

The Development Effects of the Extractive Colonial Economy: The Dutch Cultivation System in Java*

Melissa Dell and Benjamin A. Olken

Harvard University and MIT

April 2017

Abstract

Colonial powers typically organized economic activity in the colonies to maximize their economic returns. While the literature has emphasized the negative impacts of colonial extraction on institutional quality, the changes in economic organization implemented to spur production may have countervailing long-run effects. We examine these in the context of the Dutch Cultivation System, the integrated industrial and agricultural system for producing sugar that formed the core of the Dutch colonial enterprise in 19th century Java. We show that areas close to where the Dutch established sugar factories in the mid-19th century are today more industrialized, have better infrastructure, are more educated, and are richer than nearby counterfactual locations that would have been similarly suitable for colonial sugar factories. We also show, using a spatial regression discontinuity design on the catchment areas around each factory, that villages forced to grow sugar cane have more village owned land and also have more schools and substantially higher education levels, both historically and today. The results suggest that the economic structures implemented by colonizers to facilitate production can continue to promote economic activity in the long run.

Keywords: long-run development, colonialism, economic organization

*We thank Abdulhamid Haidar, Mateo Montenegro, Roman Zarate and especially Peter Hickman for providing excellent research assistance in Cambridge, Timo Stibbe for assistance in the Netherlands, and Muhammad Abduh for assistance in Indonesia. Contact email: melissadell@fas.harvard.edu, address: Harvard University Department of Economics, Cambridge MA 02138.

1 Introduction

Extractive colonial institutions are often thought to be an important reason why some places remain persistently poorer than others. Colonizers whose focus was on extracting resources from the colonies plausibly set up weak institutions with poor property right protections to facilitate this extraction (Acemoglu et al., 2001). These institutions may in turn have persisted, continuing to be a drag on economic performance today.

Counterbalancing this effect, however, is the fact that in many cases colonial powers established complex economic systems to create the economic surplus that they wanted to extract. In the case of agricultural extraction, crops needed to be grown, harvested, and processed before being transported to the home country. This typically involved technology transfer and infrastructure construction, such as the large-scale sugar processing and transport infrastructure that the Japanese created in colonial Taiwan or the tea plantations, processing plants, and railroads that the British established in colonial India. Extracting a large surplus also often required a substantial reorganization of the indigenous local economy, and these changes likewise may have had long-run effects that differ from the macro institutional impacts.

In this paper, we examine the persistent local development effects of extractive institutions by studying the Dutch Cultivation System (*Cultuurstelsel*) in Java. The island of Java was the main population center of the vast Dutch colonial empire in the East Indies, and with a modern population of over 160 million, remains the economic and population center of Indonesia today. From the early 1830s through the 1870s, the Dutch colonial state forced peasants along Java's northern coast to cultivate sugar, which was then processed in nearby colonial factories for subsequent export to Europe. The revenues extracted from this system made Java among the world's most lucrative colonies, and at their peak accounted for over one-third of Dutch government revenue (Luiten van Zanden, 2010). Sugar production continued after the abolition of the System but collapsed during the Great Depression, as Indonesia lacked protected markets, and today Indonesia is the world's largest sugar importer.

Prior to the Cultivation System, the Javanese economy was heavily specialized in rice cultivation, largely for subsistence consumption. The System required a substantial reorganization of economic life in order to operate. At its heart were 94 water-powered sugar factories set up by the Dutch, which crushed recently harvested cane and processed it into refined sugar. Over the course of the System, millions of Javanese worked in sugar processing and transport - via both forced and free labor - a major shift since factory production had been nearly non-existent initially (Elson, 1994, p. 215). Since raw cane is heavy and needs to be refined quickly after harvest, the poor infrastructure of the time meant that the cane had to be grown nearby. Accordingly, the Dutch constructed a catchment area with a radius of

approximately four to seven kilometers around each factory, and forced all villages within the catchment area to reorganize their land to grow cane. Village officials within the catchment areas were politically empowered by the colonial government to make this happen.

Colonial extraction frequently both reorganized village life around coerced labor and introduced large-scale infrastructure investments and agglomeration forces. It can be difficult to separate the persistent legacies of these diverse impacts. The unique features of the Cultivation System, combined with a variety of highly disaggregated modern data, allow us to separately identify how the construction of a sugar processing infrastructure and forced sugar cultivation each influence disparities in economic prosperity in Java today.

To estimate the effects of creating a modern sugar processing infrastructure, we use GIS to construct the locations of alternative, counterfactual sites that would have been equally plausible locations for sugar factories and hence can serve as a comparison group. We then estimate the effects of being close to an actual factory and compute p-values by comparing these effects to the distribution of estimated effects of being near plausible counterfactual sites. To identify plausible counterfactual locations, we exploit the fact that while there were various similarly suitable locations, catchment areas were often adjacent to one another. The factories could not be too close since each required an adequately sized catchment area - once one factory was placed, it constrained the locations of other nearby factories. Hence, there were many possible equilibria for site selection (Salop, 1979).

Nearby sites are classified as feasible if they meet the following two criteria. Since state-of-the-art processing technology was water-powered, sugar factories were located along rivers; we therefore only consider sites that can be reached by moving upstream or downstream via river from the actual factory. Moreover, factories needed to have a sufficient amount of sugar-suitable land nearby; we therefore only consider sites for which the amount of suitable land, based on elevation and slope, is similar to the amount of suitable land near actual factories. We also constrain placebo factories for each draw to be at least as far apart as the 10th percentile in the distribution of actual factories. The approach is illustrated in Figure 1. To separate the effects of Dutch extraction from modern sugar processing, we focus on Dutch colonial factories where there is no modern sugar factory nearby, though most of our results are similar if we consider the entire sample.

Second, to identify the impact on a village of being subjected to forced sugar cultivation, we use a spatial discontinuity across the borders of the catchment areas. We obtained a 19th century handwritten list of the over 10,000 villages subjected to forced cultivation from Dutch archives in the Hague, and matched these villages with modern georeferenced locations. This allowed us to compute the exact location of the Cultivation System boundaries (see Figure 2). Inside, villages cultivated sugar cane for the government, whereas outside they did not. The boundaries form a multi-dimensional discontinuity in longitude-latitude space and allow

us to control for smooth geographic variation, including in proximity to the nearest historical factory. Identification requires that pre-determined characteristics change continuously at the catchment area boundaries, and we provide evidence that this assumption is reasonable.

We document that the Dutch colonial sugar industry substantially transformed economic activity in contemporary Java, and agglomeration and infrastructure investments appear to be important channels of persistence. SUSENAS household surveys, collected between 2001 and 2011, show that people living within a few kilometers of historical sugar factories are much less likely to be employed in agriculture and more likely to be employed in manufacturing or retail than people living further away. Effects are similar in Census data from 1980, when Indonesia was more agricultural, suggesting that impacts are not just driven by recent industrialization. A similar analysis using distance to counterfactual locations provides strong evidence that these differences are not driven by geography or chance.

The composition of industries in these areas suggests agglomeration through input-output linkages as a channel of persistence. Using national input-output tables, we classify manufacturing employment in the 2006 Economic Census into employment that is upstream and downstream from sugar processing. Upstream manufacturing industries - those whose outputs serve as direct or indirect inputs to sugar processing - include farm and capital machinery. Downstream industries - those that use processed sugar as an input - include most other food processing industries. Areas near historical factories have a higher share of employment in upstream and downstream industries, as compared to places further away. When we restrict to locations with no modern sugar factory, the upstream effect disappears, but we still find substantially more employment in downstream industries. While the Dutch claimed the high quality processed sugar, the factories were permitted to sell the low quality sugar, which was costly to transport, on local markets. Industries that used sugar as an input plausibly sprung up around colonial sugar factories. Even when the original source of sugar disappeared, those industrial centers persisted, sourcing sugar from elsewhere.

The construction of a sugar processing infrastructure also impacted public investments. The Dutch built road and rail infrastructure to transport processed sugar to the ports, and this infrastructure has persisted long after the abolition of the System, plausibly promoting trade and economic activity through the present. Villages located within a few kilometers of a historical sugar factory were more likely to have a paved road in 1980 and today have a much higher density of intercity and local roads, as well as railroads, than places just a couple of kilometers further away. It is hard to imagine a scenario where the Dutch would have made these infrastructure investments in the absence of extractive colonial institutions, as they invested explicitly to extract a large colonial surplus.

Over time the more industrialized and integrated areas near historical factories may have gained differential access to public goods more generally. We find that villages near a

historical factory were much more likely to have electricity in 1980. They were also more likely to have a high school in 1980, a time when high schools were very rare. Using the complete 100% sample microdata from the 2000 Census, we document that people in villages within a few kilometers of colonial factories are more educated. This is true both for cohorts born in the 1920s, who completed their education during the Dutch era, and for more recent cohorts who completed their education under an independent Indonesia.

The historical construction of sugar processing plants promoted agglomeration - at least in part through input-output linkages - and public investments, which have persistently influenced the structure of the economy. Consistent with this, household survey data show that households living within a few kilometers of a historical factory location have per-capita consumption that is about 14 percent higher than those living more than 10 kilometers away.

It is difficult to say what would have happened if the Dutch had not shown up at all, as this counterfactual did not occur and instead existing variation allows us to examine the impacts of being more intensively affected by colonial extraction. Nevertheless, it seems reasonable to speculate that while industrial production centers may have arisen in rural Java in the absence of outside intervention, this would have likely occurred much more slowly. The economy in 1830 was primarily agricultural, landless peasants were customarily attached to land-owning households and hence not readily available for factory labor, and transport infrastructure and external finance were nearly non-existent, all major barriers to industrial production that the Cultivation System had to alter in order to create an extractable surplus.

The discussion thus far has focused on one aspect of the Cultivation System, namely the establishment of the factories. The system also subjected villages to forced cane cultivation, fundamentally altering their economic and political organization. We examine the impacts on these villages using a spatial discontinuity design that compares them to villages just outside the cultivation catchment areas. In principal, effects on subjected villages could undo the impacts of the factories, but instead we find that they tend to reinforce them.

There were relatively few Dutch officials in Java, especially in rural areas. Instead of directly administering the Cultivation System, the Dutch empowered Javanese village heads to manage it, giving them considerably greater command over land and labor than they had exercised previously. In particular, village heads were tasked with choosing the location of the cane fields, which had to be rotated each year. This required regularly reassigning land between public and private use, and village heads were free to manage this as they saw fit as long as cane quotas were met. As discussed in Section 2, the historical literature suggests that village heads used this opportunity to redistribute land to the village - where they would firmly control it - and to their cronies.

We therefore begin by examining how this process changed the modern distribution of land, focusing on two types of land. First, Javanese villages typically have some amount of

land set aside permanently for use by the village for public purposes. Temporary use rights over some of this land (known as *tanah bengkok*) are given to village heads and other village officials as compensation for their service; revenues from the use of other land enters the village treasury for other purposes. We find that there is about 10 to 15 percent more village public-use land in Cultivation System villages. This has remained fairly constant over time, with virtually identical effects in 1980 and 2003. Second, we examine the distribution of private agricultural land. Effects on private land inequality are if anything positive but fall short of statistical significance.

We then turn to economic structure, documenting that households in Cultivation System villages are more educated. Echoing the previous results, educational effects for households in these villages go all the way back to the 1920s cohort, which was educated during the Dutch period. Since we control flexibly for distance to the nearest historical factory and geographic location, these effects isolate the impacts of being subjected to forced cultivation, and are above and beyond the effects of being closer to a factory. During the early 20th century, villages that wanted a school needed to fund the school building themselves. To the extent that collective action problems were endemic - which various scholars argue tends to be the case in low-income rural settings (Popkin, 1979; Foster, 1965) - the revenue from the extra village land may have been essential to funding schools. Indeed these places appear to have had more school infrastructure prior to the nationally-funded school construction program begun in the 1970s and studied by Duflo (2001). The village heads, who receive more land as compensation, are also more highly educated in cultivation village today, to an even greater extent than we would expect given the overall higher education levels.

Differences in human capital - both of village inhabitants and the village head - may have impacted economic structure more broadly. Indeed, cultivation villages have a higher percentage of households working in manufacturing and retail and fewer in agriculture, both in 1980 and today. They also have more manufacturing firms on average than villages located just outside the discontinuity, and are more densely populated. People living in these villages are no poorer than those living just outside.

The long-run development effects of forced cultivation do not undo the impacts of the sugar processing infrastructure, and if anything go in the same direction. Economic changes induced by the Cultivation System led to a more modern economy, with more manufacturing and retail and greater education levels - both for the areas in the immediate vicinity of the historical sugar factories and further out where the village was reorganized to grow cane.

These results inform three different debates about the historical origins of modern economic structure. First, they speak to a classic literature on the role of input-output linkages in development. During the 1950s and 60s, development economists hypothesized that linkages between sectors provided a key mechanism for the propagation of structural change

and debated the merits of targeting industrial policy based on input-output linkages. (Rasmussen, 1956; Myrdal, 1957; Hirschman, 1960). More recent evidence documents the importance of input-output linkages in transmitting macroeconomic shocks (Acemoglu et al., 2016; Baqaee, 2015; Acemoglu et al., 2012; Carvalho, 2009) and in development (Liu, 2017; Lane, 2017). However, there is little empirical evidence on the role that they play in the long run. The Cultivation System provided local markets with low grade sugar that was costly to transport, generating benefits from agglomeration with downstream producers. Relative to other agricultural products, sugar has very dense linkages, since it is added to most processed foods. We find substantial effects on downstream firms that persisted even after the original sugar factories disappeared, plausibly because of agglomeration of downstream producers and the endogenous concentration of population near initial production sites.

Second, to the extent that not all returns from village lands are extracted for private consumption, the persistently larger public landholdings in Cultivation System villages provided the local government with additional ability to raise revenue, a classic measure of state capacity, which we might expect to increase public goods provision. A classic literature (Hayami, 1980; Scott, 1977), as well as more recent empirical evidence (Dell et al., 2017), argue for the importance of village level organization for public goods provision and long-run economic development. Our results suggest that village government is indeed important, even as the central state has become increasingly involved in public goods provision in recent years. Public goods traditionally provided by villages, and schooling in particular, are plausibly central to development (Gennaioli et al., 2013).

Third, the study contributes to an extensive literature on the long-run impacts of colonial institutions and relates to a large body of work that highlights the relevance of persistence for economic development.¹ In a series of innovative cross-country studies, Acemoglu et al. (2002, 2001) hypothesize that extractive colonial institutions are at the heart of much of the developing world's poor long-run growth performance. A subsequent literature aimed to refine identification and isolate channels of persistence by focusing on within country variation, with several studies examining forced labor of various types. In particular, Dell (2010) uses a spatial discontinuity to document that forced labor in the Potosi silver mines during Spanish colonization lowered economic prosperity in the long-run, a finding echoed in an examination of forced rubber cultivation in the Congo by Lowes and Montero (2016).

On the other hand, historical evidence highlights that extractive colonial institutions of-

¹See Becker et al. (2016); Guiso et al. (2016); Lowes et al. (2015); Acemoglu et al. (2015); Bukowski (2016); Oto-Peralías and Romero-Ávila (2014); Grosjean (2014); Michalopoulos and Papaioannou (2014); Valencia Caicedo (2014); Acemoglu and Robinson (2013); Spolaore and Wacziarg (2013); Michalopoulos and Papaioannou (2013); Alesina et al. (2013); Voigtländer and Voth (2012); Acemoglu et al. (2011); Nunn and Wantchekon (2011); Luttmmer and Singhal (2011); Grosjean (2011); Tabellini (2010); Fernández and Fogli (2009); Tabellini (2008); Nunn (2008); Guiso et al. (2008); Gennaioli and Rainer (2007); Giuliano (2007); Banerjee and Iyer (2005); Glaeser and Shleifer (2002).

ten had major effects on the local organization of economic production. To the extent that these systems were accompanied by major infrastructure investments or agglomeration, a variety of evidence suggests that long-run economic effects could be positive. For example, Jedwab and Moradi (2016) and Jedwab et al. (2015) provide evidence that colonial railroads persistently affected the distribution and aggregate level of economic activity in Africa, and Donaldson (2010) documents that railroads in colonial India significantly improved welfare at the time. More generally, various evidence points to the benefits of infrastructure construction (Donaldson and Hornbeck, 2016; Banerjee et al., 2012; Michaels, 2008; Baum-Snow, 2007; Fogel, 1964) and agglomeration (see Glaeser and Gottlieb (2009) for a review).

While colonial states frequently extracted revenues via forced labor (Van Waijenburg, 2015), *how* these institutions impacted the local economy varied considerably. Labor coercion might occur in an individual's community or in a single location hundreds of miles away, with infrastructure built only around that location, and processing might happen on site or raw materials might instead be exported directly. The intensity of technology transfers and linkages to other sectors varied, as did the counterfactual to colonial extraction. All of these features plausibly matter in determining how colonial extraction influences disparities in economic prosperity today. This study documents how in-situ changes in production, combined with nearby industrial processing, increased economic prosperity in the long run. The next section discusses evidence that the Dutch in Java were not the only colonial power to implement an economic system that encouraged a more modern economy, and the Javanese are unlikely to be unique in utilizing colonial structures to generate economic activity long after the colonizers had left.

The remainder of the study is organized as follows. Section 2 discusses the historical context, and Section 3 describes the data. Section 4 examines the long-run impacts of the sugar processing infrastructure established by the Cultivation System, whereas Section 5 estimates the effects of being subjected to forced cultivation. Section 6 concludes.

2 Historical Background

2.1 Historical Overview

The Cultivation System (*Cultuurstelsel*), in force from the early 1830s through the 1870s, compelled Javanese villagers to produce export crops for the Dutch colonial government. At its peak, the Cultivation System provided over one third of Dutch government revenues and four percent of Dutch GDP (Luiten van Zanden, 2010; Ricklefs, 2008, p. 159), making Java one of the world's most financially lucrative colonies. While a variety of crops were grown, from the 1850s onward sugar and coffee accounted for more than 96% of profits (Elson, 1994,

p. 135).² We focus on sugar cultivation, which took place on the populated plains of Java's Northeast Coast. Coffee, while profitable, was primarily grown on forested mountain slopes that were uninhabited initially (Elson, 1994, p. 65). Extant data on coffee are at a higher level of aggregation that do not permit the detailed analyses that we perform for sugar.

The sugar system was a major enterprise, as Java was the world's second largest sugar producer at the time, surpassed narrowly only by Cuba. The use of coercion in sugar production was typical at this time, see Dippel et al. (2015) for a theoretical and empirical examination of how returns from sugar were invested in coercive institutions in the Caribbean. The system started in the early 1830s with over 70,000 sugar growers cultivating cane for 59 factories, and over time grew to encompass 94 factories and a large share of the Javanese population (Elson, 1994, p. 55).³ A Dutch report cited by Fasseur (1992) notes that as much as 25 percent of the native Javanese population was involved in the Cultivation System (excluding coffee), of which sugar was the largest component. Dutch reports from the 1860s show that over 2.5 million workers labored in the sugar factories or related services (i.e. transporting cane and firewood to the factories), with the number of free laborers expanding significantly across time (Elson, 1994, p. 215). Sugar factories were run by private entrepreneurs with close links to the Crown, who were required to sell their high quality sugar to the government. Villages surrounding each factory had to allocate up to one-fifth of their land to grow cane, though in practice this ratio was often significantly exceeded (Elson, 1994, p. 229; Van Niel, 1992, p. 137). Europeans were prohibited from renting or purchasing land and from establishing their own private factories, and hence sugar production occurred almost exclusively through the Cultivation System.⁴

The Cultivation System was primarily administered by Javanese officials, who gained considerably greater command over land and labor than they had exercised previously (Van Niel, 2005, Elson, 1994, p. 183). Village heads were in charge of allocating sugar cultivation land and assigning labor to cultivation services. The village received incentive payments (*kultuurporcenten*) for cultivating cane, and the village head was also tasked with distributing these and with collecting the land tax owed to the Dutch government.

The Dutch phased out the Cultivation System in the 1880s, and sugar expanded rapidly under free enterprise (Elson, 1984, p. 131). However, the industry collapsed during the Great Depression, and subsequently Java's primary competitors enjoyed large protected markets, whereas Java did not.⁵ Moreover, much of Europe and North America subsidized and pro-

²Indigo was grown initially but the Dutch started phasing it out in 1834 due to low profitability, and indigo cultivation declined sharply throughout the 1840s (Elson, 1994, p. 83, 110, 131; Van Niel, 1992, p. 112). Cinnamon, tea, and tobacco were grown by a small number of peasants in localized areas.

³The factories were primarily established in the 1830s and after that changed little. (Elson, 1984, p. 71).

⁴The exceptions were the indirectly ruled Principalities, which had private estate agriculture and were exempt from the Cultivation System.

⁵For example, the Philippines and Hawaii had protected access to the U.S., and Taiwan had protected

tected new domestic beet sugar producers. Today Indonesia is a major sugar importer, with sugar contributing only 0.05% of their agricultural exports.

Sugar cultivation exerted a variety of impacts, on subjected villages as well as on villages in proximity to sugar factories. Much of the historical literature has focused on distributional effects on subjected villages. An older literature argues that the System led to an equalizing of the land distribution, as villagers aimed to share the System's burdens (Geertz, 1963). This equalization is hypothesized to have stifled modernization by preventing the emergence of a class of prosperous entrepreneurial farmers. A subsequent literature likewise hypothesizes that the Cultivation System had equalizing effects, through increasing the amount of communal village land or through leading communal land to be divided more equally.⁶

A revisionist literature, initiated by Elson (1994, 1984), argues instead that the Cultivation System empowered Javanese village elites, increasing hierarchy within the village. There were only a handful of Dutch officials in Java, concentrated in major cities, so the Dutch relied on Javanese authorities to ensure that resources were devoted to government cultivation. In order to guarantee their cooperation, the Dutch made village authorities accountable to the colonial state and ensured that they benefited from furthering Dutch objectives.⁷ In particular, village heads were responsible for assigning land for government cultivation - which had to be rotated on a regular basis - and they often used this as a pretext to redistribute land to the village (where they controlled the output) or to their cronies (Van Niel, 1992, p. 139). By 1857, 20% of all the irrigated land in Probolinggo Residency - a center of sugar cultivation - was controlled by officials, in Kedu village chiefs had taken over more than half of the land, and similar patterns obtained in Pekalongan-Tegal and other sugar-intensive areas (Elson, 1984, p. 94). Ricklefs (2008, p. 159) likewise argues that the increase in village landholdings - regardless of whether the motivation was to share the burdens of the system - increased the power of village heads.

Despite these distributional effects, historians such as Elson (1994) have argued that the Cultivation System positively impacted economic outcomes for the typical villager. Subjected communities received incentive payments based on the amount of sugar produced. Historians hypothesize that crop payments, by injecting currency into what had previously been a non-monetized subsistence economy, promoted modern exchange (Elson, 1994, p. 261, Ricklefs, 2008, p. 158). Residencies (provinces) more involved in the Cultivation System experienced a greater growth in markets while the System was in force. Moreover, the

access to Japan.

⁶See Elson (1994, p. 162) for a review.

⁷Traditionally, the head's term was limited and required consent from landholders. Dutch officials increasingly intervened to install village authorities friendly to their interests and resisted the periodic rotation of village heads (Ricklefs, 2008, p. 158; Elson, 1994, p. 172-174). Following an 1854 reform, succession of the village head was made hereditary.

payments often exceeded the land tax owed to the Dutch (Elson, 1994, p. 311, Van Niel, 1992). In places without forced cultivation, money for taxes had to instead be obtained by selling rice at low prices to Chinese traders, who colluded with officials to monopolize the rice trade. It should be noted though that government sugar land was always assessed at the highest rate, so if anything it increased the village's tax burden.

Historians have also pointed to modernizing economic effects that depend on proximity to the sugar processing centers. Extensive road and rail infrastructure were constructed to connect sugar producing regions to ports, plausibly promoting market integration (Elson, 1994, p. 251-252, Ricklefs, 2008, p. 158). Moreover, the Cultivation System increased economic specialization by creating a variety of economic opportunities outside forced cultivation (Elson, 1994, p. 207-208). Factory and transport workers were often paid a wage, and there were also employment opportunities making baskets and matting to transport the cane, making clay pots for use in the factories, and transporting firewood to the factories. Moreover, while the high grade refined sugar was exported, the low grade sugar could not be exported due to a high water content. The factories were permitted to sell this sugar locally, potentially spurring other food processing establishments to locate nearby.

The Dutch were not the only colonial power that attempted to institute major economic changes in order to extract more surplus. The Japanese in Taiwan are a classic example of a developmental colonial state that made massive investments in increasing output in order to maximize extraction, including the construction of a large-scale, modern sugar-processing infrastructure. The establishment of tea plantations and a network of tea processing plants in India - including the formation of the Assam Tea Company under the auspices of the colonial British government - likewise bears similarities to the Cultivation System, as does the establishment of a colonial Indian jute production and processing infrastructure (Tomlinson, 2014). Van Waijenburg (2015) documents the widespread nature of labor conscription by colonizers in Africa, highlighting many different types of economic arrangements including some that involved labor in agricultural export and processing enterprises.

2.2 Assignment to Sugar Cultivation

Understanding why sugar factories were located in particular places and why certain villages were subjected is central to identifying the long-run impacts of the Cultivation System. During the 18th century, sugar cultivation in Java was limited to a small number of estates in the hinterland surrounding Batavia (now Jakarta), the capital of the Netherlands Indies and the location of most European settlement. These estates crushed cane using wooden rollers pulled by water buffalo and the output was consumed in local markets. During the 19th century, Batavian cane production nearly disappeared, as the technology was primitive

and the land marginally suitable (Van Niel, 2005, p. 133, 139).

Dutch official correspondence - analyzed by van Schaik (1986, p. 183) - highlights several factors used to locate sugar factories. Sugar grows well in plains and valleys and does not grow well on mountain slopes, and factories had to be near locations suitable for cultivation since raw cane was costly to transport. However, sugar was not grown in all agronomically suitable locations, as illustrated by the large-scale expansion of cultivation following the end of the government system. Proximity to water was crucial because sugar processing technology in 1830 used iron crushing cylinders powered by water wheels located along rivers. It was not until the late 19th century that steam power became dominant. Proximity to timber was also important but is unobserved, since rapid deforestation occurred throughout the 19th century and forest maps prior to the Cultivation System have not survived. Finally, there needed to be nearby population to cultivate the cane and the factories could not be located in areas where transport costs were prohibitive; hence, the remote valleys of southern Java were not suitable, despite their agronomic conditions. The government factories were created in the 1830s, and were relatively fixed for the remainder of the period.

Each factory was allocated a certain area from which to draw the land and labor it required. Villages were typically within four to seven kilometers of their factory. The Dutch referred to this as the ‘circle system’. Catchment areas tended to be adjacent to each other and avoided major pre-existing cities, which primarily consisted of residency (province) capitals.⁸ Once one factory was sited, another one could not be located too close even if there were many suitable places along a river, since each factory required a sufficiently sized catchment area. We will exploit this fact to create suitable counterfactual factory locations that can be used as a comparison group.

3 Data

3.1 The Cultivation System

Data on the Cultivation System are drawn from handwritten manuscript archival records held by the Hague (Commissie Umbgrove, 1858). See Appendix Figure X-1 for an example manuscript. Constitutional reforms in the Netherlands in 1848 placed colonial affairs under partial parliamentary control, and the Minister of Colonies subsequently ordered an inquiry into government sugar cultivation. The commission was appointed in 1853 and spent four years collecting detailed data. Handwritten documents list which villages contributed to

⁸Colonial Java was divided into 24 residencies, which were the center of colonial administration and the major cities of the time. Modern Indonesian provinces are larger, with only 6 provinces on Java, but the colonial residency capitals by and large remain the largest cities today.

each sugar factory and how much land and labor each village provided. They also contain extensive qualitative information about the social situation in sugar producing areas and sometimes show sketches of the catchment areas.

Historical villages are matched with coordinates from the U.S. National Geospatial Intelligence Agency’s extensive Geonames database, which provides a detailed list of populated places in Indonesia, including many sub-village units. Matches are made using the historical village’s name and the location of its historical district.⁹ The manuscripts also list the distance between the village’s fields and the factory, which allows us to distinguish between multiple matches and provides an additional check. We match 6,383 historical villages with Geonames coordinates, which are located in 2,519 modern village polygons. Villages are small - there are over 30,000 in Java today - and hence the monographs allow us to know the precise location of government sugar cultivation. The sugar factories are matched using the same procedure. The combination of the historical sub-village, plus the fact that the factory had to be adjacent to a river, means that the factories can be precisely located.

Not all historical sugar villages can be matched, as some disappeared or changed their names, leaving gaps in the matched catchment areas. The actual catchment areas were contiguous, and hence we infer a contiguous set of subjected villages by drawing straight lines from each sugar factory to the coordinates of contributing villages, assigning any unmatched village along the lines as treated. This method is illustrated in Figure 3, which shows the factories, matched village points, straight lines, and catchment boundaries for an example catchment area. Modern village polygons are shown in the background. 82% of modern villages in the constructed catchment areas are matched to at least one village in the historical dataset. We designate all villages inside these catchment areas as treated. The appendix shows that results are highly robust to instead designating matched villages as treated and instrumenting these with being inside the catchment areas.¹⁰

3.2 Outcome Data

Outcome data are drawn primarily from population, industrial, agricultural, and village censuses and household surveys. These data were collected between 1980 and the present by the Indonesian government’s Central Bureau of Statistics (BPS) and are described in more detail in the data appendix (Figure X-2). All of these datasets contain village identifiers that allow us to match the microdata to modern village boundaries.¹¹

⁹Data are available for download at <http://geonames.nga.mil/gns/html/index.html>.

¹⁰The first stage is very strong, with an F-stat of around 1,450 when standard errors are clustered at the sub-district level, making these estimates similar by construction.

¹¹Note that by “village”, we refer to the lowest administrative level in Indonesia. In rural areas these are known as *desa*; in more urban areas, these are known as *kelurahan*. We do not distinguish between *desa* and *kelurahan* in the analysis, and they are treated identically in modern Indonesian data.

The variation that we exploit to identify the impacts of being a subjected village and being near a historical factory is extremely fine. Anything above the village level aggregates away the relevant variation, and villages are very small, with over 30,000 of them in Java. While a number of older data sources are available, they are too aggregated to be useful for the analysis. We can, however, trace some effects further back in time, using cohort analysis from more recent datasets.

4 Impacts of Sugar Factories

4.1 Empirical Strategy

4.1.1 Overview

In order to extract a surplus from the Javanese economy, the Dutch constructed a network of water-powered industrial sugar processing plants to transform raw cane into processed sugar that could be exported. Prior to the Cultivation System, rural Java was overwhelmingly focused on subsistence production, and hence this extensive network of factories represented a major reorganization of economic life. We might expect the factories to have long-run effects - through agglomeration, infrastructure investments, or other channels - even in areas where sugar processing has long since disappeared.

Estimating the effects of proximity to a historical factory requires identifying a plausible set of counterfactual factory locations, since various characteristics could vary with distance to areas suitable for factories. We exploit the fact that since sugar catchment areas were typically adjacent and each needed to produce enough cane for the factory to operate, factories had to be spaced several kilometers apart. A variety of locations along rivers between actual factories would have been suitable, but factories were not placed there due to the spacing constraints. This suggests that there were many possible equilibria for where factories could have been located (Salop, 1979).

Therefore, we conduct the following counterfactual exercise. For each historical factory, we take 1000 independent random draws that shift the factory anywhere from 5 to 20 kilometers to a suitable counterfactual location along the river. To be suitable, the point must have at least as much sugar suitable land within a 5 kilometer radius as the 10th percentile of the actual factory distribution. Sugar grows in the plains and not on steep hillsides or in mountainous areas. We infer suitability by observing where sugar was grown historically, first computing the 90th percentile of slope and elevation in sugar villages and then assigning anywhere with slope or elevation less than these cutoffs as suitable. We choose the 90th percentile because some sugar villages contain both plains and part of a mountain, leading

the right tail of the elevation and slope distributions to be highly skewed.¹² We also constrain placebo factories for each draw to be at least as far apart as the 10th percentile in the distribution of actual factories (this restriction makes little difference to the estimates).

This approach is illustrated in Figure 1. Panel a) shows an actual factory, surrounded by a 5 kilometer radius, with sugar suitability and rivers in the background. Panel b) shows a suitable counterfactual factory location, which has a sufficient amount of nearby sugar suitable land, as compared to the distribution of actual factories. Panel c) shows a larger set of counterfactual factories.

We then estimate the following specification, both for the actual factory and for the 1000 sets of counterfactual factory locations:

$$out_v = \alpha + \sum_{i=1}^{20} \gamma_i dfact_v^i + \beta \mathbf{X}_v + \sum_{j=1}^n fact_v^j + \epsilon_v \quad (1)$$

where out_v is an outcome of interest in village v , and the $dfact_v^i$ are indicators equal to one if village v is 0-1 km from the nearest (placebo) factory, 1-2 km, ..., 18-19 km. The omitted bin is 19-20 km, which is the maximum distance included in the sample. \mathbf{X}_v includes elevation, slope, distance to the coast, distance to the closest natural harbor, distance to the nearest river, log flow accumulation, and a spline in distance to the nearest 1830 residency capital, with kink points every three kilometers. The $fact_v^j$ are nearest factory fixed effects, which ensure that villages are compared to other villages near the same (placebo) factory.

The baseline specification limits the sample to sugar suitable villages, in order to ensure that the sample is geographically similar and because the analysis of the impacts of being a subjected sugar village will focus on this set of places.¹³ Our main set of results also restricts the sample to historical factories that are not in close proximity to a modern factory, to ensure that effects are not primarily driven by the persistence of sugar processing itself, and we discuss explicitly the few cases where this restriction matters (primarily: manufacturing of products used as inputs to sugar processing).

We use the position of the $dfact_v^i$ coefficients for the actual factory in the absolute value distribution of the 1000 counterfactual $dfact_v^i$ coefficients to compute a p-value. Small p-values imply that patterns near the actual factories would have been unlikely to arise in the factories' absence.

The study also conducts an additional counterfactual exercise, shifting all the historical factories up or down the river by the same distance. Then it estimates equation (1) for each

¹²An alternative would be to use FAO suitability data, but unfortunately these are too aggregated to be useful when exploiting village level variation.

¹³To be consistent with the analysis of subjected village, we also exclude places located in cities that were pre-period residency capitals. As documented below, effects do not change when these places are included.

of these common shifts. We would expect the γ_i to be largest for shifts around 0 km - which do not substantially change the location of the factories - and to dissipate the further the placebo factories are moved away from the actual factories.

4.1.2 Specification Robustness

The results are highly robust to a variety of alternative specifications. In the main text, we compute p-values using the placebo distributions constructed from independent random shifts, whereas Appendix A shows that patterns are broadly similar when we instead compute p-values using the placebo distributions constructed from the common shifts. In the baseline, the sample includes villages up to 20 kilometers from the nearest historical factory, and Appendix B documents that similar patterns obtain when only villages within 10 kilometers of the historical factories are included.

Moreover, the baseline - except where explicitly noted - does not include historical factories that are near modern sugar factories, as the focus is on effects that do not go directly through the persistence of sugar production. Appendix C shows that most conclusions do not change when these factories are included. The effects that emerge are also very similar when we include all villages in the sample, rather than limiting to sugar suitable villages, as documented in Appendix D. Finally, Appendix E shows that it makes little difference if we omit the factory fixed effects, comparing all places near historical factories to each other rather than only exploiting variation within places near the same factory.

4.1.3 Pre-Characteristic Balance

We begin by examining whether geographic characteristics are similar in places near versus far from actual factories, and in places in similar proximity to actual versus counterfactual factories. Geographic characteristics may well vary with proximity to actual factories, but we would expect places in similar proximity to actual versus counterfactual factories to look very similar.

We illustrate the patterns in the data by plotting the $dfact^i$ coefficients for proximity to the actual factories. We denote the significance of each of the $dfact^i$ coefficients relative to the counterfactual $dfact^i$ distributions constructed from randomly shifting the factories to plausible counterfactual locations. Crosses indicate coefficients that are above the 95th percentile of the counterfactual distributions, solid dots denote coefficients that are above the 90th percentile, and hollow dots indicate coefficients that are below the 90th percentile. An alternative to plotting the actual $dfact^i$ coefficients would be to plot the difference between the actual coefficients and the median placebo coefficients, which is done in Appendix F. This typically makes little difference, as the means of the placebo $dfact^i$ distributions tend

to be close to zero. The figures also highlight the general shape of the relationship by plotting a linear spline with kink points every 3km, which is fitted on the raw data controlling for nearest factory fixed effects.

Figure 4 documents that geographic characteristics vary similarly with distance to the actual and counterfactual factories, suggesting that the counterfactual locations were indeed similar to the actual ones prior to the construction of the factories. The characteristics considered are elevation, slope, distance to the coast, flow accumulation, distance to the nearest natural harbor, distance to the nearest river, and distance to the nearest 1830 residency capital. Flow accumulation is a measure constructed by the USGS Hydrosheds project that calculates how many cells are uphill from the cell under question. The higher the number, the more water we would expect to flow through the cell. While a few coefficients are significant, most likely by chance, there is no consistent pattern. Moreover, the magnitudes are very small – for example, moving from 1km away from a factory location to 10km away from a factory is associated with a difference in elevation of about 2 meters.

Places near historical factories are closer to 1830 residency capitals (panel g). There are a small number of these - only 24 in all of Java - and Appendix G shows that dropping factories within 10 kilometers of a residency capital flattens out this relationship, whereas the other results are broadly robust to this variation. Moreover, when common instead of random shifts are used to compute p-values (Appendix A), the *d_{factⁱ}* coefficients are less significant in comparison to the placebo distributions. Finally, in other robustness checks - such as dropping the fixed effects (Appendix E) or not requiring villages in the sample to be sugar suitable (Appendix D), distance to the nearest residency capital is also balanced relative to the placebos, whereas other results remain similar. Since distance to the nearest 1830 residency capital differs, we control for a flexible spline in it - with kink points every three kilometers - throughout the analysis.

4.2 Results

4.2.1 Economic Structure

To the extent that agglomeration economies are important and highly persistent, government sugar factories could plausibly influence industrialization in the long-run, even in places where sugar has not been processed for many decades. We examine this using a variety of data on sectoral employment. We focus here on places where there is not a modern sugar factory nearby, in order to isolate effects that do not go directly through the persistence of sugar processing. Subsequently, we will explicitly examine modern sugar production.

We begin by illustrating the methodology in more detail, using as a dependent variable an indicator for whether the individual works in the agricultural sector, taken from the

SUSENAS 2001-2011 household surveys.¹⁴ The sample includes prime age males aged 18 to 55 to avoid confounding labor market participation.¹⁵ As described in Section 4.1, for each factory we take 1,000 independent random draws that shift the factory to a suitable counterfactual location along the river. We estimate equation (1) using distance to the actual factories, as well as distance to each of these 1,000 sets of counterfactual factory locations.

The sub-plots in Figure 5, panel a) show the counterfactual distributions of absolute coefficients for each of the $d\text{fact}^i$ bins in distance to the nearest factory: 0-1 km, 1-2 km, etc. The coefficients that measure the impacts of proximity to *actual* factories are denoted by a red line. Each sub-plot also reports a p-value, given by the fraction of the absolute value counterfactual coefficients to the right of the absolute value of the actual coefficient. For bins near a factory, the actual coefficients fall far in the tails of the counterfactual distributions, indicating that patterns of agricultural employment near government factories would have been very unlikely to arise in the factories' absence. For bins further away, the actual estimates fall in the center of the counterfactual distributions, documenting that agricultural employment is not different from what we would have expected in the factories' absence. Effects dissipate within five kilometers.

Panel b) plots the $d\text{fact}^i$ coefficients for proximity to the actual factories. Crosses indicate coefficients that are above the 95th percentile of the counterfactual distributions shown in panel a), solid dots denote coefficients that are above the 90th percentile, and hollow dots indicate coefficients that are below the 90th percentile. These figures also plot a linear spline with kink points every 3 kilometers. The spline is fitted on the raw outcome data, controlling for geographic characteristics, nearest factory fixed effects, and survey year fixed effects.

Individuals within a few kilometers of a historical factory are 20 to 25 percentage points less likely to work in agriculture than those 10 to 20 kilometers away. Appendix Figure F-2 shows a version of this plot where the mean of each placebo $d\text{fact}^i$ distribution has been subtracted from each coefficient. Relative to the placebos, individuals in the immediate vicinity of a historical factory are around 17 percentage points less likely to work in agriculture.

We also conduct a placebo exercise that shifts all the historical factories up or down the river by the same distance and then estimates equation (1) for each of these common shifts. Each sub-plot in panel c) shows the γ_i coefficients for a $d\text{fact}^i$ bin in distance to the nearest factory: 0-1 km, 1-2 km, etc. The x-axis plots the magnitude of the common shift upstream (left side of the plot) or downstream (right side of the plot), in kilometers. The y-axis plots the γ_i , with the length of each bar indicating the magnitude of the counterfactual estimate. The patterns that emerge using this set of placebos are similar.

¹⁴When we use the SUSENAS datasets, for which we pool a number of different years of the data together in order to have sufficient number of observations in each village, we augment equation (1) to include a survey-year dummy.

¹⁵Estimates, available upon request, are similar using all individuals.

Panel d) is analogous to panel b), plotting the coefficients for proximity to the actual factories and computing p-values using the counterfactual distributions constructed with the specific shifts. Relative to the specific shifts counterfactuals, individuals in the immediate vicinity of a historical factory are around 20 percentage points less likely to work in agricultural (Appendix Figure F-2).

Figure 6 examines the industrial structure near government sugar factories in more detail, examining both different sectors and different time periods. Due to space constraints, only the coefficient plots are shown in the main text, with the statistical significance estimated using the counterfactual distributions from the independent random shifts. The complete set of counterfactual distributions for the independent and common shifts are shown in Appendices H and I, respectively. Coefficient plots with p values constructed using the common shifts are shown in Appendix A.

For comparison, panel a) repeats the plot for agriculture from the 2001-2011 household survey data. Panel b) likewise examines whether the individual works in agriculture, using data from the 1980 Population Census, again limiting the sample to prime age males. The patterns are similar to those from 2001-2011, but the effects are even larger than in the more recent period. Places in close proximity to a factory are 30 to 35 percentage points less agricultural than those just ten to twenty kilometers away, and effects are similar when the counterfactual mean is differenced out (Figure F-3). Indonesia was much more agricultural in 1980, and these results suggest some convergence with more recent industrialization.

Next, panels c) and d) consider whether the individual works in manufacturing, again using data from 2001-2011 SUSENAS and the 1980 Population Census. The pattern for manufacturing during 2001-2011 is the inverse of what we see for agriculture, with around six percentage points more individuals working in manufacturing. The 1980 Census likewise reveals that manufacturing employment is higher within a few kilometers of the historical factory. Places near a historical factory had around seven percentage points more individuals working in manufacturing, a very large effect given that Indonesia was a primarily agrarian economy at the time – only 11 percent of the population in 1980 was employed in manufacturing. The effects remain unchanged when the counterfactual means are differenced out (Figure F-3), further indicating that it is not simply the case that places near suitable factory locations are more industrial - i.e. because rivers are good for industrialization.

We also examine employment in retail. For 2001-2011, places in the immediate vicinity of a factory have around 10 percentage points more employment in retail relative to places ten to twenty kilometers away, and the coefficient for the closest bin is above the 90th percentile of the counterfactual distribution (panel e). In contrast to manufacturing, the retail effects are somewhat higher today than in 1980.

Historically the sugar factories were major centers of employment, which increasingly

relied on wage labor as time went by, and population would have plausibly clustered nearby. The occupational patterns – more manufacturing, more retail, and less agricultural – are also consistent with the areas near historical factories being more urban. Panels g and h examine log population density, using data from 2003 and 1980 (respectively), defined as the log of population in each village per square kilometer. Areas within 1km or so of a factory are much more densely populated than area 10 km away in both 1980 and today.

4.2.2 Sugar and Linked Industries

The above results suggest that the historical location of production impacted current economic activity through agglomeration, and we now examine this mechanism in more detail. Firms that used sugar as an input may have located near historical sugar processing plants - creating additional value through input-output, labor skill, and technological linkages (Ellison et al., 2010; Marshall, 1890). The locations of linked firms may have persisted even after the original sugar factories disappeared. Final goods consumers or workers may have also clustered near historical production centers.

Before examining linkages in detail, we consider whether sugar production itself has persisted. Figure 7, panel a) shows using the full sample of historical factories and data from the 2006 Economic Census that indeed, places near a historical factory - particularly those within 0 to 1 kilometers - produce substantially more processed sugar today than places further away. In contrast, panel b) documents that once we drop the 19% of historical factories within 2 kilometers of a modern sugar factory - as we do in all of our baseline results - the relationship between distance and modern sugar production flattens out.

Panel c) shows that there is little relationship in the full sample between raw cane production, drawn from the Podes 2003 village census, and distance to a historical factory, presumably because improvements in transport have obviated the need to grow cane in the immediate vicinity of processing plants. The relationship between modern sugar cultivation and distance to a historical factory remains flat in the restricted sample (panel d).

Input-output linkages were plausibly an important driver of agglomeration around sugar factories historically, since the low grade sugar that factories sold on local markets was costly to transport. To test whether such agglomeration has persisted, we construct weighted average employment shares for industries upstream and downstream from sugar processing. Sectoral employment is drawn from the 2006 Economic Census and total employment is from the 2000 Census. The weights are from the Leontief inverse of the 2006 Indonesian Input-Output Table. The input-output table specifies how many dollars of sector i 's output are needed by sector j to produce one dollar of its own output. The higher this number, the stronger the linkages between the sectors. The Leontief inverse of this matrix captures

not just direct linkages but also indirect ones. If sector k uses inputs from sector j , which in turn uses inputs from sector i , sector k is indirectly linked to sector i via sector j . The Leontief inverse measures how much of sector i 's output is used both directly and indirectly by sector k . The Leontief weights are described in more detail in Appendix Y.

The main manufacturing sectors upstream from sugar processing are farm machinery, used to harvest cane, and capital equipment, used to process cane. (Raw cane is, of course, the largest upstream sector but is not included in our measure since the data only include manufacturing establishments.) The main manufacturing industries downstream are in food processing, as sugar is an additive to most processed foods. Many types of services - restaurants, hotels, schools, and hospitals, to name a few - are also downstream from sugar processing, since sugar is consumed at these establishments, but these are not included in the downstream measure since they are not in manufacturing. These downstream linkages could though contribute to the retail impacts documented above.

Figure 7, panel e) shows that when all historical factories are included in the analysis, upstream industries - i.e. farm and processing machinery - are about three times as prevalent relative to the mean within one kilometer of historical factories. This relationship flattens out when we exclude historical factories that are located near modern ones (panel f).

In contrast, even when we limit the sample to historical factories that are not near modern ones, employment in manufacturing industries downstream from sugar is much higher near the historical factories (panel h). This is particularly true for places within 0 to 1 kilometer of a historical factory, and the effect is large. The difference in employment share in downstream industries between the 0-1 km bin and places 10 to 20 km away is about 50 percent greater than the sample mean. Sugar was used as an input in other foods historically, and other processed foods also tend to be used as inputs into each other. This suggests a particular channel for manufacturing persistence: even after the original sugar factories disappeared, there were still agglomeration advantages for the remaining downstream firms to continue to locate in the same place.

4.2.3 Public goods and the public sector

The analysis thus far has focused on the private sector, but there may be important linkages that occurred through the location of public investments. The historical literature emphasizes that the Dutch government constructed road and rail networks to transport sugar to ports. The configuration of transport infrastructure tends to be very persistent, and hence these impacts may have well persisted through the present. In the analysis of infrastructure - and for the remainder of the study - we drop historical factories near modern sugar factories to isolate impacts that do not go directly through sugar production, with Appendix C

documenting similar effects when these places are included.

Figure 8, panels a) and b) document that intercity and local road density, respectively, are higher today in the immediate vicinity of factories, and these effects are atypical of those near counterfactual factories. When comparing the 0-1 km bin to places 5-20 km away, the difference in intercity road density is about twice as large as the sample mean. Differences in local road density are similarly large. Moreover, data from the 1980 Podes, a census of village governments, show that villages in the immediate vicinity of the historical factories were less likely to only be accessible via a dirt road historically (panel c). Magnitudes remain similar when the counterfactual means are differenced out (Figure F-5).

Railroad density - most of which was constructed during the Dutch period - is also substantially higher in the immediate vicinity of the historical factories, and the effect is non-zero only for the 0-1 and 1-2 kilometer bins, as we would expect given that rail infrastructure was initially built in Java explicitly to transport processed sugar to ports. The Dutch made these large and costly infrastructure investments in order to extract a surplus.

The 1980 Podes data also reveal that places within a few kilometers of the factories were substantially more likely to have electricity, and the effects fall far in the right tails of the counterfactual distributions (panel e). In the immediate vicinity of the historical factories, villages were 50 percentage points more likely to have electricity than places ten to twenty kilometers away. In 1980 relatively few households had electricity, and hence higher rates of electrification likely indicate a greater presence of manufacturing.

The historical literature suggests that the main public goods that the Cultivation System directly impacted were road and rail infrastructure, but over time the more industrialized and integrated areas near historical factories may have gained better access to public goods more generally. We find that in 1980, areas in the immediate vicinity of a factory were four percentage points more likely to have a high school than places ten to twenty kilometers away, relative to a sample mean of only two percent of villages with high schools (panel g). The effects on high schools are also positive using the pooled 1996-2011 Podes Village Censuses but are not statistically significant relative to the placebo distributions, and high schools had become more prevalent by this period (panel i).

High schools were mostly built in administrative centers, and indeed places near factories are around two kilometers closer to the nearest subdistrict capital, relative to a sample mean distance of 3.8 kilometers (panel h). Note that differences in proximity to subdistrict capitals remain statistically significant and of nearly identical magnitude when we limit the sample to be within 10 km of a historical factory (Appendix B), when we include all villages in the sample (Appendix D), and when we omit factory fixed effects (Appendix E), all specifications in which distance to the nearest 1830 residency capital is balanced. The effect is not large in absolute terms, given that everything in densely populated Java is in fairly close proximity,

but we will show that there are important differences in access to education.

Figure 9 examines long-run impacts on schooling in more detail, using cohort level data from the 2000 Population Census. We focus in the main text on three representative cohorts, with the remainder shown in the appendix: the 1920-1929 cohort (educated during the Dutch period), the 1950-1954 cohort (educated following independence), and the 1970-1974 cohort (educated during Indonesia's large-scale school building campaign). Appendix J shows results for the complete set of cohorts. Panel a) documents that individuals within a few kilometers of a historical factory on average have between 0.75 and 1.25 more years of schooling than those located ten to twenty kilometers away, relative to a sample mean of 4.9 years of schooling. These effects are atypical relative to the counterfactual distributions and hold across all three cohorts (panel b). We see a similar pattern for primary completion (panels c and d). It is much higher in close proximity to the historical factories, and this is particularly true for the two older cohorts, whose schooling occurred at a time when primary access was far from universal. High school completion again shows a similar pattern (panels e and f). Effects on high school are largest for the younger cohorts, who received schooling at a time when high school was more common, but impacts are still positive and statistically significant even for the oldest cohort, educated during the Dutch era.

4.2.4 Household Consumption

We conclude by examining the effects on household consumption, using pooled data from the SUSENAS household survey. Following Deaton (1997), we assume that children aged 0 to 4 are equal to 0.4 adults and children aged 5 to 14 are equal to 0.5 adults. All regressions control for survey year fixed effects and the number of household members aged 0-4, 5-14, and 15 and older. We find that consumption levels in areas immediately adjacent to the historical factories are around 14 percent higher than areas even just 5 kilometers further away. When the mean of the counterfactual distribution is differenced out, the difference remains at around 10 percent (Figure F-7). Consistent with these being relatively integrated areas, this effect is about what we would expect given the differences in education. People living adjacent to a historical factory have about 1.25 more years of education, and an 8-10 percent return to schooling (Duflo, 2001) would yield the observed consumption differences.

4.3 Discussion

In summary, the Dutch colonial sugar industry substantially transformed economic activity in contemporary Java. Sectoral composition suggests that industries that used sugar as an input sprung up surrounding colonial sugar factories, creating food processing centers with population clustered nearby. Even when the original source of sugar disappeared, those

industrial centers persisted, sourcing sugar from elsewhere. Of course, industrial production centers may have arisen in rural Java in the absence of the Cultivation System. However, given that rural Java in 1830 was a heavily agricultural economy governed by traditional agricultural labor norms in which landless peasants were customarily attached to landowning households and most production was for subsistence, without outside intervention such an industrial transition would have likely occurred much more slowly.

Moreover, infrastructure had to be constructed to transport Cultivation System sugar to ports, and we show that there are still impacts on public infrastructure today. There is broad consensus that transport infrastructure is important for economic development. Given that maximizing colonial profits was a central part of the Dutch objective, it is hard to imagine a scenario where they would have made these investments in the absence of a large extractable surplus. These results highlight how colonial extraction can have a range of effects, some of which may encourage agglomeration and infrastructure investments that promote economic activity in the long run.

5 Impacts on Subjected Villages

5.1 Empirical Strategy

5.1.1 Overview

The discussion thus far has focused on one aspect of the Cultivation System, namely the establishment of sugar factories. The second part of our analysis turns to the impact on the villages that were forced to grow sugar cane. To estimate the effects of forced cultivation, we exploit the discontinuous change in exposure at the borders of the subjected catchment areas (as illustrated by Figure 2). Inside, villages cultivated sugar for the government, whereas outside they did not. The boundaries form a multi-dimensional discontinuity in longitude-latitude space, and regressions take the form:

$$out_v = \alpha + \gamma cultivation_v + f(\text{geographic location}_v) + g(dfact_v) + \beta \mathbf{X}_v + \sum_{i=1}^n seg_v^i + \epsilon_v \quad (2)$$

where out_v is an outcome in village v . $cultivation_v$ is an indicator equal to 1 if the village grew cane for the Cultivation System and equal to zero otherwise. $f(\text{geographic location}_v)$ is the RD polynomial, which controls for smooth functions of geographic location. Following Gelman and Imbens (2014), we use a local linear RD polynomial for the baseline.¹⁶ The RD

¹⁶Regressions use a triangular kernel such that the weight given to each observation decays with distance from the threshold.

polynomial is estimated separately for each of the 17 contiguous catchment areas.

To ensure that the effects we estimate are due to a village being subjected to growing cane, rather than merely being close to a factory, $g(dfact_v)$ controls for a linear spline in distance to the nearest historical sugar factory, with kink points estimated every three kilometers; Section 4 shows that this functional form captures the impacts of proximity to a factory well. \mathbf{X}_v contains exogenous geographic characteristics: elevation, slope, distance to the coast, distance to the closest natural harbor, distance to the nearest river, log flow accumulation, and a spline in distance to the nearest 1830 residency capital.

The seg_v^i split each catchment area boundary into 10 kilometer segments, equaling one if village v is closest to segment i and zero otherwise. They ensure that the specification is comparing nearby villages. For regressions examining household consumption, we also include a vector of demographic variables giving the number of infants, children, and adults in the household. The baseline specification limits the sample to villages within 10 kilometers of the threshold. Standard errors are clustered at the sub-district level.

Villages where sugar could not grow are not a suitable counterfactual for areas with government cultivation. Hence, we limit the sample to areas that are suitable for sugar, where suitability is defined using the same elevation and slope requirements described in Section 4. For the same reason, we exclude places in cities that were residency capitals.

RD plots are shown in Appendix K. Legibly showing multi-dimensional RD graphs where latitude and longitude are the running variables, as in Dell (2010), would require zooming in separately for each of the 17 contiguous catchment areas. This results in a very large number of plots when repeated for all the study's outcomes. In the interest of concision, we instead show single dimensional RD graphs where the running variable is distance to the nearest catchment boundary. Negative values are used for places outside the catchment areas. Border segment fixed effects, geographic controls, a spline in distance to the nearest sugar factory, and a linear polynomial in latitude and longitude - estimated separately by catchment area - have been partialled out.

5.1.2 Specification Robustness

There are a variety of specification choices inherent in an RD design, and we document that results are very robust to alternative choices along the following dimensions: 1) RD bandwidth, 2) RD polynomial, 3) inclusion of additional controls, 4) sample inclusion criteria, and 5) additional features of the specification.

We are not aware of a well-accepted multi-dimensional RD optimal bandwidth, so Appendix L examines robustness to the full range of feasible alternative bandwidths. Each plot shows point estimates and confidence intervals of γ using equation (2) and different band-

widths between 2 and 10 kilometers, with the bandwidth under consideration denoted on the x-axis and the point estimate on the y-axis. Since the catchment areas are small, treated villages are never more than 10 kilometers from the nearest catchment boundary. Results are broadly robust to the choice of bandwidth, with narrower bandwidths by construction tending to lead to somewhat noisier estimates.

Next, we consider robustness to the specification of the RD polynomial. Appendix M shows that results change little when a single multi-dimensional RD polynomial is used - rather than estimating the polynomial separately by catchment area - and Appendix N documents that estimates are also very similar when the RD polynomial is specified as quadratic. Appendix O shows that results are broadly similar when a linear polynomial in distance to the boundary is included instead of the multi-dimensional RD polynomial.

We also examine robustness to varying additional control variables. Appendix P shows that results are robust to changing the length of the boundary segment fixed effects to 25 kilometers, and Appendix Q documents robustness to excluding the geographic controls.

In the baseline we limit the sample to villages suitable for sugar cultivation - since non-suitable places are not a plausible counterfactual - but estimates are robust to varying the sample inclusion criteria. An alternative would be to limit to boundary segments that are similarly suitable for sugar cultivation on either side of the boundary segment. Appendix R shows that estimates are broadly robust to this alternative sample restriction. Moreover, Appendix S shows that estimates are very similar when all villages are included in the sample, regardless of their sugar suitability.

Finally, we document robustness to some additional aspects of the specification. The baseline specification uses a triangular kernel such that the weight given to each observation decays with distance from the threshold, and Appendix T shows that estimates change little when no weighting is used. Finally, recall that not all villages within the constructed catchment zones can be matched to historical villages, since some villages have disappeared or changed their names. Appendix U designates matched villages as treated and instruments these with being inside the catchment areas. The first stage F-stat is over 1300, and estimates remain quite similar.

5.1.3 Pre-characteristic Balance

The key regression discontinuity identifying assumption is that all relevant factors besides treatment vary smoothly at the Cultivation System boundaries. That is, letting c_1 and c_0 denote potential outcomes under treatment and control, x denote longitude, and y denote latitude, identification requires that $E[c_1|x, y]$ and $E[c_0|x, y]$ are continuous at the discontinuity thresholds. This assumption is needed for observations located just outside the catchment

areas to be an appropriate counterfactual for observations located just inside.

To assess the plausibility of this assumption, Table 1 examines a variety of geographic characteristics, using gridded geographic data and regressions of the form described in equation (2). The unit of observation is the village (Appendix Table V-1 shows that results change little when the unit of analysis is a grid cell centroid).

Column (1) examines elevation. The point estimate on *cultivation* is negative and statistically significant but is only 2 meters, a very small difference. Indeed, the entire sample is in the plains, close to sea level (the mean elevation in the sample is 31 meters). In some of the alternative specifications examined in the appendices, this difference is smaller and not statistically significant, yet other results remain similar. None of the remaining characteristics – slope; flow accumulation; a dummy for being on the coast; or distance to the coast, nearest river, nearest natural harbor, and nearest 1830 residency capital – show statistically significant or economically meaningful differences.¹⁷

Another RD identifying assumption is that individuals did not selectively sort around the threshold while the Cultivation System was in force, in order to exploit its design. Typically in this context, one would worry that a substantial number of productive individuals moved just outside the subjected areas to escape forced cultivation. However, as will be shown below, we find positive economic impacts of the Cultivation System on subjected villages, and it appears implausible that high productivity individuals would have moved to regions subjected to forced cultivation. In historical Java, individuals who migrated to an already established village were not eligible to hold land, and disempowered movers would have plausibly borne the brunt of forced labor.¹⁸ Alternatively, low productivity individuals may have fled subjected villages.¹⁹ However, population density today is if anything greater in treated villages, suggesting that mass out-migration is unlikely to drive results.

5.2 Results

5.2.1 Land

The historical literature has largely focused on how the Cultivation System impacted land tenure in subjected villages. In particular, village heads under the Cultivation System were empowered by the Dutch to distribute land in order to facilitate the growing of sugar cane. As discussed in Section 2, the literature emphasizes that land was redistributed to the village,

¹⁷We do not examine characteristics such as weather or soil quality since the lower resolution of these data mean treated and non-treated villages typically fall in the same cell.

¹⁸While in theory landowners were responsible for cultivation services, in practice landless peasants attached themselves to landholding families and often performed the labor.

¹⁹Elson (1984, p. 60) indeed argues that the growth of frontier regions of interior Java during the Cultivation System may have been due to individuals fleeing forced labor, though it is not obvious that the lowest productivity individuals would have been the ones who moved.

whether the motive was to directly benefit officeholders, to facilitate sharing the burdens of the System, or both (Elson, 1984, p. 94; Ricklefs, 2008, p. 159). Village land is used for multiple purposes, but an important one is as *tanah bengkok*, which is land owned by the village that the village head is allowed to use (or rent out) as compensation for his service.

We therefore begin by examining the allocation of land in more contemporary periods, and in particular focusing on village-owned land, as well as overall land inequality among privately held land. Indonesian PODES village censuses from 2003 and 1980 collected information on village-owned land, and land inequality for private agricultural land can be measured using the 100% microdata sample from the 2006 Agricultural Census.

The estimates show that the Cultivation System left a substantial mark on the amount of village-owned land. Table 2, columns (1) through (4) show that in both 2003 and 1980, Cultivation System villages had substantially more village-owned land, in absolute terms (columns 1 and 3) and as a percentage of total land (columns 2 and 4). The differences are highly statistically significant. In 2003, about 1.4 percentage points more land was owned by the village in Cultivation System areas, relative to a sample mean of 9% of total land being owned by villages. In 1980 these villages owned 1.2% more land, relative to a sample mean of 11%. RD plots are shown in Appendix Figure K-2.

Columns (5) through (8) examine inequality amongst agricultural households, considering the 99-90 ratio, the 90-10 ratio, the 90-50 ratio, and the 50-10 ratio of agricultural land usage.²⁰ Coefficients are positive but fall short of statistical significance. Across various specifications examined in the appendix, the coefficients tend to be positive and either marginally or not quite significant, suggesting land inequality is if anything higher in treated villages, again consistent with the historical impacts emphasized in the qualitative literature.

5.2.2 Public goods

The village's ability to raise revenue, as well as land tenure arrangements more generally, could impact human capital accumulation and the provision of local schooling (Dell et al., 2017; Acemoglu et al., 2009; Banerjee and Somanathan, 2007). This may have been particularly true in the Dutch period and in independent Indonesia prior to the massive INPRES school expansion, studied by Duflo (2001), as prior to INPRES village-level school construction was financed locally by the villages themselves (Aritonang, 1994).

We begin by looking at the presence of schools in 1980, the earliest date for which systematic village level data on different types of schools are available. Table 3 examines the presence of schools using data from the Podes 1980 village census. 1980 was around the time of Indonesia's massive primary school building campaign - INPRES - which largely equalized

²⁰Ownership data do not distinguish between agricultural and non-agricultural land, and hence we unfortunately cannot examine inequality in land ownership.

access to basic primary schooling. The data break down primary schools into those that were built by the INPRES campaign and those that were not. Columns (1) and (2) focus on non-INPRES schools. There is not an impact on the number of buildings, whereas the effect on the number of teachers is if anything positive but falls just short of statistical significance. In the various appendix specifications, these coefficients are sometimes marginally significant.

Columns (3) and (4) show that cultivation villages received significantly fewer INPRES schools and teachers, almost surely an indicator of greater education beforehand since INPRES was targeted based on the availability of pre-existing primary education (Duflo, 2001). Subjected villages were around four percentage points less likely to receive an INPRES school, relative to a sample mean of 0.36. Moreover, subjected villages were 2 percentage points more likely to have a junior high school, relative to a sample mean of 6 percent of villages that had them (junior high schools were not covered by INPRES). There is not a statistically significant impact on the presence of a high school, which only two percent of villages had.

These impacts are consistent with the fact that during the Dutch period, if communities wanted a school they were responsible for raising a share of the required resources. In particular, under the *volkschool* (village school) program, the colonial government would provide resources towards hiring a teacher only if the community constructed and maintained the school (Aritonang, 1994).

We then look at whether there are persistent effects on educational attainment, using the 100% sample of the 2000 Population Census. Table 4, column (1) documents that individuals in subjected villages have around 0.22 years more schooling, relative to a sample mean of 5 years, and the effect is statistically significant at the 1% level. They are also more likely to complete primary school and junior high (columns 2 and 3). There is not a statistically significant impact on high school completion, though only 13% of the sample completed high school (column 4).

Columns (5) and (6) examine the 1980 Population Census. Individuals in subjected villages are 1.7 percentage points less likely to have no schooling, relative to a sample mean of 41%. There is not an effect on primary completion, with only 19% of the sample in 1980 completing primary school.

To examine impacts on education even further back in time, we estimate effects by cohort. For the 1980 Census, data are only available for a relatively small sample, and hence cohort estimates are quite noisy. In contrast, the 2000 Census provides a 100% sample, so we can estimate village-level impacts on each cohort, beginning with the cohort born between 1920 and 1930 and continuing through the cohort born between 1975 and 1980, the youngest cohort to have reached adulthood by 2000.²¹ The left panel in Figure 11 plots estimates from equation (2) by cohort. Each point represents a γ coefficient from a separate estimation of

²¹We use ten year age groupings to increase power for the oldest cohorts, since these are much smaller.

equation (2) for a given cohort, with the 95% confidence interval indicated by a line. For ease of interpretation, the right panel plots sample means for each cohort and outcome variable. In general, schooling levels were initially very low and increased over time.

The left panel of Figure 11 shows that impacts on years of schooling are large and positive across cohorts, even going as far back as the cohort born in the 1920s who would have completed their primary education under the Dutch. Effects on primary completion peak for cohorts born when primary was rapidly expanding - in particular for cohorts born in the 1950s and 60s. The impact decreases somewhat as primary completion becomes more universal, though is still present in the most recent cohorts. On the other hand, impacts on junior high completion are zero for the older cohorts, whose junior high completion rates were close to zero, and then become large and positive as cohorts born from the 1950s onward begin to complete junior high at more appreciable rates. Impacts on high school completion show a similar pattern.

Finally, columns (7) and (8) of Table 4 consider the education of village heads, pooling data from Podes village censuses collected between 1996 and 2011. The regressions include survey year fixed effects. While there is not a statistically significant impact on years of schooling, village heads in subjected areas are three percentage points more likely to have completed high school, relative to a sample mean of 74%. The effect on village heads is about double the average effect for cohorts of similar age shown in Figure 11, panel g). One possibility is that the more generous compensation for village heads, in the form of more village land they are allowed to use *ex-officio*, attracted a relatively more educated pool of village heads. Empirical evidence from Indonesia in turn shows that more educated village heads increase public goods provision (Martinez-Bravo, 2017). Appendix Table W-1 considers various public goods. Effects of forced sugar cultivation tend to be positive but not statistically significant.

5.2.3 Economic structure

We finally turn to the overall structure of the economy. Table 5, columns (1) through (3) examine data from the SUSENAS household survey, collected in annual waves from 2001 through 2011. All regressions include survey year fixed effects.

We find that villages subjected to the Cultivation System are substantially *less* agricultural today. Specifically, in subjected villages, individuals are 4 percentage points (15 percent) less likely to work in agriculture, 3 percentage points (14 percent) more likely to work in manufacturing, and 1.3 percentage points (7 percent) more likely to work in retail. The sample is limited to prime aged males - aged 18 to 55 - to avoid confounding impacts with labor force participation, but estimates are similar when the entire labor force is included.

Columns (4) through (6) examine the structure of the economy in 1980, using data from the 1980 Population Census and again focusing on prime age males.²² The coefficient on employment in agriculture is similar in magnitude to that from more modern SUSENAS data. The effect on manufacturing, which was a much smaller share of the Indonesian economy in 1980, is close to zero and statistically insignificant. However, we cannot rule out a similar proportionate increase in manufacturing to that in the more recent data. Individuals in subjected villages were 2.6 percentage points (22 percent) more likely to be employed in commerce.

Data from Indonesia’s 2006 Economic Census, examined in Table 6, show a similar pattern. The Economic Census captures all large manufacturing firms and a random sample of small and informal firms. We find that subjected villages have 30 percent more manufacturing firms than non-subjected villages (column 1). Column (2) considers the village’s manufacturing employment, taken from the Economic Census, divided by the size of the labor force, drawn from the 2000 Population Census. The estimate, which indicates a 2.3 percentage point (14 percent) increase, is similar to what we obtain from the SUSENAS household surveys, though more noisily estimated (not surprisingly since it is drawn from multiple datasets). Population density is also around seven percent higher in subjected villages, both in 2003 and in 1980.

Finally, column (5) examines equivalent consumption from the SUSENAS household survey. There is not a statistically significant difference in household consumption. However, the confidence intervals include the possibility of consumption being around 2 percent higher, about what we would expect given that individuals in subjected villages on average have around 0.2 years additional schooling.

5.3 Discussion

These estimates show that if anything, the persistent local development effects of forced cultivation reinforce the positive long-run economic impacts of creating a sugar processing infrastructure. This is surprising at first glance in the context of other studies that find negative development effects of colonial forced labor (Lowe and Montero, 2016; Dell, 2010), but there are various clues about what drives the impacts on cultivation villages.

The historical literature emphasizes that while in force, the Cultivation System led to increases in communal village landholdings and the share of village lands apportioned to village officeholders (Elson, 1984, p. 94; Ricklefs, 2008, p. 159). We find that still today there is about 10 to 15 percent more public-use land in Cultivation System villages. This has

²²We do not consider sectoral information from the 2000 Census, as it was collected in the midst of a severe economic downturn – the Asian Financial Crisis – and hence is less informative about long-run economic conditions.

remained fairly constant over time, with similar effects in both 1980 and 2003. Households in Cultivation System villages are also more educated going all the way back to the 1920s cohort, which received its education during the Dutch period. While there are a variety of channels that could link the Cultivation System to long-run human capital accumulation, the presence of village land is a particularly plausible mechanism. Historically, villages that wanted a school needed to fund the school construction and upkeep themselves, and the revenue from village-owned lands may have facilitated this construction.

We also find that cultivation villages continue to have a different economic structure, with a greater percentage of households working in manufacturing and retail and fewer in agriculture, both in 1980 and today. Higher human capital levels - and the disproportionately higher schooling levels of the village head - could plausibly lead to less agriculture and more manufacturing. Moreover, the historical literature emphasizes that the influx of crop incentive payments into Cultivation System villages incentivized modern production and exchange (Elson, 1994, p. 261, Ricklefs, 2008, p. 158)

6 Conclusion

The establishment of a sugar processing infrastructure in colonial Java persistently increased industrialization, education, and household consumption in areas near government sugar factories, even after the factories themselves had disappeared. Similarly, villages forced to grow sugar cane for the Cultivation System have more schooling and manufacturing today. Infrastructure, agglomeration via input-output linkages, and public goods provision are particularly plausible channels of persistence.

Extractive colonial institutions can have a myriad of impacts, including on agglomeration, infrastructure, and the organization of economic and political relationships within the village. In the case of the Cultivation System, the positive impacts on economic activity plausibly dominated in the long-run because of a constellation of factors: processing had to be done on site, and a modern transport infrastructure was built to connect the extensive network of sugar factories to international ports. Processed sugar, some of which was sold on local markets, had dense linkages to industries that remained important after Indonesian independence. In subjected villages, the higher share of village land plausibly raised the village's capacity to fund schools, as education became increasingly important over time.

The Dutch were not the only colonial power to reorganize colonial economies in order to create the surplus that they wished to extract, nor are the Javanese unique in creatively utilizing colonial structures to generate economic activity long after the colonizers had left. These results plausibly inform a variety of contexts in which foreign powers attempted to extract surplus by reorganizing the economy along more modern lines.

References

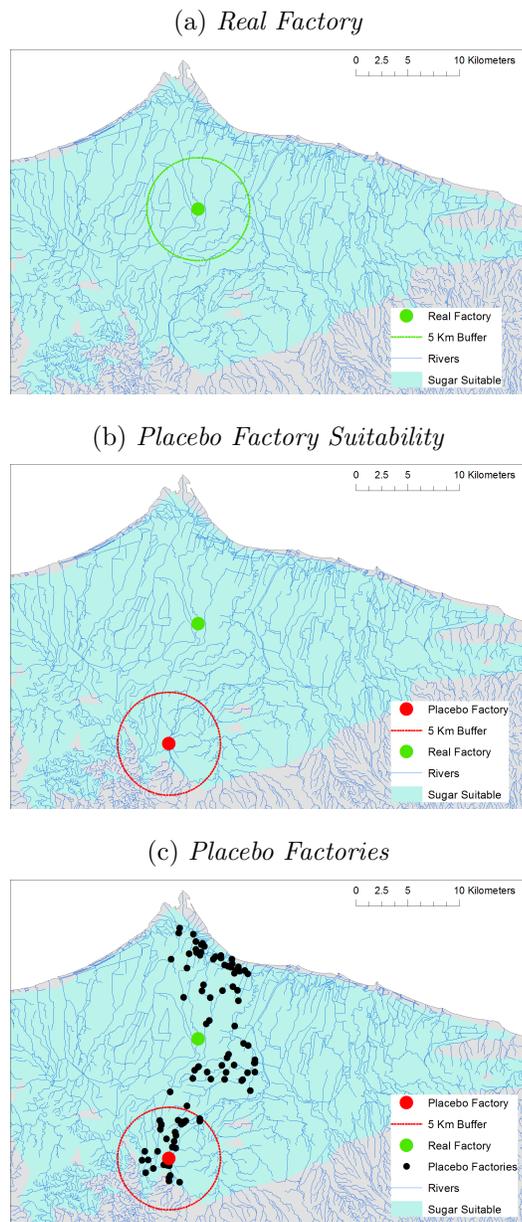
- ACEMOGLU, D., U. AKCIGIT, AND W. KERR (2016): “Networks and the Macroeconomy: An Empirical Exploration,” *NBER Macroeconomics Annual*, 30, 273–335.
- ACEMOGLU, D., M. A. BAUTISTA, P. QUERUBIN, AND J. A. ROBINSON (2009): “Economic and Political Inequality in Development: The Case of Cundinamarca, Colombia,” in *Institutions and Economic Performance*, ed. by E. Helpman, Harvard University Press.
- ACEMOGLU, D., D. CANTONI, S. JOHNSON, AND J. A. ROBINSON (2011): “The Consequences of Radical Reform: The French Revolution,” *American Economic Review*, 101, 3286–3307.
- ACEMOGLU, D., V. M. CARVALHO, A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2012): “The Network Origins of Aggregate Fluctuations,” *Econometrica*, 80, 1977–2016.
- ACEMOGLU, D., C. GARCIA-JIMENO, AND J. A. ROBINSON (2015): “State Capacity and Economic Development: A Network Approach,” *American Economic Review*, 105, 2364–2409.
- ACEMOGLU, D., S. JOHNSON, AND J. A. ROBINSON (2001): “The Colonial Origins of Comparative Development: An Empirical Investigation,” *American Economic Review*, 91, 1369–1401.
- (2002): “Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution,” *The Quarterly Journal of Economics*, 117, 1231–1294.
- ACEMOGLU, D. AND J. A. ROBINSON (2013): *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*, Crown Business.
- ALESINA, A., P. GIULIANO, AND N. NUNN (2013): “On the Origins of Gender Roles: Women and the Plough,” *The Quarterly Journal of Economics*, 128, 469–530.
- ARITONANG, J. S. (1994): *Mission Schools in Batakland (Indonesia): 1861-1940*, vol. 10, Brill.
- BANERJEE, A., E. DUFLO, AND N. QIAN (2012): “On the Road: Access to Transportation Infrastructure and Economic Growth in China,” Tech. rep., National Bureau of Economic Research.
- BANERJEE, A. AND L. IYER (2005): “History, Institutions, and Economic Performance: The Legacy of Colonial Land Tenure Systems in India,” *American Economic Review*, 95, 1190–1213.
- BANERJEE, A. AND R. SOMANATHAN (2007): “The Political Economy of Public Goods: Some Evidence from India,” *Journal of Development Economics*, 82, 287–314.
- BAQAEE, D. (2015): “Labor Intensity in an Interconnected Economy,” Working paper, Harvard.
- BAUM-SNOW, N. (2007): “Did Highways Cause Suburbanization?” *The Quarterly Journal of Economics*, 122, 775–805.
- BECKER, S. O., K. BOECKH, C. HAINZ, AND L. WOESSMANN (2016): “The Empire Is Dead, Long Live the Empire! Long-Run Persistence of Trust and Corruption in the Bureaucracy,” *The Economic Journal*, 126, 40–74.
- BUKOWSKI, P. (2016): “How History Matters for Student Performance. Lessons from the Partitions of Poland,” Working paper, Central European University, http://www.personal.ceu.hu/staff/repec/pdf/2015_3.pdf.
- CARVALHO, V. M. (2009): “Aggregate Fluctuations and the Network Structure of Intersectoral Trade,” Working paper, Pompeu Fabra University.
- COMMISSIE UMBGROVE (1858): “Archief van de Commissie voor de Opname van de Verschillende Suikerfabrieken op Java [1853–1858],” Nationaal Archief, The Hague, Access Number 2.10.11.

- DEATON, A. (1997): *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*, World Bank.
- DELL, M. (2010): “The Persistent Effects of Peru’s Mining Mita,” *Econometrica*, 78, 1863–1903.
- DELL, M., N. LANE, AND P. QUERUBIN (2017): “The Historical State, Local Collective Action, and Economic Development in Vietnam,” Working paper, Harvard.
- DIPPEL, C., A. GREIF, AND D. TREFLER (2015): “The Rents from Trade and Coercive Institutions: Removing the Sugar Coating,” Tech. rep., National Bureau of Economic Research.
- DONALDSON, D. (2010): “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure,” Tech. rep., National Bureau of Economic Research.
- DONALDSON, D. AND R. HORNBECK (2016): “Railroads and American Economic Growth: A Market Access Approach,” *The Quarterly Journal of Economics*, qjw002.
- DUFLO, E. (2001): “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment,” *American Economic Review*, 91, 795–813.
- ELLISON, G., E. L. GLAESER, AND W. R. KERR (2010): “What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns,” *American Economic Review*, 100, 1195–1213.
- ELSON, R. E. (1984): *Javanese Peasants and the Colonial Sugar Industry: Impact and Change in an East Java Residency, 1830–1940*, Oxford University Press.
- (1994): *Village Java Under the Cultivation System, 1830–1870*, Asian Studies Association of Australia.
- FASSEUR, C. (1992): *The Politics of Colonial Exploitation: Java, the Dutch, and the Cultivation System*, Cornell University Southeast Asia Program.
- FERNÁNDEZ, R. AND A. FOGLI (2009): “Culture: An Empirical Investigation of Beliefs, Work, and Fertility,” *American Economic Journal: Macroeconomics*, 1, 146–177.
- FOGEL, R. W. (1964): *Railroads and American Economic Growth*, Books on Demand.
- FOSTER, G. M. (1965): “Peasant Society and the Image of Limited Good,” *American Anthropologist*, 67, 293–315.
- GEERTZ, C. (1963): *Agricultural Involution: the Process of Ecological Change in Indonesia*, vol. 11, University of California Press.
- GELMAN, A. AND G. IMBENS (2014): “Why High-Order Polynomials Should Not be Used in Regression Discontinuity Designs,” Tech. rep., National Bureau of Economic Research.
- GENNAIOLI, N., R. LA PORTA, F. LOPEZ-DE SILANES, AND A. SHLEIFER (2013): “Human Capital and Regional Development,” *The Quarterly Journal of Economics*, 128, 105–164.
- GENNAIOLI, N. AND I. RAINER (2007): “The Modern Impact of Precolonial Centralization in Africa,” *Journal of Economic Growth*, 12, 185–234.
- GIULIANO, P. (2007): “Living Arrangements in Western Europe: Does Cultural Origin Matter?” *Journal of the European Economic Association*, 5, 927–952.
- GLAESER, E. L. AND J. D. GOTTLIEB (2009): “The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States,” *Journal of Economic Literature*, 47, 983–1028.

- GLAESER, E. L. AND A. SHLEIFER (2002): “Legal Origins,” *The Quarterly Journal of Economics*, 117, 1193–1229.
- GROSJEAN, P. (2011): “The Weight of History on European Cultural Integration: A Gravity Approach,” *American Economic Review*, 101, 504–508.
- (2014): “A History of Violence: The Culture of Honor and Homicide in the US South,” *Journal of the European Economic Association*, 12, 1285–1316.
- GUIO, L., P. SAPIENZA, AND L. ZINGALES (2008): “Alfred Marshall Lecture: Social Capital as Good Culture,” *Journal of the European Economic Association*, 6, 295–320.
- (2016): “Long-Term Persistence,” *Journal of the European Economic Association*, 14, 1401–1436.
- HAYAMI, Y. (1980): “Economic Approach to Village Community and Institutions,” *Journal of Rural Development*, 3, 27–49.
- HIRSCHMAN, A. O. (1960): *The Strategy of Economic Development*, Yale University Press.
- JEDWAB, R., E. KERBY, AND A. MORADI (2015): “History, Path Dependence and Development: Evidence from Colonial Railroads, Settlers and Cities in Kenya,” *The Economic Journal*.
- JEDWAB, R. AND A. MORADI (2016): “The Permanent Effects of Transportation Revolutions in Poor Countries: Evidence from Africa,” *Review of economics and statistics*, 98, 268–284.
- LANE, N. (2017): “Manufacturing Revolutions: Industrial Policy and Networks in South Korea,” Tech. rep., IIES.
- LIU, E. (2017): “Industrial Policies and Economic Development,” Tech. rep., MIT.
- LOWES, S. AND E. MONTERO (2016): “Blood Rubber: The Effects of Labor Coercion on Institutions and Culture in the DRC,” Working paper, Harvard, <http://www.saralowes.com/research.html>.
- LOWES, S., N. NUNN, J. A. ROBINSON, AND J. WEIGEL (2015): “The Evolution of Culture and Institutions: Evidence from the Kuba Kingdom,” Working Paper 21798, National Bureau of Economic Research.
- LUITEN VAN ZANDEN, J. (2010): “Colonial State Formation and Patterns of Economic Development in Java, 1800–1913,” *Economic History of Developing Regions*, 25, 155–176.
- LUTTMER, E. F. P. AND M. SINGHAL (2011): “Culture, Context, and the Taste for Redistribution,” *American Economic Journal: Economic Policy*, 3, 157–79.
- MARSHALL, A. (1890): *Principles of Economics*, Macmillan and Co.
- MARTINEZ-BRAVO, M. (2017): “The Local Political Economy Effects of School Construction in Indonesia,” *American Economic Journal: Applied Economics*, 9, 256–289.
- MICHAELS, G. (2008): “The Effect of Trade on the Demand for Skill: Evidence from the Interstate Highway System,” *The Review of Economics and Statistics*, 90, 683–701.
- MICHALOPOULOS, S. AND E. PAPAIOANNOU (2013): “Pre-Colonial Ethnic Institutions and Contemporary African Development,” *Econometrica*, 81, 113–152.
- (2014): “National Institutions and Subnational Development in Africa,” *The Quarterly Journal of Economics*, 129, 151–213.
- MYRDAL, G. (1957): *Economic Theory and Under-Developed Regions*, G. Duckworth.

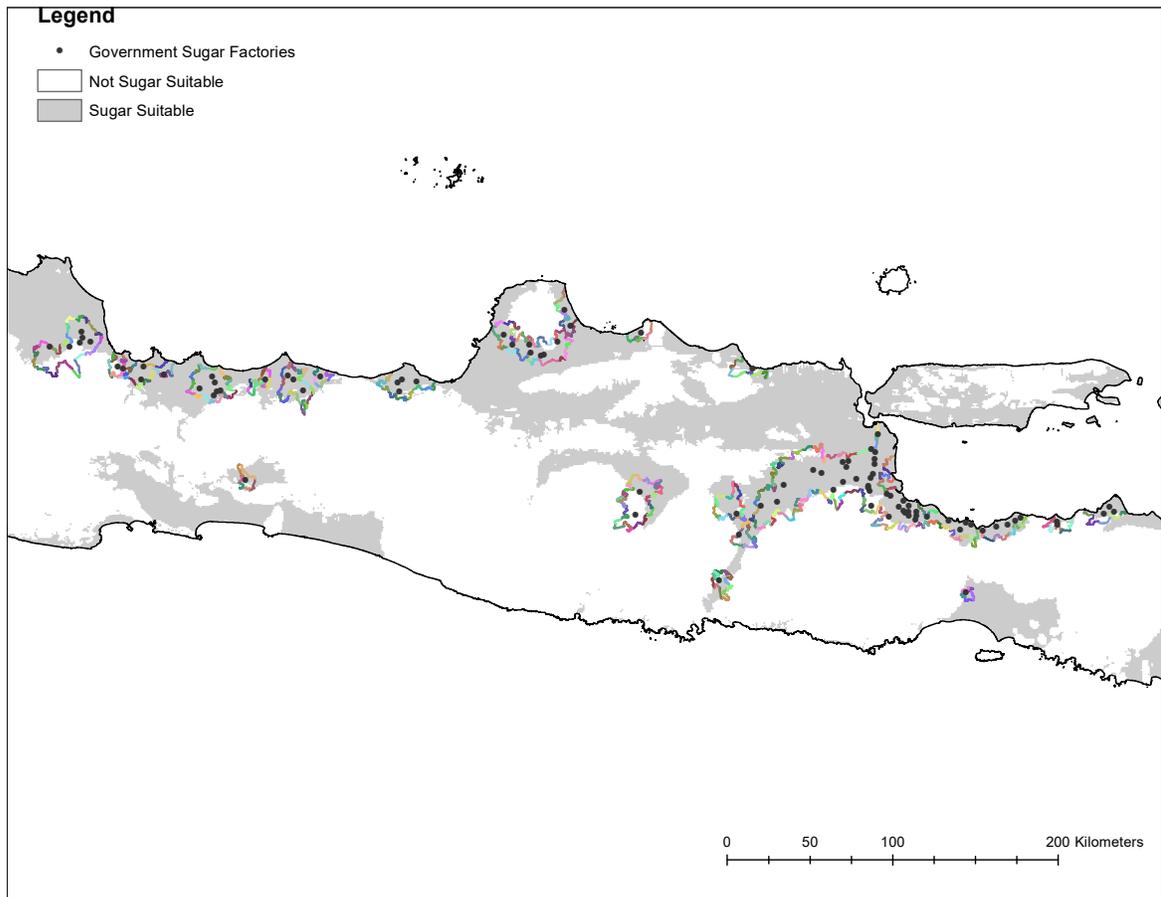
- NUNN, N. (2008): “The Long-Term Effects of Africa’s Slave Trades,” *The Quarterly Journal of Economics*, 123, 139–176.
- NUNN, N. AND L. WANTCHEKON (2011): “The Slave Trade and the Origins of Mistrust in Africa,” *American Economic Review*, 101, 3221–3252.
- OTO-PERALÍAS, D. AND D. ROMERO-ÁVILA (2014): “The Distribution of Legal Traditions Around the World: A Contribution to the Legal-Origins Theory,” *Journal of Law and Economics*, 57, 561–628.
- POPKIN, S. L. (1979): *The Rational Peasant: The Political Economy of Rural Society in Vietnam*.
- RASMUSSEN, P. N. (1956): *Studies in Inter-Sectoral Relations*, E. Harck.
- RICKLEFS, M. C. (2008): *A History of Modern Indonesia Since c. 1200*, Penerbit Serambi.
- SALOP, S. C. (1979): “Monopolistic Competition with Outside Goods,” *The Bell Journal of Economics*, 141–156.
- SCOTT, J. C. (1977): *The Moral Economy of the Peasant: Rebellion and Subsistence in Southeast Asia*, Yale University Press.
- SPOLAORE, E. AND R. WACZIARG (2013): “How Deep Are the Roots of Economic Development?” *Journal of Economic Literature*, 51, 325–369.
- TABELLINI, G. (2008): “Presidential Address: Institutions and Culture,” *Journal of the European Economic Association*, 6, 255–294.
- (2010): “Culture and Institutions: Economic Development in the Regions of Europe,” *Journal of the European Economic Association*, 8, 677–716.
- TOMLINSON, J. (2014): *Dundee and the Empire: ‘Juteopolis’ 1850-1939*, Edinburgh University Press.
- VALENCIA CAICEDO, F. (2014): “The Mission: Economic Persistence, Human Capital Transmission and Culture in South America,” Tech. rep., Mimeo.
- VAN NIEL, R. (1992): *Java Under the Cultivation System*, KITLV Press.
- (2005): *Java’s Northeast Coast, 1740–1840: A Study in Colonial Encroachment and Dominance*, CNWS Publications.
- VAN SCHAİK, A. (1986): *Colonial Control and Peasant Resources in Java: Agricultural Involution Reconsidered*, Amsterdam: Instituut Voor Sociale Geografie.
- VAN WAIJENBURG, M. (2015): “Financing the Colonial State: The Revenue Imperative and Forced Labor,” Working paper, African Economic History Network.
- VOIGTLÄNDER, N. AND H.-J. VOTH (2012): “Persecution Perpetuated: The Medieval Origins of Anti-Semitic Violence in Nazi Germany,” *The Quarterly Journal of Economics*, 127, 1339–1392.

Figure 1: Placebo Factories Example



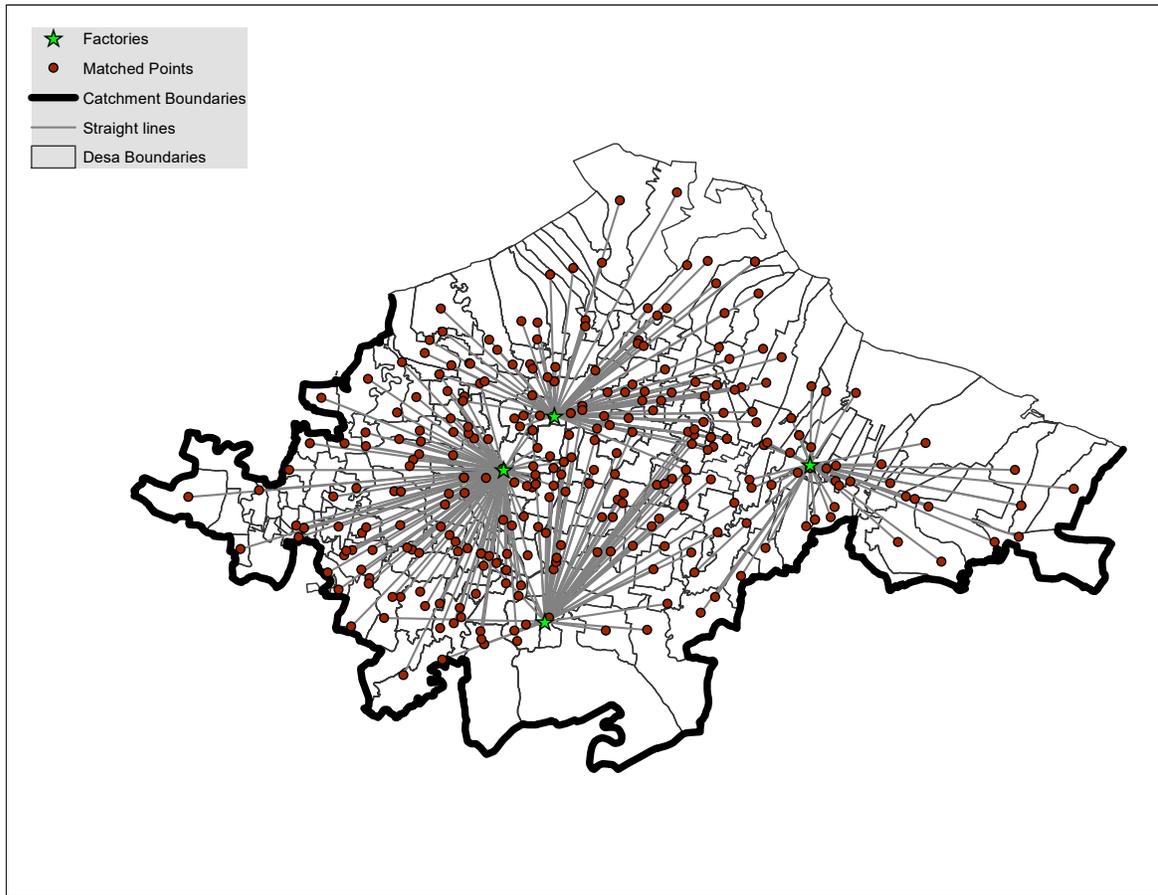
This figure illustrates the construction of the placebo factories, as described in Section 4.1.

Figure 2: The Cultivation System



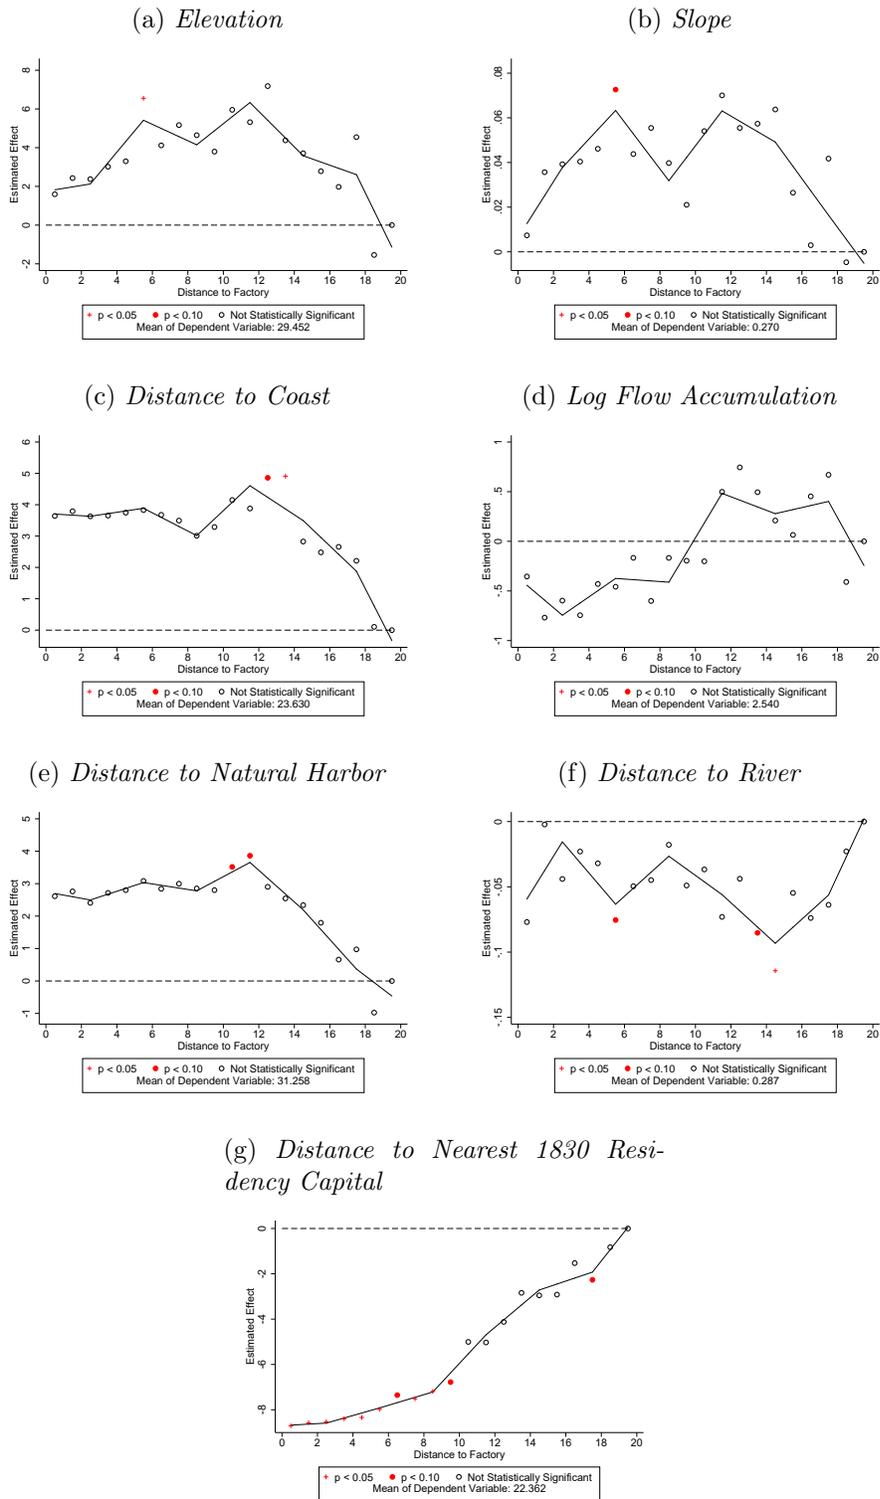
Sources: Commissie Umbgrove (1858).

Figure 3: Construction of Catchment Areas



This figure illustrates the construction of the catchment areas, as described in Section 3.1.

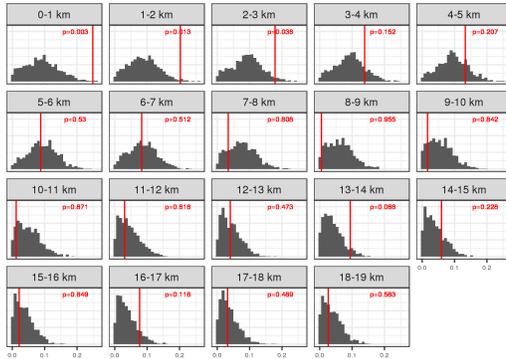
Figure 4: Geography



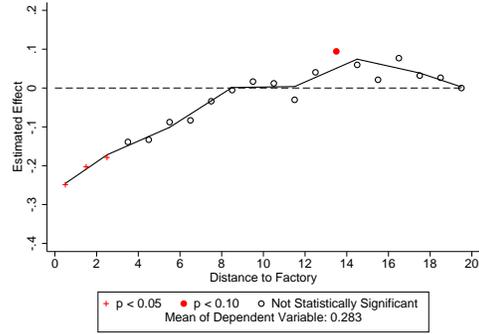
Notes: Points plot coefficients estimated from regressing the outcome variable on 1-km bins of distance to the nearest historical factory, controlling for nearest-factory fixed effects. The data are fit with a linear spline. p-values compare the impact of proximity to actual factories to the impact of proximity to 1,000 counterfactual factory locations.

Figure 5: Share in Agriculture (2001-11): Illustration of Methodology

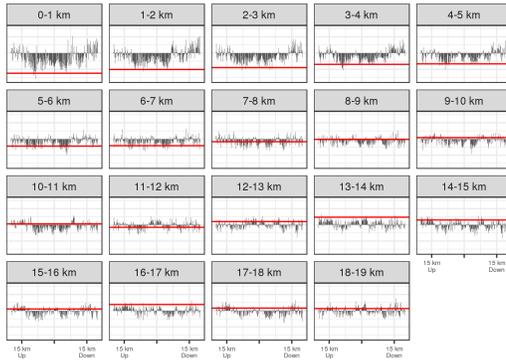
(a) *Independent Shifts: Counterfactuals*



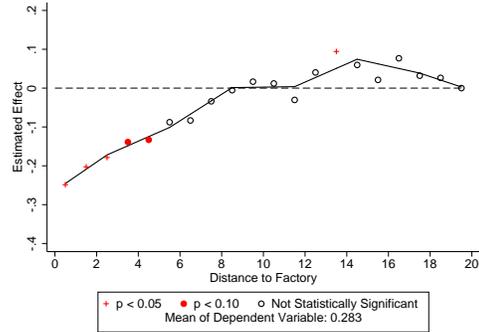
(b) *Independent Shifts: Plotted Coefficients*



(c) *Common Shifts: Counterfactuals*

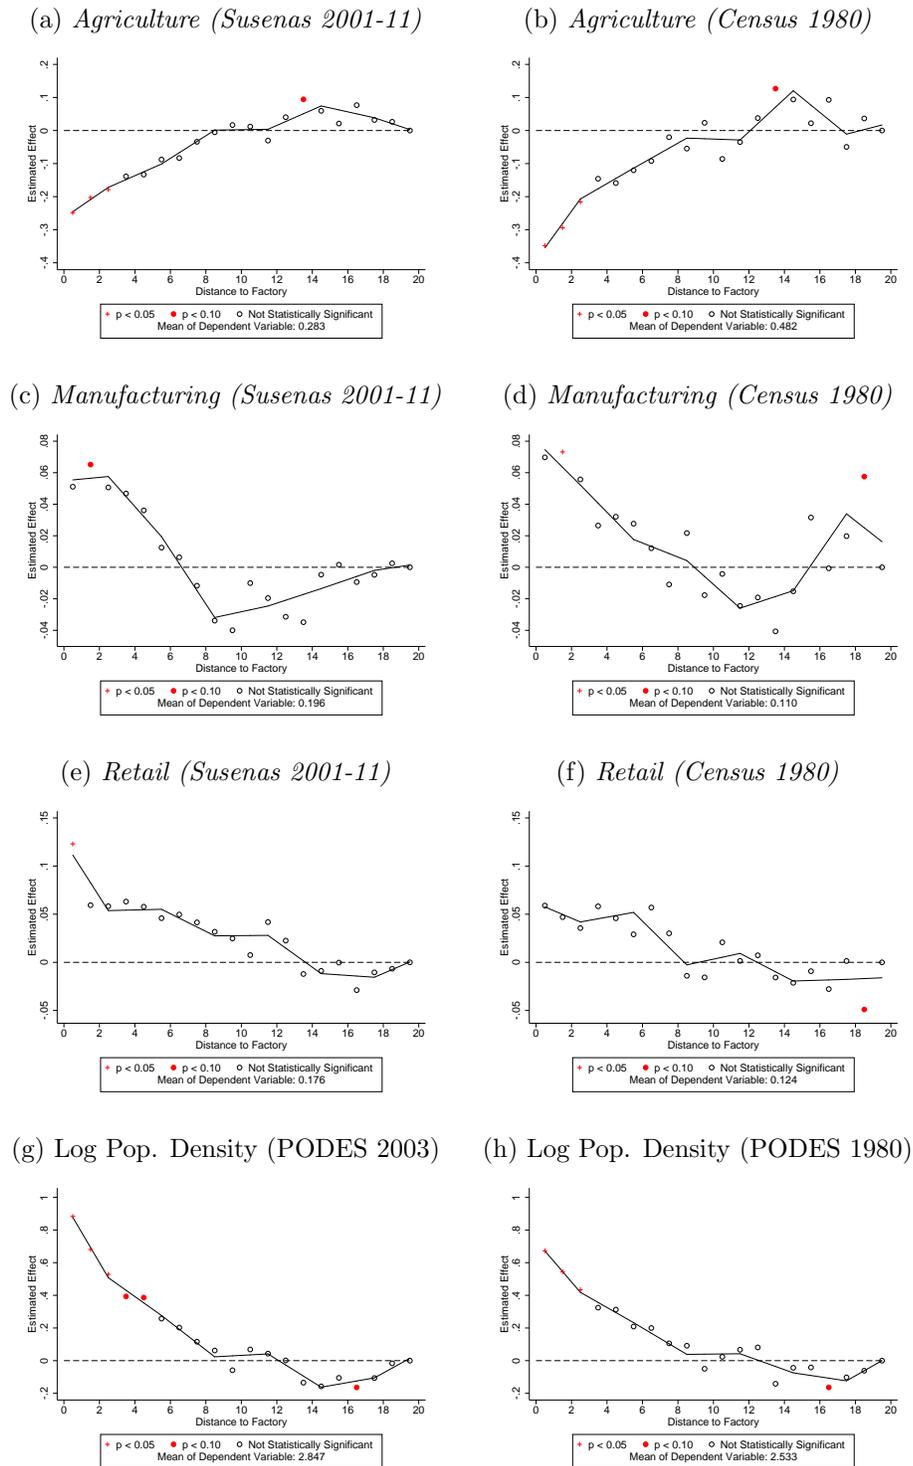


(d) *Common Shifts: Plotted Coefficients*



Notes: Panel (a) plots histograms of absolute coefficients from a regression of the outcome variable on bins in distance to counterfactual factories, controlling for nearest-factory fixed effects, geographic controls, a linear spline in distance to the nearest 1830 residency capital, and survey year fixed effects. The sample is restricted to men aged 18 to 55. For each factory, a counterfactual was selected at random from the region of the river network that was sugar-suitable and within 5-20 km via river from the real factory. This procedure was repeated to construct 1000 sets of counterfactual factories. The coefficients for distance to the real factories are shown as vertical lines. Panel (b) plots the real coefficients for each bin, with the symbols indicating their position in the distribution of counterfactual coefficients shown in panel (a). Panel (c) plots coefficients on distance to counterfactual locations, where here placebos were chosen to be a specific distance upstream or downstream from the real factories. Real coefficients are shown as horizontal lines. Panel (d) plots real coefficients for each bin, with the symbols indicating their position in the distribution of counterfactual coefficients shown in panel (c).

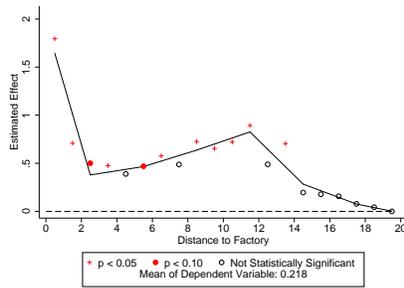
Figure 6: Industry and Agglomeration



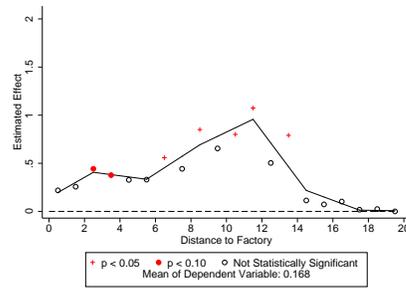
Notes: These figures plot coefficients estimated from regressing the outcome variable on 1-km bins of distance to the nearest historical factory, controlling for nearest-factory fixed effects, geographic controls, and a linear spline in distance to the nearest 1830 residency capital. Panels a), c), and e) include survey year fixed effects. In panels a) through f), the sample is restricted to men aged 18 to 55. The data are fit with a linear spline. p-values compare the impact of proximity to actual factories to the impact of proximity to 1,000 counterfactual factory locations.

Figure 7: Sugar and Linked Industries

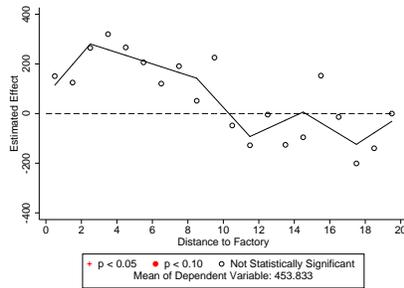
(a) *Log Value Sugar Processed (Full Sample, Economic Census 2006)*



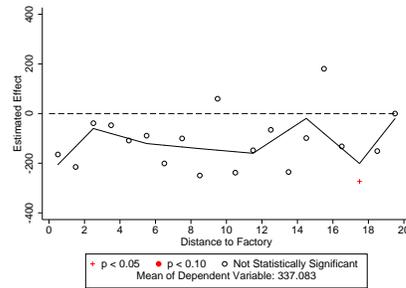
(b) *Log Value Sugar Processed (No Modern Factories, Econ Census 2006)*



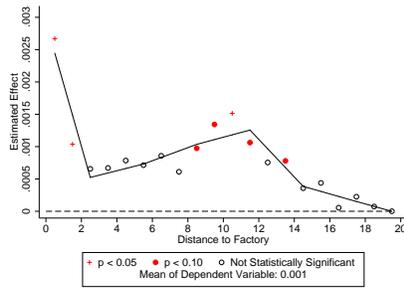
(c) *Tons of Cane Grown (Full Sample, PODES 2003)*



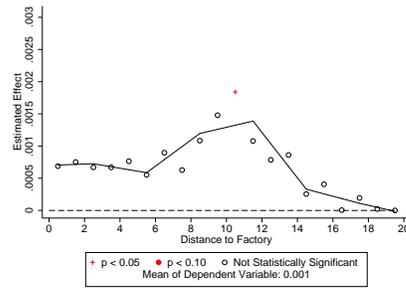
(d) *Tons of Cane Grown (No Modern Factories, PODES 2003)*



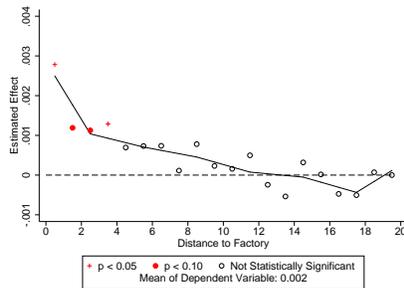
(e) *Employment Share Upstream (Full Sample, Economic Census 2006)*



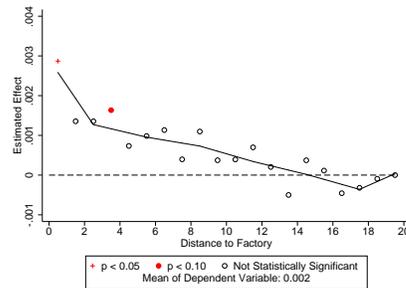
(f) *Emp Share Upstream (No Modern Factories, Economic Census 2006)*



(g) *Employment Share Downstream (Full Sample, Economic Census 2006)*

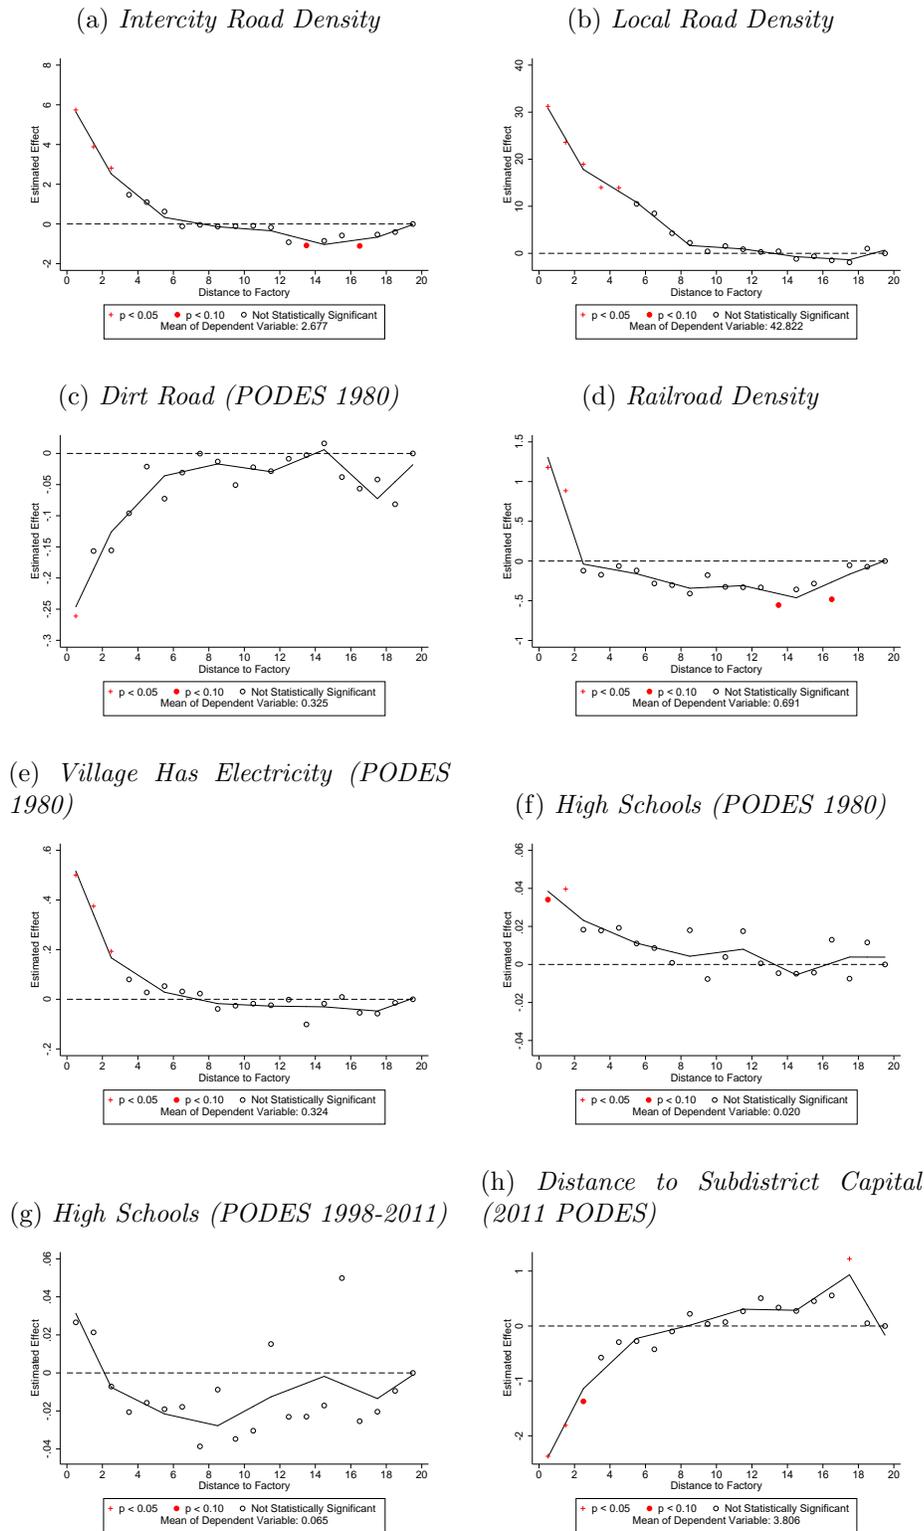


(h) *Emp Share Downstream (No Modern Factories, Economic Census 2006)*



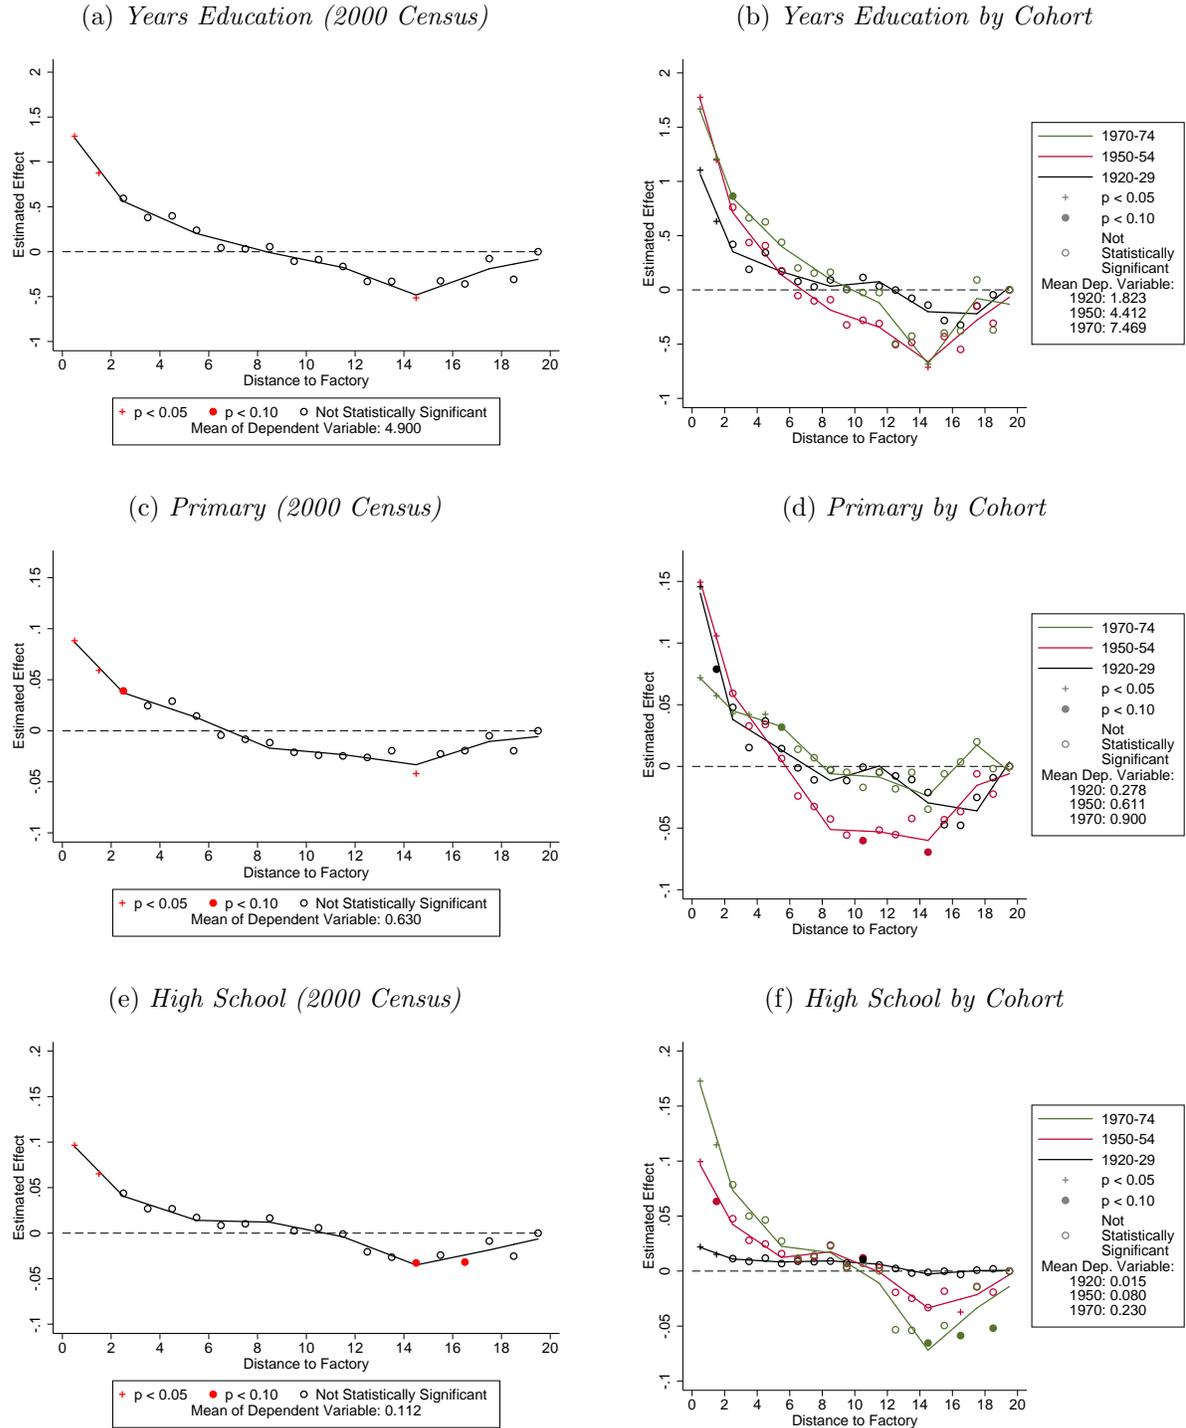
Notes: These figures plot coefficients estimated from regressing the outcome variable on 1-km bins of distance to the nearest historical factory, controlling for nearest-factory fixed effects, geographic controls, and a linear spline in distance to the nearest 1830 residency capital. The data are fit with a linear spline. p-values compare the impact of proximity to actual factories to the impact of proximity to 1,000 counterfactual factory locations.

Figure 8: Public Goods



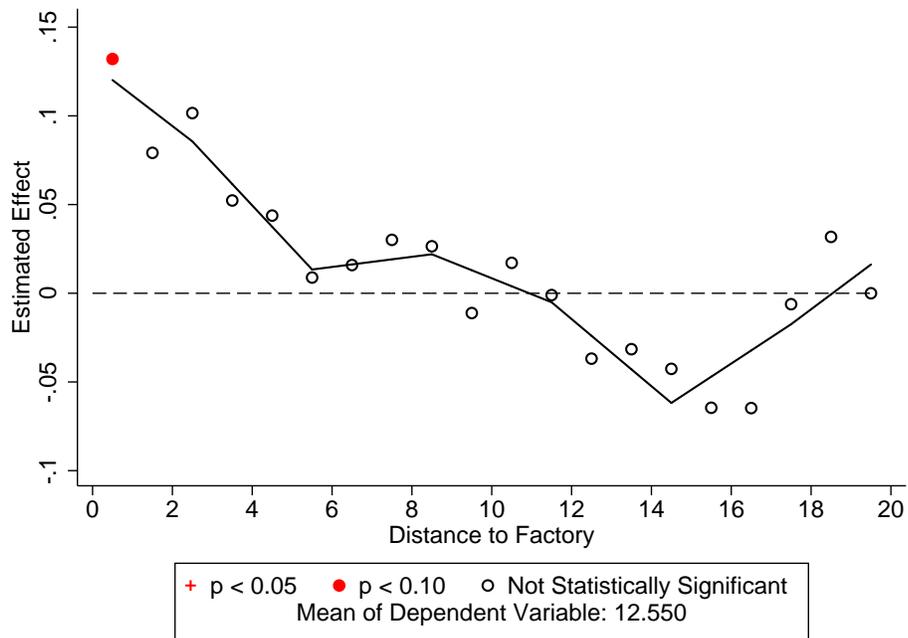
Notes: These figures plot coefficients estimated from regressing the outcome variable on 1-km bins of distance to the nearest historical factory, controlling for nearest-factory fixed effects, geographic controls, and a linear spline in distance to the nearest 1830 residency capital. Panel g) includes survey year fixed effects. The data are fit with a linear spline. p-values compare the impact of proximity to actual factories to the impact of proximity to 1,000 counterfactual factory locations.

Figure 9: Education



Notes: These figures plot coefficients estimated from regressing the outcome variable on 1-km bins of distance to the nearest historical factory, controlling for gender, nearest-factory fixed effects, geographic controls, and a linear spline in distance to the nearest 1830 residency capital. Left panels pool all birth cohorts and right panels plot separate coefficients for three birth cohorts. The data are fit with a linear spline. p-values compare the impact of proximity to actual factories to the impact of proximity to 1,000 counterfactual factory locations.

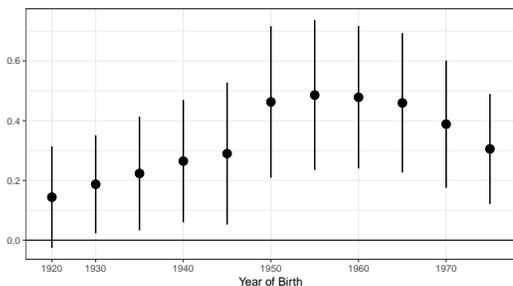
Figure 10: Expenditure (2001-11)



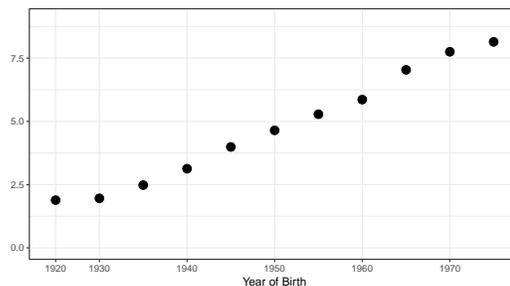
Notes: This figure plots coefficients estimated from regressing the outcome variable on 1-km bins of distance to the nearest historical factory, controlling for demographic variables, survey year fixed effects, nearest-factory fixed effects, geographic controls, and a linear spline in distance to the nearest 1830 residency capital. The data are fit with a linear spline. p-values compare the impact of proximity to actual factories to the impact of proximity to 1,000 counterfactual factory locations.

Figure 11: Education by Cohort: Subjected Villages (2000 Census)

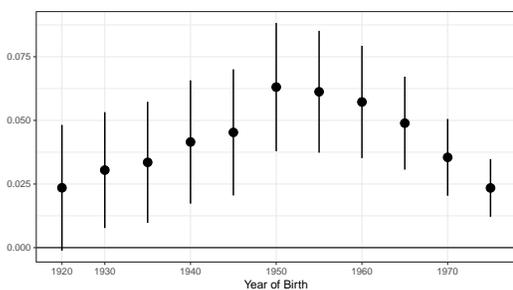
(a) *Years of Schooling*



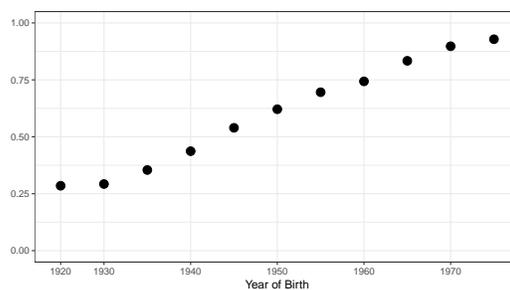
(b) *Years of Schooling Levels*



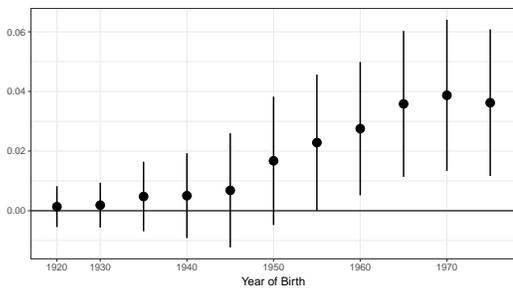
(c) *Primary Completion*



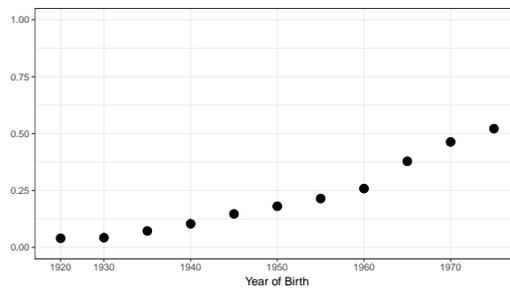
(d) *Primary Completion Levels*



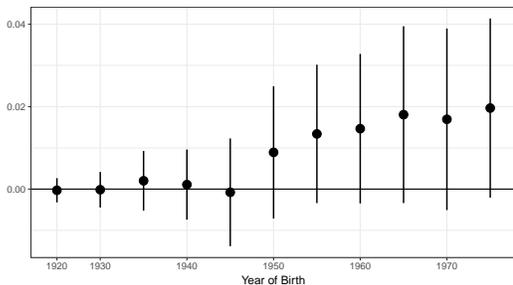
(e) *Junior High Completion*



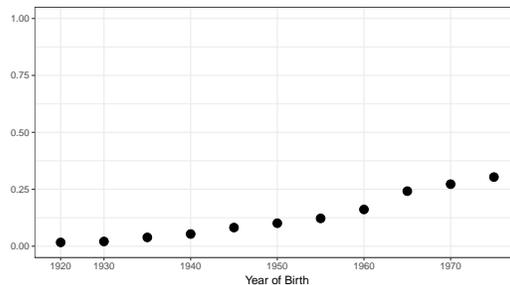
(f) *Junior High Completion Levels*



(g) *High School Completion*



(h) *High School Completion Levels*



Notes: In the left panels, each point plots a separate regression coefficient for different birth cohorts (1920-1929, 1930-1934, 1935-1939, ..., 1975-1979). Lines show 90% confidence intervals. In the right panels, points plot means. The unit of analysis is the individual, and the specification includes gender dummies, geographic controls, boundary segment fixed effects, linear splines in distance to the nearest historical factory and residency capital, and a linear polynomial in latitude and longitude estimated separately for each catchment area. Robust standard errors are clustered by subdistrict.

Table 1: Geographic Characteristics: Subjected Villages

	Elevation (1)	Slope (2)	Log Flow Accumulation (3)	On Coast (4)	Distance To Coast (5)	Distance To River (6)	Distance To Natural Harbor (7)	Distance To 1830 Residency Capital (8)
Cultivation	-2.209 (0.774)***	-0.019 (0.014)	0.014 (0.164)	-0.012 (0.016)	0.132 (0.114)	0.001 (0.012)	0.074 (0.105)	0.059 (0.141)
Obs	4,209	4,209	4,205	4,209	4,209	4,209	4,209	4,209
Clusters	380	380	380	380	380	380	380	380
Mean	31.60	0.26	2.54	0.06	25.50	0.29	32.94	24.89

Notes: The unit of observation is the village. Regressions include boundary segment fixed effects, a spline in distance to the nearest historical factory with kinks each 3km, and a linear polynomial in latitude and longitude estimated separately for each catchment area. Robust standard errors, clustered by subdistrict, are in parentheses.

Table 2: Land Tenure: Subjected Villages

	Village Land 2003 Total Land (1)	Land Share (2)	Village Land 1980 Total Land (3)	Land Share (4)	99th Pctile ÷ 90th Pctile (5)	90th Pctile ÷ 10th Pctile (6)	90th Pctile ÷ 50th Pctile (7)	50th Pctile ÷ 10th Pctile (8)
Cultivation	2.235 (0.852)***	0.014 (0.004)***	3.473 (1.538)**	0.012 (0.005)**	0.018 (0.262)	13.289 (8.469)	0.378 (0.272)	1.793 (2.007)
Obs	4,550	4,550	4,205	4,107	4,089	4,080	4,088	4,080
Clusters	383	383	380	380	381	381	381	381
Mean	18.61	0.09	23.95	0.11	3.53	38.58	4.34	7.84

Notes: The unit of observation is the village. Regressions include geographic controls, boundary segment fixed effects, linear splines in distance to the nearest historical factory and residency capital, and a linear polynomial in latitude and longitude estimated separately for each catchment area. Robust standard errors, clustered by subdistrict, are in parentheses.

Table 3: Schools (1980): Subjected Villages

	Public Non-INPRES Buildings (1)	Primary Teachers (2)	INPRES Buildings (3)	Primary Teachers (4)	Junior High Schools (5)	High Schools (6)
Cultivation	0.009 (0.019)	0.169 (0.104)	-0.035 (0.020)*	-0.207 (0.080)**	0.019 (0.009)**	0.007 (0.006)
Obs	4,205	4,205	4,205	4,205	4,205	4,205
Clusters	380	380	380	380	380	380
Mean	0.43	2.81	0.36	1.37	0.06	0.02

Notes: The unit of observation is the village. Regressions include geographic controls, boundary segment fixed effects, linear splines in distance to the nearest historical factory and residency capital, and a linear polynomial in latitude and longitude estimated separately for each catchment area. Robust standard errors, clustered by subdistrict, are in parentheses.

Table 4: Education: Subjected Villages

	Years Education (1)	2000 Population Census			1980 Census		Village Head	
		Primary School (2)	Junior High (3)	High School (4)	No School (5)	Primary School (6)	Years Education (7)	High School (8)
Cultivation	0.238 (0.075)***	0.026 (0.006)***	0.017 (0.007)**	0.008 (0.005)	-0.017 (0.009)**	-0.003 (0.010)	0.106 (0.087)	0.029 (0.013)**
Obs	16,125,747	16,125,747	16,125,747	16,125,747	653,313	653,188	26,630	26,630
Clusters	383	383	383	383	358	358	383	383
Mean	5.10	0.64	0.27	0.13	0.41	0.19	11.87	0.74

Notes: The unit of observation is the individual. Regressions include boundary segment fixed effects, linear splines in distance to the nearest historical factory and residency capital, geographic controls, and a linear polynomial in latitude and longitude estimated separately for each catchment area. Columns (1) through (6) include gender dummies, and columns (7) and (8) include survey year fixed effects. Robust standard errors, clustered by subdistrict, are in parentheses.

Table 5: Industrial Structure: Subjected Villages

	SUSENAS (2001-11)			1980 Population Census		
	Ag. (1)	Manuf. (2)	Retail (3)	Ag. (4)	Manuf. (5)	Commerce (6)
Cultivation	-0.042 (0.012)***	0.030 (0.009)***	0.013 (0.007)*	-0.039 (0.021)*	0.008 (0.012)	0.026 (0.010)***
Obs	130,335	130,335	130,335	127,873	127,873	127,873
Clusters	381	381	381	358	358	358
Mean	0.27	0.21	0.18	0.48	0.11	0.12

Notes: The unit of observation is the individual. The sample is restricted to men age 18-55. Regressions include geographic controls, boundary segment fixed effects, linear splines in distance to the nearest historical factory and residency capital, geographic controls, and a linear polynomial in latitude and longitude estimated separately for each catchment area. Columns (1) through (3) include survey year fixed effects. Robust standard errors, clustered by subdistrict, are in parentheses.

Table 6: Firms, Population, and Consumption: Subjected Villages

	Num. Manuf. Firms (1)	Manuf. Emp. Share (2)	Log Population Density (2003) (3)	Log Population Density (1980) (4)	Log Equiv. Consumption (5)
Cultivation	21.433 (8.513)**	0.023 (0.018)	0.067 (0.034)*	0.064 (0.032)**	0.007 (0.011)
Obs	4,549	4,549	4,550	4,107	144,046
Clusters	383	383	383	380	381
Mean	71.72	0.16	2.87	2.54	12.55

Notes: The unit of observation is the village in columns (1) through (4) and the household in column (5). Regressions include geographic controls, boundary segment fixed effects, linear splines in distance to the nearest historical factory and residency capital, and a linear polynomial in latitude and longitude estimated separately for each catchment area. Column (5) includes year fixed effects and household demographic controls. Robust standard errors, clustered by subdistrict, are in parentheses.