixed effects are included in all specifications. Robust standard errors clustered at the district level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table A2: District-Level Analysis: First Stage Robustness Tests

	(1)	(3)	(5)	(4)
Outcome Variable:	Agricultural Productivity (kcal/ha)			
Predicted Agricultural Productivity, 5 Year Lead			-0.00727	
			(0.00942)	
Predicted Agricultural Productivity	0.111***	0.114***	0.0891***	0.219***
	(0.0214)	(0.0243)	(0.0178)	(0.0697)
Inst. Calculated Using 5 Crops (+ Barley, Sorghum)				Yes
Lag log Ag. Wage, lag Literacy Rate		Yes		
Latitude x $\sum_t I_t$	Yes	Yes	Yes	Yes
Longitude x $\sum_t I_t$	Yes	Yes	Yes	Yes
Coastline Indicator x $\sum_t I_t$	Yes	Yes	Yes	Yes
1961 log Ag. Wages x $\sum_t I_t$	Yes	Yes	Yes	Yes
1961 Population Density x $\sum_t I_t$	Yes	Yes	Yes	Yes
1961 Literacy Rate x $\sum_t I_t$	Yes	Yes	Yes	Yes
District & Year Fixed Effects	Yes	Yes	Yes	Yes
Districts	271	271	271	271
Observations	1,862	1,591	1,591	1,862
R-squared	0.625	0.629	0.855	0.632

Notes: Each regression uses a seven-year panel of 271 between 1957 and 1987. Each observation is separated by five years (1957, 1962, 1967, 1972, 1977, 1982, 1987). In Panel A the outcome variable is actual agricultural productivity measured in hectograms per hectare and in Panel B it is actual agricultural productivity measured in kilocalories per hectare. Predicted productivity is calculated using maximum potential yield for rice, wheat, and maize, except in columns 4 and 8. The five-crop version of predicted productivity used in column 4 is calculating using maximum potential yield for rice, wheat, maize, barley, and sorghum. The F-statistic for the instrument, predicted agricultural productivity, is reported at the bottom of the table. District and year fixed effects are included in all specifications. Robust standard errors clustered at the district level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table A3: District-Level Analysis: OLS Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Outcome Variable:	Good Roads	Med. Center	Literacy	log Ag. Wage	Farm Labor	Frac. Farm Labor	Farm Land	Frac. Non-Farm Labor	Urban Population	Manufac.
Agricultural Productivity (kcal/ha) x 10^{-6}	0.691**	0.454***	0.110**	1.610***	0.691***	0.158***	0.0202	-0.103**	-0.301*	-1.041***
	(0.268)	(0.123)	(0.0461)	(0.352)	(0.143)	(0.0516)	(0.116)	(0.0446)	(0.161)	(0.322)
R-squared	0.839	0.766	0.937	0.838	0.982	0.952	0.993	0.959	0.985	0.974
Latitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Longitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coastline Indicator x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	392	264	542	542	542	542	542	542	540	542

Notes: Each regression is an OLS model that uses two-year panel data of 271 Indian districts and the years 1961 and 1981. Agricultural productivity is measured in hectograms per hectare in Panel A and kilocalories per hectare in Panel B. District and year fixed effects are included in all specifications, along with a postyear interaction with latitude and longitude. Columns 4-5, 7, and 9-10 also include log population on the right hand side and column 7 also includes log of total land area on the right hand side. Robust standard errors clustered at the district level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table A4: District-Level Analysis: Productivity Calculated Using All Crops (Consumption and Non-Consumption)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Outcome Variable:	Good Roads	Med. Center	Literacy	log Ag. Wage	Farm Labor	Frac. Farm Labor	Farm Land	Frac. Non-Farm Labor	Urban Population	Manufac.
Productivity of Consumption & Non-Consumption Crops (hg/ha)	0.320***	0.128*	0.0440***	0.785***	0.227***	0.0607***	0.0809***	-0.0433***	-0.140***	-0.376***
	(0.117)	(0.0685)	(0.0157)	(0.114)	(0.0479)	(0.0178)	(0.0299)	(0.0157)	(0.0535)	(0.0923)
Latitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Longitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coastline Indicator x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	392	264	542	542	542	542	542	542	540	542
R-squared	0.822	0.755	0.938	0.804	0.982	0.952	0.992	0.960	0.985	0.974

Notes: Each regression is an 2SLS model that uses two-year panel data of 271 Indian districts and the years 1961 and 1981. In all specifications, agricultural productivity is calculated using a broader set of crops than in the main results, incorporating both crops intended for consumption and crops not intended for consumption (all crops included in Sanghi et al. 1998). Agricultural productivity is measured in hectograms per hectare since applying calorie content to cash crops is less logical. District and year fixed effects are included in all specifications, along with a postyear interaction with latitude and longitude. Columns 4-5, 7, and 9-10 also include log population on the right hand side and column 7 also includes log of total land area on the right hand side. Robust standard errors clustered at the district level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively. Robust standard errors clustered at the district level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table A5: District-Level Analysis: Reduced Form Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Outcome Variable:	Good Roads	Med. Center	Literacy	log Ag. Wage	Farm Labor	Frac. Farm Labor	Farm Land	Frac. Non-Farm Labor	Urban Population	Manufac.
Predicted Agricultural Productivity (kcal/ha)	0.208**	0.115*	0.0300	0.646***	0.168***	0.0480**	0.0609*	-0.0355**	-0.108*	-0.308***
	(0.103)	(0.0642)	(0.0187)	(0.115)	(0.0596)	(0.0208)	(0.0351)	(0.0174)	(0.0641)	(0.116)
R-squared	0.973	0.836	0.762	0.936	0.845	0.980	0.953	0.993	0.960	0.985
Latitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Longitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coastline Indicator x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	392	264	542	542	542	542	542	542	540	542

Notes: Each regression is the reduced form regression that uses two-year panel data of 271 Indian districts and the years 1961 and 1981. Agricultural productivity is measured in hectograms per hectare in Panel A and kilocalories per hectare in Panel B. District and year fixed effects are included in all specifications, along with a postyear interaction with latitude and longitude. Columns 4-5, 7, and 9-10 also include log population on the right hand side and column 7 also includes log of total land area on the right hand side. Robust standard errors clustered at the district level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table A6: District-Level Analysis: The Role of Migration

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Outcome Var., parameterized as log(1+x)	Farm Labor	Manufac.	Commerce	"Other" Service	Transport. &	Household Industry	Migrants
	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
Agricultural Productivity (kcal/ha) x 10^{-6}	1.170***	-0.984*	-0.414	-0.326	0.360	-0.261	-1.276***
	(0.323)	(0.572)	(0.344)	(0.364)	(0.482)	(0.224)	(0.347)
log Incoming Migrants	0.000869	0.453***	0.338***	0.294***	0.354***	0.0989***	
	(0.0668)	(0.0968)	(0.0597)	(0.0645)	(0.0820)	(0.0291)	
Latitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Longitude x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Coastline Indicator x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes
1961 log Ag. Wages x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes
1961 Population Density x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes
1961 Literacy Rate x I_{1981}	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	542	542	542	542	542	542	542
R-squared	0.981	0.977	0.988	0.983	0.980	0.992	0.969

Notes: Each regression is a 2SLS model that uses two-year panel data of 271 Indian districts and the years 1961 and 1981. The independent variable of interest is agricultural productivity measured in kilocalories per hectare. Columns 1-8 also include (log of) income migrants from other Indian districts on the right hand side—the coefficient is displayed. In column 8, it is the outcome variable. Other outcome variables are measures of employment from the previous two tables. District and year fixed effects are included in all specifications. Robust standard errors clustered at the district level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table A7: Country-Level Analysis: First-Stage Robustness Checks

	(1)	(2)	(3)	(4)
	Outcome is Agricultural Productivity (kcal/ha)			
Predicted Productivity (kcal/ha), 5 year lead			0.0722	
			(0.151)	
Predicted Productivity (kcal/ha)	0.159***	0.267**	0.234***	0.221**
	(0.0499)	(0.108)	(0.0818)	(0.0861)
Instrument calculated using 3 crops (rice, wheat, maize)	No	No	No	Yes
Lag log p.c. GDP, lag urbanization	No	Yes	Yes	No
Country & Year Fixed Effects	Yes	Yes	Yes	Yes
log 1960 p.c. GDP x $\sum_t I_t$,	Yes	Yes	Yes	Yes
Observations	798	678	678	798
R-squared	0.967	0.971	0.971	0.967

Notes: All columns use a 7-year panel (years included are 1961, 1965, 1970, 1975, 1980, 1985, and 1990). Agricultural productivity measured in kilocalories per hectare is the outcome variable in all columns. The independent variable of interest is the predicted agricultural productivity instrument. The three-crop version of predicted productivity used in column 4 is calculating using maximum potential yield for rice, wheat, and maize. Country and year fixed effects are included in all specifications, along with the characteristics listed at the bottom of each panel. Robust standard errors clustered at the country level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table A8: Country-Level Analysis: First Stage Falsification Exercise

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Outcome Variable is Actual Agricultural Productivity (kcal/ha)						
	Panel A: Green Revolution Crops						
Crop Pair:	Rice and Wheat			Cassava and Potato		Bean and Barley	
Years:	1961/1990	1961/1990	1990/2000	1961/1990	1990/2000	1961/1990	1990/2000
Predicted Productivity of crop pair noted in each column	0.109**	0.115**	0.162**	0.115**	0.172**	0.184**	0.193*
	(0.0524)	(0.0529)	(0.0716)	(0.0549)	(0.0749)	(0.0829)	(0.114)
p-value of income itneraction		[0.124]					
Country & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	230	230	230	230	230	230	230
R-squared	0.962	0.963	0.936	0.963	0.937	0.962	0.934
	Panel B: Placebo Crops						
Crop Pair:	Rye and Oats			Onion and Carrot		Cowpea and Tomato	
Years:	1961/1990	1961/1990	1990/2000	1961/1990	1990/2000	1961/1990	1990/2000
Predicted Productivity of crop pair noted in each column	0.00934	-0.0391	0.142	0.134	-0.0189	0.130	-0.0703
	(0.0797)	(0.105)	(0.239)	(0.104)	(0.158)	(0.120)	(0.349)
p-value of income itneraction		[0.295]					
Country & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	230	230	230	230	230	230	230
R-squared	0.961	0.961	0.933	0.962	0.932	0.962	0.932

Notes: Each regression uses two-year panel data of 115 countries. The years in the panel—either 1961 and 1990 or 1961 and 2000—are listed at the top of each column. In all columns the outcome variable is agricultural productivity (kcal/ha). In each specification, the independent variable of interest is a measure of predicted agricultural productivity calculated using the crop pair listed at the top of the column. Crops used in Panel A are also used to construct the main version of the predicted productivity instrument in the first and second stage results. Crops used in Panel B, I argue in the main text, were not significantly influenced by the Green Revolution. Country and year fixed effects are included in all specifications. Log per capita GDP interacted with a dummy that equals 1 if the year is 1990 is included on the right hand side in column 2. The p-values of the corresponding F-tests for the controls are reported in brackets and robust standard errors clustered at the country level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table A9: Cross-Country Analysis: OLS Estimates

	(1)	(2)	(3)	(4)	(5)
	Frac. Land Devoted to	Agricultural Employment			
Outcome Variable:	Ag.	Share	Urbanization	log GDP	log p.c. GDP
Estimation:	OLS	OLS	OLS	OLS	OLS
Agricultural Productivity (kcal/ha) x 10^{-6}	-0.0188	-0.00146*	0.0141	0.00626**	0.00830***
	(0.0947)	(0.000827)	(0.0429)	(0.00310)	(0.00276)
Country & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
log 1960 p.c. GDP x $\sum_t I_{t,}$	Yes	Yes	Yes	Yes	Yes
Observations	798	763	791	798	798
R-squared	0.986	0.974	0.975	0.988	0.961

Notes: All columns are OLS estimates that use a 7-year panel (years included are 1961, 1965, 1970, 1975, 1980, 1985, and 1990) and the unit of observation is the country-year. The outcome variable varies in each regression and is listed at the top of the column. Country and year fixed effects, along with a full set of year indicator interaction with log per capita 1960 GDP, are included in all specifications. Robust standard errors clustered at the country level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.

Table A10: Cross-Country Analysis: Reduced Form Estimates

	(1)	(2)	(3)	(4)	(5)
	Frac. Land Devoted to	Ag. Employment			
Outcome Variable:	Ag.	Share	Urbanization	log GDP	log p.c. GDP
Predicted Productivity (kcal/ha)	0.156***	0.00292**	-0.237**	-0.0119**	-0.00472
	(0.0526)	(0.00121)	(0.0993)	(0.00546)	(0.00474)
Country & Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
log 1960 p.c. GDP x $\sum_t I_{t,}$	Yes	Yes	Yes	Yes	Yes
Observations	798	770	798	798	798
R-squared	0.987	0.975	0.978	0.988	0.961

Notes: All columns are reduced form estimates that use a 7-year panel (years included are 1961, 1965, 1970, 1975, 1980, 1985, and 1990) and the unit of observation is the country-year. The outcome variable varies in each regression and is listed at the top of the column. Country and year fixed effects, along with a full set of year indicator interaction with log per capita 1960 GDP, are included in all specifications. Robust standard errors clustered at the country level are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively.