Persistence and transformation in economic development

Melissa Dell¹

Accepted: 7 September 2021 / Published online: 2 December 2021

© The Author(s) under exclusive licence to Editorial Office, Indian Economic Review 2021

Abstract

Much work in economic history has been done to study patterns of development in the USA and Europe. However, insights from other areas and regions are also required to better inform policy today that helps elucidate how policy interacts with the broader historical and institutional context and how the effects of these policies unfold over time. In this light, the paper focuses on illustrations from previous works of the author on extractive colonial institutions and their persistent effects on development paths later on. The first example is of the Peruvian mining *mita*. Here, the natives of Peru were forced by the Spanish to work in silver mines. The study focusses on the persistence of differential land tenures and public goods provision. The second example, a joint work with Ben Olken, considers Java's cultivation system in the 19th century, where the Javanese were forced to produce sugar that was sold in the world market by the Dutch. It shows that the sugar factory areas of the 19th century are more industrialized today. They are richer with better infrastructure and education levels as compared to the nearby counterfactual locations.

Keywords Long-run development · Channels of persistence · Forced labour · Land tenure · Public goods · Colonialism

JEL classifications H41 · N26 · O43 · N0 · O1

1 Introduction

Over the past 15 years or so, research on the persistent determinants of economic development has transformed economic history. Not only has it helped to inspire the integration of modern empirical methods with the study of historical data but it

---

Indian Economic Review public lecture, delivered on June 10, 2021 by Melissa Dell, Professor of Economics, Harvard University and John Bates Clark Medal, 2020 awardee.

✉ Melissa Dell
melissadell@fas.harvard.edu

¹ Department of Economics, Faculty of Arts and Science, Harvard University, Littauer Center M-24, Cambridge, MA 02138, USA

has also had an influence on economics more generally. In particular, it has contributed to the expansion of economics beyond its traditional US and European focus. Insights from this literature also have the potential to inform economic policy today by clarifying how policy interacts with the broader historical and institutional context and how the effects of these policies unfold over time.

In thinking about this literature, I draw inspiration from one of the greats of economics, Robert Solow, who said the part of economics that is independent of history and social context is not only small, but dull (Solow, 1997); and of course none of us want our research to be boring. And so regardless of one's area of study or field, bringing in some of that history and social context can really enrich the insights that come out of our research.

In this paper, I focus on insights that have been developed over the past decade using illustrations from my research. There is a large literature on persistence and transformation in development. In this review, I focus on some of my own work in this area and in addition to looking at some key learnings of this research, I outline promising future research directions.

2 Intellectual history

In the 1940s and 1950s, development economics was a central part of the broader field of economics. Writings in this genre tended to be very rich in institutional detail and in historical knowledge. But, they were not typically disciplined by formal theory or by data. Of course, those times predated modern computing; one would have to invert a matrix by hand to run a regression.

During the 1950s there was a major shift in economics and by the 1960s most economics journals would not publish anything that was not disciplined by a mathematical model. Fields like economic history and development economics then tended to move largely outside the mainstream. With the type of mathematical models economists had in the 1960s, with their assumptions of perfect information, perfect competition, it was difficult to describe the complex processes of economic development.

Most models of economic growth had a single long run steady state. History was irrelevant in those models; one always ended up in the same place regardless of initial conditions. Development practitioners, following the seminal work by Solow, tended to emphasize accumulating physical capital; later, they brought in human capital but there was less of an emphasis on the role that institutional context potentially had to play in economic development.

Economic history was by and large focused on the historical ascendancy of Europe and to an extent the United States. The literature tended to focus on questions such as why did the Industrial Revolution happen when it did and where it did, in England? There was much less of a focus on the economic history of the rest of the world.

Along with the resurgence of development economics in the 1990s and 2000s, led by economists like Michael Kremer, Abhijit Banerjee and Esther Duflo, there also came a growing interest in the specific role history played in economic development.

I like to think of this somewhat in terms of an analogy Paul Krugman made in explaining what happened to development economics; why did it die for a while?

Krugman (1994) uses an analogy of how European maps of the African continent evolved from the 15th to the nineteenth centuries. In the fifteenth century, maps of Africa were very rich and full of detail, often based on second- or third-hand reports by travellers, but some of the things on them were absurd—there were even monsters in some places! They were not disciplined by any empirical understanding of what these places actually were like. By the eighteenth century, improvement in the art of mapmaking raised the bar for what was considered valid data. So, while coastline and cities on the coast were shown with great accuracy, the interior for which reliable, verifiable information was not available emptied out. It became an empty space, a dark continent, an acknowledgement of the limits of their knowledge. And then, slowly over time, as Europeans became more familiar with sub-Saharan Africa, they started to fill in more details and the geographic details became more accurate.

In some ways, we can think of the journey of development economics along these lines. In the 1940s and 1950s, there were many rich theories of development. While they had some insights they were not necessarily based on facts. And given the tools of that time, the limited data, and the absence of processing power, they really could not have been.

Then for a while economists seemed to have lost interest, essentially saying their tools were not suited to examine these issues and they did not have all the facts. But, over time additional knowledge has accumulated. I would say we are still at a point where much of the map is blank: there is much that we do not understand about economic development. But, a lot of people are working very hard to add in more details to our understanding of why countries grow and develop economically, or do not, over time.

3 Insights—channels of persistence

Some of my research illustrates one thing in particular and that is the importance of understanding the channels or the mechanisms of persistence and economic development. We know by now that there is a lot of persistence. But, why is that? What generates that persistence?

One of the things we learned is there can be historical events or institutions that, at first glance, to a researcher living hundreds of years later, appear to be very similar. But, actually, they can have very different effects on economic development in the long run. An anthropologist would say, of course we know that we cannot generalize, and that everything is local, everything is specific. In fact, I do not think this means we cannot generalize. Rather, to understand the long run effects of something that happened in the past and to understand it in a general way we need to focus on the transmission mechanisms: these will determine the long run effects. This suggests moving beyond the emphasis on what we would call in causal language, the

Fig. 1 The *Mita*'s extent. The mita boundary is in black and the study boundary in light gray. Districts falling inside the contiguous area formed by the mita boundary contributed to the mita. Elevation is shown in the background. Source: Dell, M. (2010)

treatment. So, beyond saying, for instance, that “it is about extractive institutions”, we must move towards an examination of the intervening variables that matter.

I illustrate this with two different examples of extractive colonial institutions, drawing from my work. The first example is the Peruvian mining *mita* in which the Spanish forced the native peoples in Peru to work in the silver mines of Potosí for 240 years. The second example is the cultivation system in Java through which the Dutch in the nineteenth century forced Javanese to produce sugar that the Dutch then sold on world markets. The latter is joint work with Ben Olken.

3.1 Mining *mita* in Peru

The mining *mita* was instituted by the Spanish colonial government in 1573 and abolished in 1812. It required over 200 indigenous communities in Peru and Bolivia to send one seventh of their adult male population to work in the Potosí silver mines. They also worked in mercury mines because mercury was needed to get the silver out of the rocks that it was trapped in.

The Potosí mines were discovered in 1545. They provided the largest deposits of silver to the Spanish empire. A large share of the silver the Spanish brought

Fig. 2 Household equivalent consumption. District Level Scatter Plot of Consumption (2001). Source: Dell, M. (2010)

back from the new world came from Potosí. The *mita* assigned over 14,000 conscripts from Peru and Bolivia to the Potosí mines, and another several thousand to the Huancavelica to mine the mercury needed to refine the silver (Bakewell, 1984).

The map in Fig. 1 shows the *mita*'s extent. It covered much of what is today Peru and Bolivia. Potosí is in the south and Huancavelica in the north.

I ask does the *mita* exert a long-run impact on development as measured by among other things household consumption and children's heights (a proxy for consumption) because many children in this region today remain malnourished. And if there is a long run effect on economic development the question is why something that was abolished a couple of hundred years ago would still matter today. I focus on a couple of channels of persistence—the impact *mita* had on land tenure as well as on public goods. And to do that I draw together data from the seventeenth century through to the present and trace as best as I can, given data limitations, the impacts that the *mita* had over time.

Figure 2 is a zoomed-in map of the region that I focus on, and I compare across the green boundaries. The places inside these boundaries were forced to send a seventh of their male population to work in the mines whereas the places outside were exempt [40:20]. I focus on this area because three characteristics we are able to measure and proxies for the prosperity of a place based on how many

Fig. 3 Children's heights. District Level Scatter Plot of Stunting (2005). Source: Dell, M. (2010)

taxes they paid before the *mita*, and the geography, all look the same across these boundaries [40:36]. The rest of the *mita* boundary in Fig. 1 follows essentially the edge of the Andean range, so when you cross the boundary geography changes a lot and so do other characteristics [40:47].

Focusing on the area in Fig. 2, the boundary of the *mita* got determined essentially by ex ante deciding how many people were needed for the mines based on a survey commissioned by the King of Spain. Then census data were used to figure out the places to subject to *mita* and they just took the places that were the closest to the mines. Communities that were barely a kilometer or two inside this boundary had to send people to work in the mines, whereas communities outside the boundary did not have to. We can precisely locate these boundaries and if we look across them today there are indeed differences, measurable differences, in household consumption taken from the household living standards survey. The dots in Fig. 2 represent the raw data; the bigger the dot, the more people that were surveyed in that village; and the darker the color the lower is the consumption. The shaded background shows the predicted values from a regression that includes a polynomial in latitude and longitude to control for everything that changes smoothly. Since characteristics vary across space, having a flexible polynomial in latitude and longitude enables us to control for the things that change smoothly and just exploit the discontinuous change that happens as you move across this boundary from being an exempt village to a village subjected to *mita*. The figure shows that household equivalent

consumption drops discontinuously; places that are inside this boundary are about 25% poorer today.

Figure 3 displays children's heights. This is actually a census of children, so there is more power than in the household consumption survey. Again, the darker the color the higher the percentage of children that have stunted growth because they have not got enough nutrition. Moving across this boundary to the villages that were subjected to *mita* children become substantially more likely to have stunted growth due to the fact that the subjected places are poorer and have worse nutrition.

These figures show that forced labour during the colonial period still matters in a very concrete way for the people and for the children that are living in these places that used to be subjected to forced labor. It is natural to ask why on earth would this still matter? Forced labor was abolished 200 years ago. That is a really long time ago; do not we expect there to be convergence? Why is that not happening? That is the question that most of my paper (Dell, 2010) tried to focus on.

Through a reading of the historical literature I settled on land tenure as a major channel. This literature emphasises that the *mita* had a major effect on land tenure. The *mita* was instituted fairly early in Spanish colonial history before many Spaniards had arrived to form a landed elite in the region (Brisseau, 1981; Glave and Remy, 1978). There are a few reasons why we would expect the *mita* to impact land tenure.

When land owners arrived from Spain to settle this region, the state did not want to have competition from them for labour (Larson, 1988): in order for the Spanish king to make money from the silver, (20% of the silver was paid in tax to the King of Spain) he had to have people to work in the mines. And because there was population collapse during this period if there was free labour wages would have been too high for there to be any profit left. So, the king needed to be able to coerce these labourers to work in the mines. But, if Spanish settlers had been allowed to set up estates next to where the indigenous people who worked in the mines were living, the landowners would have wanted these people to work on their fields and resisted their being taken to work in the mines. The state did not want to have that competition. So, they did not allow the Spanish who showed up to settle the region to settle in areas that were subjected to the *mita*. They were only granted estates in areas that were exempt. Another reason for this was the workers in the mines were paid less than subsistence wages (Garrett, 2005; Tandeter, 1993; Cole, 1985). So, if the Spaniards were allowed to come and take away the workers' land the state would have to figure out how to feed them, whereas if it permitted the workers' communities to keep their land, then they could feed themselves.

Allowing them to keep their land also provided a means to compensate indigenous leaders who cooperated in administering the *mita*. The Spanish crown had very few people on the ground in Peru; people who worked in the mines were rounded up and taken to Potosí by indigenous leaders and the Spaniards needed a way to co-opt them. In short, there were a lot of reasons why the crown did not let Spanish landowners come into the areas that were subjected to the *mita*.

Land ownership, in turn, is very persistent. Even in the United States, work by Bleakley and Ferrie (2016), shows land ownership patterns persist over a hundred years. In a country like Peru where well-functioning land markets were absent, it is

Fig. 4 Haciendas 1689. District Level Scatter Plot of Haciendas (1689). Source: Dell, M. (2010)

very persistent even though the *mita* was abolished long back. So, in the non-*mita* districts private ownership by a Hispanic elite predominated.

The sixteenth and seventeenth centuries were a period of massive population collapse for the indigenous population. The Spanish settlers therefore did not face much resistance in seizing the land. It is very different than colonialism in Asia which did not see the same population collapse.

So Spanish settlers in non *mita* districts had secure property rights and title and faced little conflict over land. And these land ownership norms remained stable after the abolition of the *mita*. Thus, even after Peru gained independence Spanish settlers still had their land titles and these were stable and respected. Whereas, in the *mita* districts there was communal land tenure. At independence, communal land tenure was abolished but the Peruvian state did not replace it with any form of property rights; the people living in the *mita* areas did not have any enforceable title to their land. Spanish/Hispanic elites tried to come in and take the land. But, by this time the population had recovered and the indigenous people put up a lot of resistance. This led to a lot of instability and fighting over land.

Figure 4 shows *haciendas* (estates) in 1689. Inside the *mita* there are practically no *haciendas* (dark blue) with a couple of exceptions whereas in the area outside, there are lots of estates. There is no data for the full region, but this is as much data as has been preserved. Moving forward to 1845 (Fig. 5), there are some changes, reflecting the limited success of Spanish settlers in getting land. Even in 1940

Fig. 5 Haciendas 1845. District Level Scatter Plot of Haciendas (1845). Source: Dell, M. (2010)

(Fig. 6), by which time some Spanish elites have succeeded in seizing land in the former *mita* areas, there is still a discontinuity: it is still more likely for there to be *haciendas* outside the former *mita* areas than inside them.

Land reform was undertaken after 1940 so today there actually are no differences in land distribution. The elites largely left and their land was split up. But, there are still lingering effects. Long after the abolition of the *mita*, it still affects who owns land and how much stability there is over land, which in turn affects other downstream outcomes that matter for economic development since this is an agricultural region where land is the foundation of the economy.

We can see from Fig. 7 that pertains to 1876, that education historically was lower in the *mita* areas. It was actually really low everywhere; only children of the landed elites got an education. Coming all the way forward to very recently, as in Fig. 8 (2006), we see that there are fewer roads in the *mita* areas. And the reason for this is that when the Peruvian government built most of the modern roads in the early twentieth century, the owners of the largest estates lobbied for the roads to come to their estates so they could take their produce to market. Since there was stability for the *hacienda* owners outside the *mita*, they also had higher returns to having those roads. The roads were thus built to go to the estates of these large landowners who were also politically connected. Peru has not built a massive number of new roads in the second half of the twentieth century. Where they put the roads initially therefore matters a lot for where there are roads today.

Fig. 6 Haciendas 1940. District Level Scatter Plot of Haciendas (1940). Source: Dell, M. (2010)

In the *mita* areas, the roads that exist are by and large in very bad condition. And therefore, rates of market participation are much lower (Fig. 9). The agricultural census confirms that the vast majority of agricultural producers in the *mita* areas grow for subsistence. I have been to these places and talked to people, and they have said things like, “(W)ell, yes, I could grow more corn, but there’s no point because I don’t have a road to get it to the market in the nearest city.” So, because these areas did not have as many landowners historically due to the persistent effect of the *mita*, they do not have access to roads today which are a really important part of economic development. This provides a kind of a link to what happened hundreds of years ago to what life looks like for these people today.

It perhaps does not correspond completely to the stereotype in the traditional literature. That literature would say extractive institutions are bad because then the government does not protect property rights. To be sure, this was certainly true historically in the *mita*. But, even though the non *mita* districts were also under the Spanish crown, the long term presence of large landowners in them provided a stable land tenure system that encouraged public goods provision; in particular, the difference in road infrastructure across these two areas is stark and persistent.

In the literature on Latin America, the hypothesis is that it is poor relative to the US because of these large landowners. But, this research suggests that at least in Peru, the alternative to large landowners was not, say, something like Massachusetts with small, enfranchised landowners. Rather, the institutional structures in place did

Fig. 7 Education 1876. District Level Scatter Plot of Education (1876). Source: Dell, M. (2010)

not respect the rights of the indigenous people that held the vast majority of the land in the *mita* districts. Without the Spanish landowners, what would have obtained instead was just the state exploiting people initially, and then subsequently, poorly defined property rights and an under-provision of public goods.

Certainly, these exempt districts were not organised for the benefit of the masses, and of course, today they are still much poorer than the developed world. But, relative to the *mita* districts where there was exploitation by the state, they are actually doing better. To summarise, the evidence suggests this is because these large landowners did actually serve as a stabilising force and they did provide some public goods that persisted after the landowners were removed. And the populace today, now that the landowners are gone, does derive some benefits from them.

3.2 *Cultuurstelsel* in Java

I now turn an alternative example of extractive institutions, which is the Dutch Cultivation System (*Cultuurstelsel*) in Java. In the studies here, and the much broader literature starting with the famous work by Acemoglu et al. (2001), people perhaps identify the common *motif* of extractive colonisers setting up weak institutions, leading to poor property rights protection and lower long-run economic development.

Fig. 8 Roads 2006. District Level Scatter Plot of Road Density (2006). Source: Dell, M. (2010)

However, the history of different extractive colonial institutions suggests that while the goal of the colonisers was always the extraction of resources, and while these institutions were undoubtedly horrible for the people that were exploited by them, the colonisers created a quite diverse set of economic institutions across time and contexts, and these could potentially have different effects for long-run development. There are some examples where colonial powers wanted to extract some surplus for which they needed to establish complex economic systems: systems that could have had a counteracting positive effect on long-run development. With Ben Olken, I look at this in the context of the Dutch cultivation system in Java (Dell & Olken, 2019).

The island of Java was the center of the Dutch colonial empire. It has a modern population of over 160 million people. 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 sugar factories for export to Europe. The revenues at the peak of this system accounted for one-third of Dutch government revenue and 5 percent of its GDP (Luiten van Zanden, 2010). Revenues of this magnitude from this cultivation system actually comprise the largest example of colonial exploitation in history.

Sugar production later collapsed during the Great Depression and Indonesia is now the world's largest sugar importer. The places that were exploited to produce sugar historically do not produce sugar anymore.

Fig. 9 Market participation (1994). District Level Scatter Plot of Ag Market Participation (1994). Source: Dell, M. (2010)

To extract the sugar, the Dutch established 94 water-powered sugar-processing factories that made use of what at the time was very modern technology. Because raw sugarcane is very heavy to transport, it had to be grown nearby. So, the Dutch created a catchment area with a radius of approximately 5 km around each of the factories that they built. Since these factories were water-powered they were along rivers. The Dutch forced all villages in the catchment area to reorganize their land to grow sugarcane. Since there were not many Dutch officials on the ground in rural Java, this system was largely organized by local officials who were empowered and given financial incentives by the Dutch. Reports from the 1860s show that over two and a half million workers were forced to labour in the sugar factories or related services, and they were joined by free labourers whose number expanded over time (Elson, 1994).

There are two distinct types of impacts this system could have had on the local development: the effects of creating a sugar-manufacturing infrastructure by building the factories; and the effects of subjecting villages to forced cultivation of sugarcane. In principle, the long-run effects could be very different.

To estimate the causal impacts of this system we exploit the fact that places that grew sugarcane were typically adjacent to each other, all in the coastal plain of Java. This is a geographically homogenous place, very flat, with similar soils and climate. The fact that they tend to be adjacent to each other means that there

Fig. 10 Sugar factory location example. This figure illustrates the construction of the placebo factories, as described in Sect. 4.1 of Dell, M., & Olken, B. A. (2019). Source: Dell, M., & Olken, B. A. (2019)

Fig. 11 Cultivation system catchment areas. Source: Commissie Umbgrove (1858)

are many different equilibria for site selection of factories (which could not be too close since each required a sugarcane cultivation catchment area of adequate size). If one factory moves two kilometers upstream, the factory next to that will also have to shift two kilometers upstream to maintain the minimal production size required for it to operate. Similarly, the factory downstream from that will have to shift by two kilometers, and so on for the entire set of factories so that they can all maintain the minimum amount of land and minimum production capacity. This is similar to Salop's (1979) model of spatial competition.

We estimate the effects of being close to an actual factory using highly disaggregated data, looking at villages that are one, two, three kilometers away. We compare these effects to the distribution of estimated effects of being near plausible, *counterfactual* factory sites. For a counterfactual, we essentially randomly move all the factories a little bit, so that they are at a locational equilibrium where they all have space to meet their minimum production capacity. We use GIS to identify counterfactual sites using two criteria: they must be along the same rivers as actual factories, and they have to have sufficient nearby land for sugarcane cultivation.

An example is in Fig. 10, which is a map of an area of the Javanese coast. The blue area is suitable for sugar, the red is not. The green dot is a factory, it is along a river and there is a 5-km radius that a factory needs to meet its capacity. If we could

Fig. 12 Share of agriculture. **a** Plots the real coefficients for each bin, with the symbols indicating their position in the distribution of counterfactual coefficients. **b** Plots real coefficients for each bin, with the symbols indicating their position in the distribution of counterfactual coefficients. Source: Dell, M., & Olken, B. A. (2019)

potentially shift that factory a few kilometers upstream, one can see that there is still enough land suitable for cane cultivation to meet the factory capacity.

In fact, there are many different places where this could be shifted, a few kilometers upstream or downstream, and it would be able to meet its capacity. But, rather than shift one factory at a time, we shift all 94 factories at once, to yield these counterfactual arrangements. One can find thousands of such arrangements and have distributions of counterfactual effects. This enables us to identify the first of the effects: the impacts of being near this factory infrastructure that the Dutch constructed.

To look at the second effect, the impacts of a village being forced to cultivate sugarcane, we match the over 10,000 villages that contributed to the system, using handwritten records with modern geo-referenced sub-village locations to construct the precise catchment boundaries, and in turn combine these with highly disaggregated outcome data. Similar to the *mita* study, we undertake a spatial comparison across the boundaries of the catchment area.

Because Java is so densely populated, there are tens of millions of people living within a few kilometers of these boundaries. And this allows us to control for smooth variation, including controlling for the distance to the nearest historical factory.

Thus places inside and outside the boundaries are both similarly near the factory, but one had to contribute its land for forced cultivation and the other did not. In Fig. 11, we see the boundaries of the different catchment areas. The black dots are the factories; the areas around them had to contribute sugarcane; we divide the boundaries of those catchment areas into short segments, and compare places near the same segment, just barely on either side of the boundary.

To summarise the results, first of all, for the impacts of sugar factories: 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. That was true in 1980, and it is true today. In particular, in the immediate

Fig. 13 Share of manufacturing. 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. **a** include survey year fixed effects. The sample is restricted to men aged 18–55. The data are fit with a linear spline. p values compare the impact of proximity to actual factories to the impact of proximity to 1000 counterfactual factory locations. Source: Dell, M., & Olken, B. A. (2019)

vicinity of the historical factories, there is more employment today in industries downstream from sugar processing (i.e., industries that use processed sugar as an input—including food processing). This holds even when we exclude the few places where there are modern sugar factories. So, by and large, even though the sugar disappeared in the 1930s, the places that were close to where a sugar factory built by the Dutch used to be still industrial in manufacturing products that are downstream from sugar. This includes a lot of food processing.

There is also more public infrastructure, roads, electricity and schools. Historically the Dutch had to build roads and railroads to the sugar factories to take the sugar to the ports so they could sell it on world markets. Roads and rail networks today in Java are basically the places where the Dutch built initially. As we might expect in places with more manufacturing and better infrastructure, household consumption is also higher.

Figure 12 describes the share of people working in agriculture, in relation to the distance from the nearest historical factory; each dot represents a bin in terms of this distance, in kilometers. So, for instance, people living one kilometer away from the historical factory are much less agricultural than people would be 20 km away. We have differenced out the effects of being close to a counterfactual factory. It is only when you look near the actual historical sugar factories built by the Dutch, that these places are substantially less agricultural, 20 percent to 30 percent less, in 1980 and today, relative to places further away. Those effects fade out after a few kilometers. But, that affects a lot of people given how densely populated Java is.

The converse of being much less agricultural is having more manufacturing (Fig. 13). The places around these historical factories are the manufacturing centers of this region today.

Fig. 14 Modern sugar processing. 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 1000 counterfactual factory locations. Source: Dell, M., & Olken, B. A. (2019)

There are very few places that have a modern sugar factory in the same place as a historical sugar factory, which are in the top left of Fig. 14. We drop these few places in our analysis. From Fig. 14, we see then that the value of sugar processing is no different as we move away from the historical sugar factories. Nowhere are they growing raw sugarcane near the factories; thanks to improvements in transport technology, there is no need to do that anymore.

In Fig. 15, we look at employment share upstream from sugar processing, that is, workers engaged in producing agricultural machinery and capital machinery for the sugar factory. Again, we drop the few places where modern sugar factory locations coincide with locations of historical sugar factories. The employment share upstream in the production of machinery for sugar factories is found only nearby the locations of modern sugar factories.

Fig. 15 Employment share upstream. 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 1000 counterfactual factory locations. Source: Dell, M., & Olken, B. A. (2019)

Fig. 16 Employment share downstream. 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 1000 counterfactual factory locations. Source: Dell, M., & Olken, B. A. (2019)

However, there is a lot of employment in downstream industries (Fig. 16), that is, any industry that uses sugar as an input, around the places that had historical sugar factories, which disappeared many decades ago.

Figure 17 shows road and rail density in 1900; it is a well-known fact in the literature that this network was constructed to get sugar to the ports for export to Europe, and Fig. 17 reflects this. Figures 18 and 19 show, respectively, road and rail density today. Inter-city and local roads are denser today in the places where they were dense historically. These places are less likely to be accessible only by dirt roads; they also still have better railroad infrastructure.

Fig. 17 Historical road and rail density. 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 1000 counterfactual factory locations. Source: Dell, M., & Olken, B. A. (2019)

Fig. 18 Modern road density. 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 1000 counterfactual factory locations. Source: Dell, M., & Olken, B. A. (2019)

They also have better schooling infrastructure. Figure 20 shows high schools in 1980 and today. We can look in the 2000 census on years of schooling (Fig. 21), so we have data for millions of people and we can look across different cohorts. The cohort born 1920–1929 (the dark line) were educated during the Dutch period. Then we have a cohort born 1950–1954, educated right after independence, and finally a cohort born 1970–1974 that was educated after Indonesia really expanded its schooling system. Across all these different cohorts people that are living in the immediate vicinity of these historical sugar factories have higher schooling.

Not surprisingly they are also about 10% richer today (as indicated by household consumption in Fig. 22). This is what one would expect given that they have more

Fig. 19 Dirt road and modern railroad density. 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 1000 counterfactual factory locations. Source: Dell, M., & Olken, B. A. (2019)

Fig. 20 Schooling. 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. It 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 1000 counterfactual factory locations. Source: Dell, M., & Olken, B. A. (2019)

manufacturing and, better infrastructure; we would expect this also from the returns to more education.

Having discussed the effects of being close to where the Dutch built this kind of modern factory infrastructure, we turn to the effects of being forced to cultivate sugar. We see effects that are very different than in the case of the *mita*. But, it turns out that land is a really important channel of persistence because like in the case of the *mita*, the sugarcane cultivation system had a major impact on land tenure.

Fig. 21 Schooling effects by cohort (2000 census). 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 1000 counterfactual factory locations. Source: Dell, M., & Olken, B. A. (2019)

Fig. 22 Household consumption. This figures 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 1000 counterfactual factory locations. Source: Dell, M., & Olken, B. A. (2019)

In the villages subjected to the cultivation system, the Dutch had land redistributed towards the village to grow the sugarcane on. The cane was cultivated on this village-owned common land. And after the cultivation system was abolished,

all the way through to today, the village still owns more of that land, than common land that is owned by villages that were not subjected to this system.

The village-owned land was used to pay public servants in the village, including the village head and other village employees. We see that education is modestly higher in subjected villages today, and going all the way back to cohorts that were educated before and after independence.

History books talk about how the Dutch did not provide schooling for Indonesian people. That in the colonial period, for a school to be established, it had to be funded by the village; and it could be funded out of this land. After the Dutch cultivation system was abolished, cynics might say that the village-owned land would probably be used for the purposes of corruption. But, in actual fact, the villages seem to have still held the land and used it to provide at least some public goods especially education.

In these villages, a greater percentage of households work in manufacturing and retail and fewer in agriculture. There are no statistically significant differences in consumption; the magnitudes are not that large and are what one would expect from having a tenth of a year more schooling. They are a lot smaller than the effects of being near a factory.

So our interpretation of these two sets of results together is that in 1830, rural Java was heavily agricultural, and governed by feudal-like labour norms. There were large Javanese landholders and they had attached landless peasants who were forced to work on their land. This was not really an area that at this time at least was conducive to modern economic development, but the Dutch knew that they could exploit it more if they brought in modern economic production.

They, therefore, came in and built these modern factories and forced people to work in them. This was no doubt horrible for the people who were forced to work in the factories. But, if today some parts of Indonesia are richer than others it actually seems this system of forced sugar cultivation brought modern manufacturing infrastructure. It brought downstream industries.

Part of the reason for that is that the Dutch took all the good sugar. The Javanese got to keep the low quality sugar that would spoil before it got to Europe. So, they built these industries downstream from sugar around the sugar factories; this downstream industry stayed long after the exploitation had been removed. And we see effects both near the historical factories and in the subjected villages.

Of course, we are not looking at the impact of colonialism more generally; we do not have any idea what would have happened if the Dutch had never showed up in Indonesia at all. There is no real way to answer that question. But, places that were subjected to this system of colonial exploitation, in the long run, seem to have supported some economic changes that were beneficial from the standpoint of consumption.

These economic changes influenced the development trajectory and its persistence, owing to infrastructure investments, input–output linkages, agglomeration and human capital accumulation, among other channels.

The Dutch were not unique in forcing major economic changes in their colonies in an effort to maximise revenue extraction. I digitised a lot of data from Taiwan which along with Indonesia was a major sugar producer in the world. The Taiwanese

were forced to produce sugar by their Japanese colonisers and there are very similar results. But, the problem is that the Japanese spent about five years conducting very precise surveys and they found exactly the places that had the most population and were the most productive. And that is where they put their sugar factories. So, there is no way to causally identify the effect of the Japanese sugar production in Taiwan. Whereas in Indonesia, the Dutch did not do any of these surveys and put factories in places that really appear to be random. So, we are more able to identify effects.

Why did the Dutch sugar system lead to these positive long-run outcomes in the places that were subjected relative to the places that were not, in contrast to the *mita* or to other well-studied examples like the paper by Lowes and Montero (2016) on forced rubber cultivation in the Congo? We see that in particular, two factors are important: the role of manufacturing and investments in infrastructure.

The sugar cultivation system required substantial, and at the time, very sophisticated local manufacturing to process sugarcane. Raw sugarcane is heavy. It is full of water, spoils quickly, and could not be transported very far. So, the Dutch had to build the manufacturing infrastructure locally. By way of comparison, the *mita* conscripts were marched from their communities to a single location up to a thousand kilometers away, to Potosí to mine raw silver, which was then exported. Congolese rubber simply needed to be cut into cubes and dried before being exported to Europe. There was no infrastructure requirement in Peru or the Congo, at the scale that was needed in Java. Nor were linkages formed of the kind that developed between sugar and other sectors and that plausibly provide a key mechanism for this persistence for the propagation of industrialization over time in Java (Rasmussen, 1956; Myrdal, 1957; Hirschman, 1960).

Infrastructure investments are also important. Java had 94 sugar factories, and the need to build roads and rail to all of them resulted in very dense infrastructure, which still exists today. This is reminiscent of the findings of Dave Donaldson (2010) in his “Railroads of the Raj” paper in India.

In the Peruvian case, the infrastructure was built later, and outside the *mita* areas. The Spanish state did not need road infrastructure for the *mita*, but it affected settlement in a way that is not relevant for Indonesia. The Dutch by and large did not settle in rural Java. Whereas in Peru, the Spanish settled and where they settled are the places that got all the public goods later on. This is really important to the persistence there.

Because infrastructure is persistent almost by definition, it is a natural place to look to understand channels of persistence. The Dutch *Cutuurstelsel* and the Peruvian *mita* had very different impacts on infrastructure for the reasons outlined above. Again, in the case of the Congo, the Belgians had strong incentives to maximize short run extraction, before rubber took off elsewhere. They wanted to get as much rubber out as they could as fast as they could using extremely coercive practices. They did not need to invest in infrastructure; they just put the rubber on a boat and floated it to ports.

4 Moving beyond simple persistence–future directions

Before I conclude, a few words about future directions. Historical data show there is enormous persistence in development; places that were underdeveloped a long time ago, even looking within a narrow geographic region, tend to be the places that continue to be less developed today. There is now more or less widespread consensus that understanding mechanisms of persistence is important. But, the literature still largely focuses on identifying historical institutions that influence development today; the treatment of richer dynamics of persistence and transformation is much more limited.

But of course there are important examples of richer dynamics and there does not seem to be a great understanding of them. Some big open questions include the following. What factors promote takeoff? This has been studied in the context of England and the Industrial Revolution. But, outside Europe there is very little literature on takeoff despite examples like the East-Asian growth miracle where we have some of the most remarkable instances of economic growth in history. We do not have much understanding of these even though there are data that could be used to study them.

What forces lead to convergence, versus leapfrogging where places jump ahead more discontinuously? What about stagnation? Reversal? What are the micro foundations of persistence dynamics and how do historical forces interact with modern reforms? Answers to all of these questions would give us a richer sense of the dynamics of persistence, not just about something historically that matters today and not even just about why, but what those richer dynamics look like. And what about the cases where there is no persistence? I think we have a new set of tools that will allow us to make a lot more progress over the coming decade on these questions than we have been able to in the past.

Part of the reason why we do not have a very rich understanding of the dynamics of persistence is that a study of dynamics is predicated on rich data. Microeconomic data help a lot here, because more aggregated data are never going to be powerful enough to understand dynamics.

Traditional contexts such as economic growth in the US are quite specific: the US tended to grow two percent per year over a period spanning centuries. That is not what persistence looks like in most of the world. The fact that economics is a much more diverse field today than it was historically and that we have a much more diverse set of scholars who are familiar with more context, will enrich what we are able to understand. There are greater synergies with other frontiers in social science and other literatures such as on the economics of networks, behavioural economics; all of these might be relevant.

The challenge then is that outside the US and Europe the vast majority of historical data is in hard copy. There are many contexts where there is incredibly rich historical data but it would be very costly to digitise by hand. New methods in deep learning, however, can allow us to automate the digitisation of these complex historical data sources.

For many of these key questions and contexts that require disaggregated data it often does exist but is trapped in hard copy; unless it is something like the US census that has been digitized by ancestry.com. For the vast majority of countries, such data are not digitized and it is too costly to do so by hand. Also, important data can come from non-traditional sources whether from text, aerial photographs or satellites.

I have spent most of the past couple of years working on developing deep learning methods to automate the processing of economic data at scale. This includes both digitising complex historical documents like historical firm records, or tables, and also tools like natural language processing to be able to extract data that are trapped in texts.

If I just wanted to do this for my own research, it would take a lot less time. I wanted to make these tools available to the profession. We have released an open source Python library called Layout Parser that substantially lowers the barriers to performing a full document image analysis using deep learning.¹ This is not just about optical character recognition (OCR) of the content but about taking a raw document which could be a very complex table and automatically extracting the structured database needed to do quantitative analysis. The goal is to allow a researcher to build a full pipeline starting with a broad document, recognising the different types of information present in it, automatically classify and digitise it so that the end output is actually a structured database that can be analysed. I also include a link to a knowledge base on my website that gives some background and resources for understanding the deep learning methods that underlie the tools in this package.²

To conclude, I have argued that understanding history is important for elucidating processes of persistence and change in economic development. The literature to date has mostly focused on showing first that history matters and then providing some insights into why. But, moving forward, innovations in various areas including deep learning in particular, but also areas such as regional studies, economics of networks and so on have the potential to unlock significant new insights and much richer patterns of persistence and transformation.

Declarations

Conflict of interest The author has no known conflicts of interest.

References

- Acemoglu, D., Johnson, S., & Robinson, J. A. (2001). The colonial origins of comparative development: An empirical investigation. *American Economic Review*, 91(5), 1369–1401.
- Bakewell, P. J. (1984). *Miners of the Red Mountain: Indian Labor in Potosí, 1545–1650*. University of New Mexico Press.

¹ <https://layout-parser.github.io/>

² <https://dell-research-harvard.github.io/blog.html>

- Bleakley, H., & Ferrie, J. (2016). Shocking behavior: Random wealth in antebellum Georgia and human capital across generations. *The Quarterly Journal of Economics*, 131(3), 1455–1495.
- Brisseau, J. (1981). *Le Cuzco dans sa Région: Étude de l'aire d'influence d'une ville andine*. Inst. Français d'Etudes Andines.
- Cole, J. A. (1985). *The Potosi mita, 1573–1700: Compulsory Indian Labor in the Andes*. Stanford University Press.
- 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.
- Dell, M. (2010). The persistent effects of Peru's mining mita. *Econometrica*, 78(6), 1863–1903.
- Dell, M., & Olken, B. A. (2019). The development effects of the extractive colonial economy: The Dutch Cultivation System in Java. *The Review of Economic Studies*, 87(1), 164–203.
- Donaldson, D. (2010). *Railroads of the Raj: Estimating the Impact of Transportation Infrastructure (No. w16487)*. National Bureau of Economic Research.
- Elson. (1994). *Village Java under the Cultivation System, 1830–1870*. Asian Studies Association of Australia.
- Garrett, D. T. (2005). *Shadows of empire: the Indian nobility of Cusco, 1750–1825* (Vol. 90). Cambridge University Press.
- Glave, L. M., & Remy, M. I. (1978). Estructura Agraria e historia rural cuzquena: a proposito de las haciendas del Valle Sagrado: 1780–1930. *Cuzco, Archivo Historico del Cuzco*.
- Hirschman, A. O. (1960). *The Strategy of Economic Development*. Yale University Press.
- Krugman, P. (1994). The fall and rise of development economics. <http://web.mit.edu/krugman/www/dishpan.html>.
- Larson, B. (1988). Colonialism and agrarian transformation in Bolivia. Cochabamba, 1550–1900.
- Lowes, S., & Montero, E. (2016). Blood rubber: the effects of labor coercion on institutions and culture in the drc. Working Paper.
- Luiten van Zanden, J. (2010). Colonial state formation and patterns of economic development in Java, 1800–1913. *Economic History of Developing Regions*, 25(2), 155–176.
- Myrdal, G. (1957). *Economic theory and under-developed regions*. London: G. Duckworth.
- Rasmussen, P. N. (1956). *Studies in Inter-Sectoral Relations*, E. Harck.
- Salop, S. C. (1979). Monopolistic competition with outside goods. *The Bell Journal of Economics*, 10, 141–156.
- Solow, R. (1997). How did economics get that way and what way did it get? *Daedalus*, 126(1), 39–58. <http://www.jstor.org/stable/20027408>. Retrieved 15 Jul 2021.
- Tandeter, E. (1993). Coercion and Market. Silver Mining in Colonial Potosí, 1692–1826.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.