The Journal of Economic Perspectives

A journal of the
American Economic Association

Winter 2022

The Journal of
Economic Perspectives

A journal of the American Economic Association

Editorial offices:
Journal of Economic Perspectives
American Economic Association Publications
2403 Sidney St., #260
Pittsburgh, PA 15203
email: jep@aeapubs.org

The *Journal of Economic Perspectives* gratefully acknowledges the support of Macalester College.
Registered in the US Patent and Trademark Office (®).

Composed by American Economic Association Publications, Pittsburgh, Pennsylvania, USA.
Printed by LSC Communications, Owensville, Missouri, 65066, USA.

THE JOURNAL OF ECONOMIC PERSPECTIVES (ISSN 0895-3309), Winter 2022, Vol. 36, No. 1. The JEP is published quarterly (February, May, August, November) by the American Economic Association, 2014 Broadway, Suite 305, Nashville, TN 37203-2418. Annual dues for regular membership are $24.00, $34.00, or $44.00 depending on income; for an additional $15.00, you can receive this journal in print. The journal is freely available online. For details and further information on the AEA go to https://www.aeaweb.org/.
Periodicals postage paid at Nashville, TN, and at additional mailing offices.
POSTMASTER: Send address changes to the *Journal of Economic Perspectives*, 2014 Broadway, Suite 305, Nashville, TN 37203. Printed in the U.S.A.

The Journal of
Economic Perspectives

Contents

Volume 36 • Number 1 • Winter 2022

Symposium

Economies of Africa

Articles

Feature

Statement of Purpose

The *Journal of Economic Perspectives* attempts to fill a gap between the general interest press and most other academic economics journals. The journal aims to publish articles that will serve several goals: to synthesize and integrate lessons learned from active lines of economic research; to provide economic analysis of public policy issues; to encourage cross-fertilization of ideas among the fields of economics; to offer readers an accessible source for state-of-the-art economic thinking; to suggest directions for future research; to provide insights and readings for classroom use; and to address issues relating to the economics profession. Articles appearing in the journal are normally solicited by the editors and associate editors. Proposals for topics and authors should be directed to the journal office, at the address inside the front cover.

Policy on Data Availability

It is the policy of the *Journal of Economic Perspectives* to publish papers only if the data used in the analysis are clearly and precisely documented and are readily available to any researcher for purposes of replication. Details of the computations sufficient to permit replication must be provided. The Editor should be notified at the time of submission if the data used in a paper are proprietary or if, for some other reason, the above requirements cannot be met.

Policy on Disclosure

Authors of articles appearing in the *Journal of Economic Perspectives* are expected to disclose any potential conflicts of interest that may arise from their consulting activities, financial interests, or other nonacademic activities.

Journal of Economic Perspectives—Volume 36, Number 1—Winter 2022—Pages 3–32

Labor Productivity Growth and Industrialization in Africa

Margaret McMillan and Albert Zeufack

A shift from agriculture to manufacturing has been one of the hallmarks of job creation, poverty reduction, and rapid growth in low-income countries around the world. Industrialization is also one of the pillars of the African Union's Agenda 2063—the blueprint for transforming the continent into a global powerhouse (as described at https://au.int/en/agenda2063).

Some of the signs for industrialization in Africa are encouraging. The most comprehensive information about manufacturing employment in Africa only covers 18 countries, but based on those data, manufacturing employment in Africa's low- and middle-income countries increased from 6 million to more than 20 million from 2000 to 2018, raising the share of employment in manufacturing from 7.2 percent to 8.4 percent (Kruse et al. 2021). In comparison, the 1990s saw zero growth in Africa's manufacturing employment. Manufacturing exports from African nations have also grown at an annual average of 9.5 percent per year (Signé 2018). However, while employment and value-added shares of manufacturing in Africa are rising, both remain very low in comparison to the rest of the world (Diao, Harttgen, and McMillan 2017; Nguimkeu and Zeufack 2019).

The performance of the manufacturing sector varies considerably across African countries; this heterogeneity is to be expected given the vast differences in labor

■ *Margaret McMillan is Professor of Economics, Tufts University, Medford, Massachusetts. She is also Faculty Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts, and Senior Research Fellow, International Food Policy Research Institute, Washington, DC. Albert Zeufack is Chief Economist for Africa, World Bank, Washington, DC. Their email addresses are margaret.mcmillan@tufts.edu and azeufack@worldbank.org.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.36.1.3.

and natural resource abundance across Africa (Abreha et al. 2021). Nevertheless, most African countries have a low share of formal employment in the manufacturing sector. According to the INDSTAT database of the United Nations Industrial Development Organization (UNIDO 2020), only a handful of countries in Africa in 2018 had total formal employment that exceeded 100,000 in firms with ten or more employees: Ethiopia, Ghana, Kenya, Nigeria, South Africa, and Tanzania. Among these countries, Nigeria stands out, with around three million workers engaged in formal manufacturing between 2010 and 2012 (Nigerian Manufacturing Sector Report 2014). However, very little systematic information exists about the manufacturing sector in Nigeria. To the best of our knowledge, no longitudinal census of manufacturing firms in Nigeria exists. There is a World Bank Enterprise Survey for the year 2014, but it only covers 19 out of Nigeria's 36 states, and the estimated total employment generated using weights is implausibly small.

This essay considers the prospects for growth of manufacturing in Africa. We begin with a description of broad trends, drawing heavily on the results in Diao et al. (2021). Three striking patterns emerge. First, there has been a rapid increase in the number of African manufacturing enterprises with less than ten employees. We call these firms "small" and "informal" although some of these firms are likely to be formally registered; for the most part, we do not have enough information to make this distinction. In contrast, in Asian comparator countries, the share of employment in small manufacturing firms is flat or falling. Second, while Africa's large manufacturing firms appear to be productive, employment growth in these firms has not been rapid enough to decrease the share of small firm employment in total manufacturing employment, at least not yet. Third, labor productivity growth in Africa's manufacturing sector is largely accounted for by structural change (or the increase in manufacturing output associated with the increased employment share in manufacturing and the decline in activity in agriculture where labor productivity is lower than in manufacturing); conversely, within-sector labor productivity growth in manufacturing is close to zero.

The relatively nascent stage of manufacturing in Africa presents both opportunities and challenges. Probably the most important opportunity stems from the African Continental Free Trade Act which took effect on January 1, 2021. Given that the total population of Africa is roughly equal to that of China, the integration of African markets could attract the foreign investment needed to upgrade capabilities in the manufacturing sector and elsewhere. Some evidence suggests that this is already happening (Newman et al. 2016; Abebe, McMillan, and Serafinelli forthcoming). We also discuss opportunities for Africa's manufacturing in the "in-between sector," as well as manufacturing in sectors related to Africa's specific resources, green energy sources, and health needs. As Africa's manufacturing sector evolves, its biggest challenge is likely to be that global manufacturing around the world is shifting toward rising capital intensity—think "robots"—which reduces the importance of Africa's comparative advantage in unskilled labor-intensive manufacturing (Diao et al. 2021). We emphasize that the ready-made garment industry remains relatively low-skilled labor intensive. In addition, even if a shift

to capital-intensive manufacturing might limit manufacturing's direct role in job creation, there are large indirect benefits associated with manufacturing, including job creation in different but related sectors.

Industrialization spans a wide array of sectors and the continent of Africa is home to 54 countries. In this essay, we will focus on sub-Saharan Africa and exclude discussion of the countries of north Africa, in part due to data limitations and in part because neither of the authors of this piece has worked on these countries. We also focus almost exclusively on the 21st century, the period during which many African countries experienced rapid labor productivity growth (Diao, McMillan, and Rodrik 2019). Finally, we focus exclusively on manufacturing. For more on the status and potential of industries other than manufacturing, the reader is referred to the excellent work by John Sutton and his co-authors in his series of Enterprise Maps and a more recently published volume, which explores the potential for job creation in Africa's economies outside traditional manufacturing (Newfarmer, Page, and Tarp 2019).

Patterns of Labor Productivity Growth and Structural Change

The shift from agriculture to manufacturing can, in theory, raise labor productivity in several ways: growth in labor productivity in the agricultural sector, growth in labor productivity in the manufacturing sector, and growth in labor productivity from the structural change between lower-productivity agriculture and higher-productivity manufacturing. In the context of growth in African nations in the last two decades, it turns out that labor productivity growth outside agriculture is almost entirely from structural change, not from within the manufacturing or services sectors. In this section, we provide evidence for this claim; in the next section, we discuss how the patterns of Africa's manufacturing growth between smaller informal firms and larger formal firms help to explain what has been happening.

To describe these patterns, we use the recently released Economic Transformation Database (ETD) (de Vries et al. 2021). There are 18 sub-Saharan African countries included in this database; in 2018 these countries accounted for about 74 percent of sub-Saharan Africa's GDP and 64 percent of its population. It includes the two most populous countries, Ethiopia and Nigeria; the two richest countries, Botswana and Mauritius (measured using GDP per capita in 2018); and two of the poorest countries in Africa, Malawi and Mozambique (author's calculations using the World Development Indicators). For the most part, economywide value-added per worker in these countries is extremely low. For example, in 2018 employment-weighted value added per worker was only $4,689 (in current US dollars); however, the value for 10 of the 18 countries fell well below this average. Following Diao, McMillan, and Rodrik (2019), we examine the period 2001–2018, which corresponds most closely to the beginning of the "growth boom" in many African countries.

Figure 1
Labor Productivity Growth Decomposition

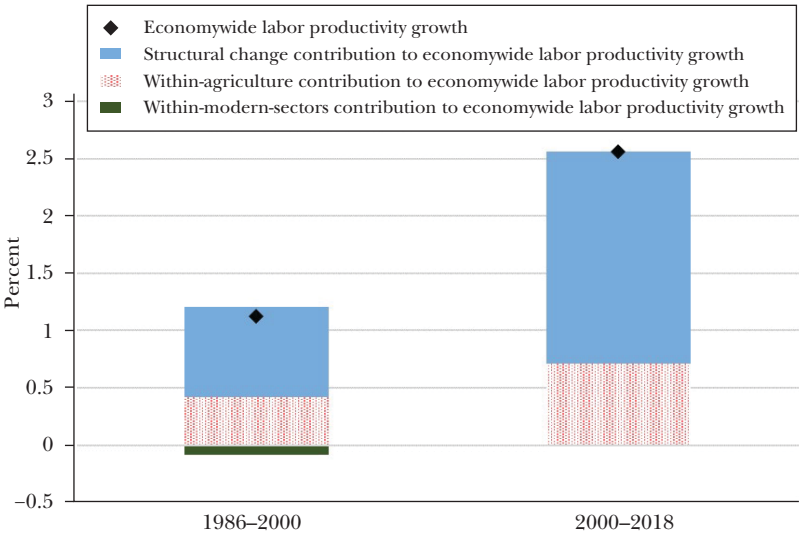

Source: Author's calculations using the Economic Transformation Database and based on 18 African countries: Burkina Faso (BFA), Botswana (BWA), Cameroon (CMR), Ethiopia (ETH), Ghana (GHA), Kenya (KEN), Lesotho (LSO), Mauritius (MUS), Mozambique (MOZ), Malawi (MWI), Namibia (NAM), Nigeria (NGA), Rwanda (RWA), Senegal (SEN), Tanzania (TZA), Uganda (UGA), South Africa (ZAF), and Zambia (ZMB).

The Economic Transformation Database data includes value added, employment, and price deflators for 51 countries for the period 1990–2018. Reliable data on capital inputs per sector are not available, which precludes calculations of total factor productivity growth. In any case, our preferred outcome measure here is value added per worker, because it corresponds most closely to GDP per capita and because it is more straightforward to interpret. Using these data for the 18 African countries, we compute value added per worker and employment shares by sector and use these values to decompose labor productivity growth into its "within" and "between" components. The within component is computed as an employment-weighted average (using initial period employment weights) of the change in value added per worker across sectors. The between component—or structural change component of productivity growth—is the sum across sectors of changes in employment shares multiplied by end-of-period productivity; it is positive when workers move from lower productivity sectors to higher productivity sectors. Because the emphasis in our work is on understanding how productivity is evolving outside of the agricultural sector, we further decompose within-sector productivity growth into its agricultural and non-agricultural components.

The bars in Figure 1 are coded according to how much of labor productivity growth comes from structural change between agriculture and non-agricultural

output (in blue) and how much comes from within-sector labor productivity growth in agriculture (in red) and in non-agriculture (in black). The annual growth rate for the nations of sub-Saharan Africa roughly doubled from the two decades before 2000 to the two decades after, rising from about 1.2 percent to 2.5 percent. Figure 1 shows that in Africa prior to the growth acceleration, average annual labor productivity growth is a little above 1 percentage point per year. After the growth acceleration, structural change contributes significantly to growth in Africa. This is not surprising because we expect the payoff to structural change to be greatest in poor countries. However, the contribution of within-sector labor productivity growth in the non-agricultural sector is close to zero.

To determine whether this broad pattern applies to the manufacturing sector, we report separately the within and between components for the manufacturing sector. Figure 2a displays these results ordering the countries from lowest within-manufacturing contribution to labor productivity growth (Burkina Faso) to highest (Mauritius). The growth from structural change is in blue, within-sector productivity growth is in green, and overall labor productivity growth in manufacturing is represented by the black diamond. The furthest bar to the right shows that, in total, manufacturing has contributed only around 0.25 percentage points to economywide labor productivity growth; on average this growth comes entirely from structural change. This is surprising since manufacturing has historically been an engine of growth.

A closer look at Figure 2a reveals some other lessons. There is considerable dispersion across countries of Africa in how manufacturing contributes to productivity. In several countries, the black diamonds are close to zero, indicating that when we combine the two sources of labor productivity growth, the manufacturing sector is not contributing to economywide labor productivity growth. This is especially surprising in the case of Ethiopia, where the government has placed a premium on the development of the manufacturing sector. But it is also surprising to see this in Ghana, Kenya, Senegal, and Tanzania, four countries whose economies have performed reasonably well since the early 2000s. The patterns for Botswana and South Africa are as expected; South Africa industrialized years ago and over time manufacturing has shed labor. Botswana is part of a customs union with South Africa and the fate of its manufacturing sector is closely tied to South Africa (McCaig and McMillan 2020).

Another lesson is that Figure 2a shows a generally negative correlation across countries between the productivity contribution from within-manufacturing and the contribution of the sectoral shift toward manufacturing. In the figure, this is illustrated by the blue and green lines going in opposite directions—one positive and the other negative. At one extreme, within-sector labor productivity growth in the manufacturing sector in Burkina Faso is negative 0.6 percentage points, while the between contribution from manufacturing is a positive 1 percentage point. At the other extreme, within-sector labor productivity growth in Mozambique's manufacturing sector is positive 0.6 percentage points while the between contribution of manufacturing is negative 0.2 percentage points. Indeed, this negative correlation

Figure 2A
Manufacturing Annual Average Labor Productivity Growth in Percentages, African Countries, 2001–2018

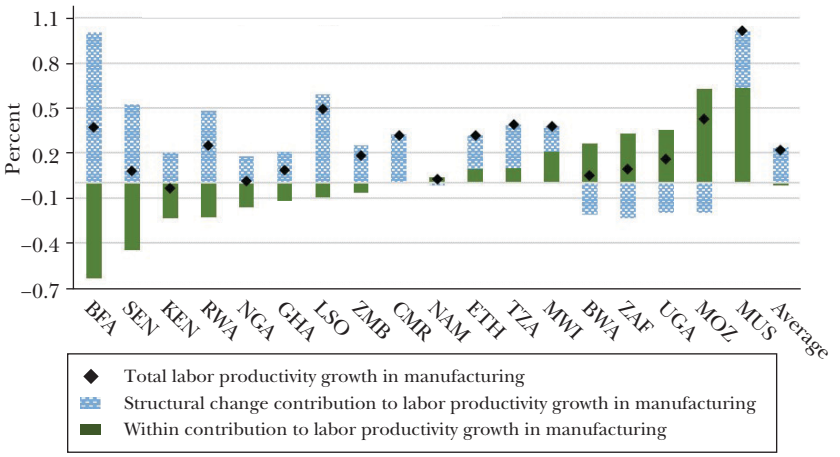

Source: Author's calculations using the Economic Transformation Database.
Note: The 18 African countries Burkina Faso (BFA), Botswana (BWA), Cameroon (CMR), Ethiopia (ETH), Ghana (GHA), Kenya (KEN), Lesotho (LSO), Mauritius (MUS), Mozambique (MOZ), Malawi (MWI), Namibia (NAM), Nigeria (NGA), Rwanda (RWA), Senegal (SEN), Tanzania (TZA), Uganda (UGA), South Africa (ZAF), and Zambia (ZMB). Data for Mauritius covers the period of 1973–2002 its period of industrialization. The average is for the 16 low and low-middle income African countries and so excludes Botswana and Mauritius.

holds for trade services, business services, construction, and transport sectors, as well as for manufacturing, but it seems especially surprising in the case of manufacturing, the canonical modern sector.

In contrast to the African countries, Figure 2b shows a comparison group of Asian countries, including early industrializers like Taiwan and Korea and late industrializers such as Bangladesh, India, Lao PDR, Cambodia, Sri Lanka, Myanmar, Nepal and Vietnam. Here, you can see a generally positive correlation between the within and structural change components of labor productivity growth for manufacturing—that is, the blue and green lines are both positive. Again, in the data for Asian industrializers, this positive correlation also holds for trade services, business services, construction, and transport sectors.

An Asian-style positive correlation between within-manufacturing growth and between-sector growth is consistent with the interpretation that rapid productivity growth in manufacturing is drawing in resources from the rest of the economy. An African-style negative correlation between within-manufacturing growth and between-sector growth is harder to understand. Apparently, the rapid growth of Africa's manufacturing sector in the last two decades is not being accompanied by rapid within-sector labor productivity growth.

Figure 2B

Manufacturing Annual Average Labor Productivity Growth in Percentages, Asian Countries, 2001–2018

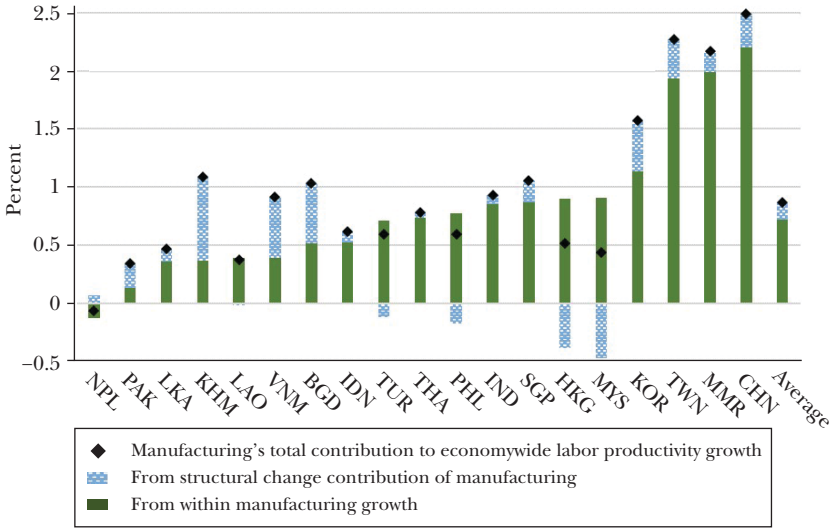

Source: Author's calculations using the Economic Transformation Database and Groningen Growth and Development Center (GGDC).
Note: The 17 Asian countries are Bangladesh (BGD), Cambodia (KHM), China (CHN), India (IND), Indonesia (IDN), Korea (KOR), Lao PDR (LAO), Malaysia (MYS), Myanmar (MMR), Nepal (NPL), Pakistan (PAK), Philippines (PHL), Singapore (SGP), Sri Lanka (LKA), Thailand (THA), and Vietnam (VNM), as well as Hong Kong, China (HKG), and Taiwan, China (TWN). Data for advanced countries/regions, including Korea, Singapore, Hong Kong, and Taiwan is from GGDC database averaged for the period of 1976–1990. For the rest of countries, data is from Economic Transformation Database averaged for the period of 2001–2018.

A Closer Look at the Manufacturing Sector in Africa

By taking a closer look at the patterns within the manufacturing sector, in particular, the role of small and informal manufacturing firms versus large formal enterprises, we can draw some inferences about the underlying reasons for the relatively poor performance of the manufacturing sector in Africa. Diao, McMillan, and Rodrik (2019) develop a model to highlight the differences between demand- and supply-driven structural change. In their model, supply-driven structural change is captured by a positive productivity shock to the modern sector (in this case, say manufacturing) allowing it to draw labor from other less productive sectors of the economy. To the extent that structural change is supply-driven, we would expect to see an expansion of modern sector (or formal) activity in the manufacturing sector. By contrast, demand-driven structural change was likely a result of positive aggregate demand shocks possibly due to some combination of factors like public investment, external transfers, or increases in rural incomes. Demand-driven structural change

is more likely to be accompanied by the entry of less-productive, smaller manufacturing firms.

Employment Growth Is Dominated by Small and Less Productive Firms

To explore this hypothesis, we again use employment data for manufacturing from two sources: the Economic Transformation Database (ETD) (de Vries et al. 2021) and the manufacturing employment data from the INDSTAT2 2020 database produced by the United Nations Industrial Organization (UNIDO 2020). The manufacturing employment data from the Economic Transformation Database is largely based on population census data, and so covers manufacturing in both the formal and informal sectors (Timmer, de Vries, and de Vries 2015). By contrast, INDSTAT2 records manufacturing employment data for formal firms in the manufacturing sector. Although country statistics sometimes vary in terms of the size of establishments covered, typically INDSTAT2 covers firms with ten or more employees. For several countries, we compute small and informal sector employment in the manufacturing sector as the difference between total employment (from the ETD data) and formal sector employment (from the INDSTAT2 data). We then plot total small informal and formal sector manufacturing employment for these countries.[1]

Figure 3a shows the results of this exercise for eight African countries: the two most industrialized countries in sub-Saharan Africa, Mauritius and South Africa, along with six African countries which experienced a relatively recent growth acceleration—Ethiopia, Ghana, Kenya, Nigeria, Senegal, and Tanzania. The only country for which UNIDO has data, which has been excluded from Figure 3a, is Botswana where the trend looks like that in South Africa. The vertical lines in each graph mark the start of the country-specific growth acceleration documented in Diao et al. (2018). The most striking trend in all but two of these countries is the upward-sloping curves for employment in small and informal manufacturing employment and the relatively flat lines for formal sector manufacturing employment. The patterns in Ghana, Kenya, Nigeria and Senegal are similar to those in Ethiopia and Tanzania as presented in Diao et al. (2021). In short, the beginning of the rise in small/informal sector employment largely coincides with the beginning of the growth acceleration in the African countries. However, the relatively higher-income African countries South Africa and Mauritius have a different pattern. In more-developed economies, formal employment is a much higher share of total manufacturing employment—although

[1]We gauged the accuracy of these data with comparisons to other data sources. For the recent total manufacturing employment numbers reported in the Economic Transformation Database, we looked at estimates of manufacturing employment based on firm level datasets and living standards measurement surveys (LSMS) for Ethiopia and Tanzania. This exercise leads to a reduction in total employment in manufacturing in recent years but no change in the aggregate patterns. A summary of these results can be found in Diao et al. (2021). For the INDSTAT2 employment data, we also plot formal sector employment data in Ethiopia, Tanzania and Vietnam—the three countries for which we have longitudinal census data for the formal manufacturing sector (these data are described extensively in Diao et al. 2021). The INDSTAT2 series and the series of formal sector employment data coincide almost perfectly. This is not surprising since UNIDO obtains its manufacturing employment data from national statistical agencies, but it is nevertheless reassuring.

Figure 3A
Manufacturing Employment in Africa

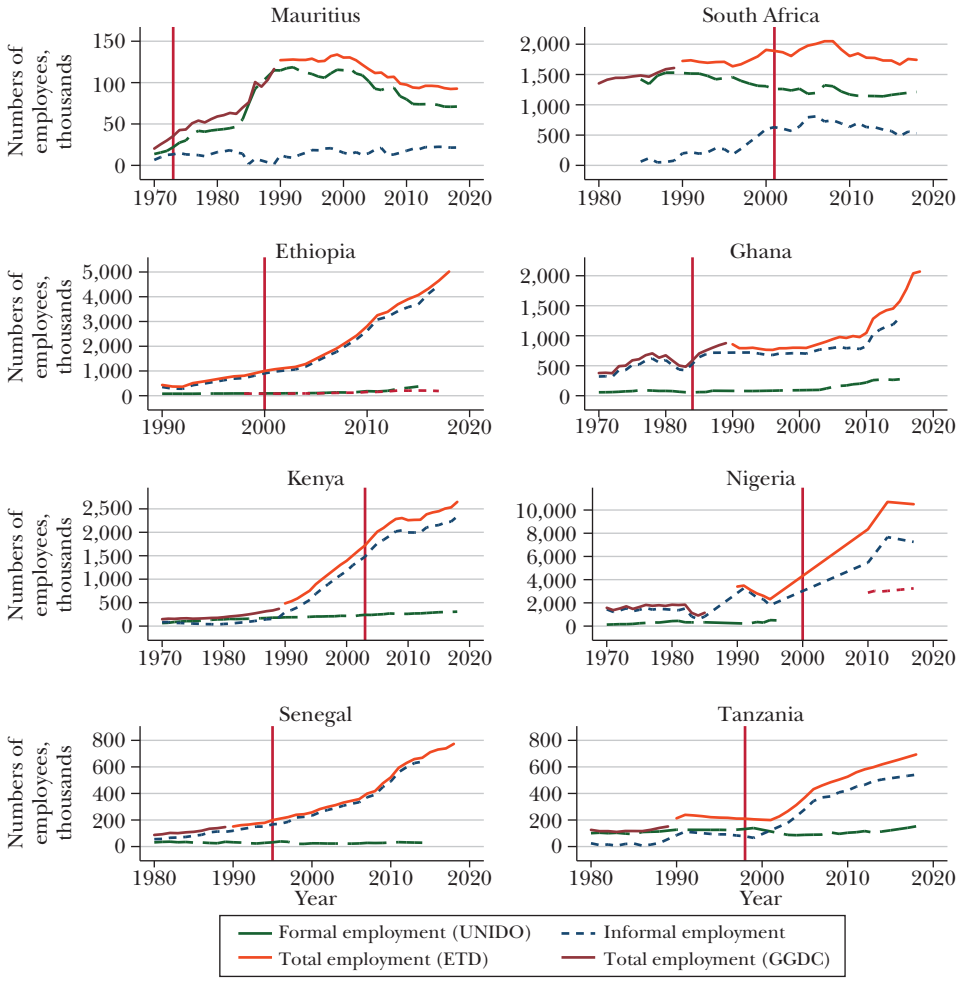

Source: CSA LMSM, ETD (de Vries et al. 2021), GGDC 10-sector (Timmer et al. 2015), NBS (2013, 2014, 2019), UNIDO Indstat2.
Note: The vertical red line indicates the start of the country's growth acceleration. Informal employment is calculated as the difference between GGDC (total) and UNIDO (formal) employment pre-1990, and ETD (total) and UNIDO (formal) employment from 1990 onward. We augment the UNIDO data on formal employment with census data for Ethiopia and Nigeria, indicated by the dashed red line. In Ethiopia the data from the Large and Medium Scale Manufacturing (LMSM) census covers two additional years (2016–2017), and this is the same source of UNIDO INDSTAT2. Nigeria does not have UNIDO employment data post-2000, so we instead measured formal employment in manufacturing using numbers from the national census for 2010–2017.

one can see a modest surge in employment in firms with less than ten employees starting in the mid-1990s South Africa's manufacturing sector.

Figure 3B
Manufacturing Employment Growth in Asia

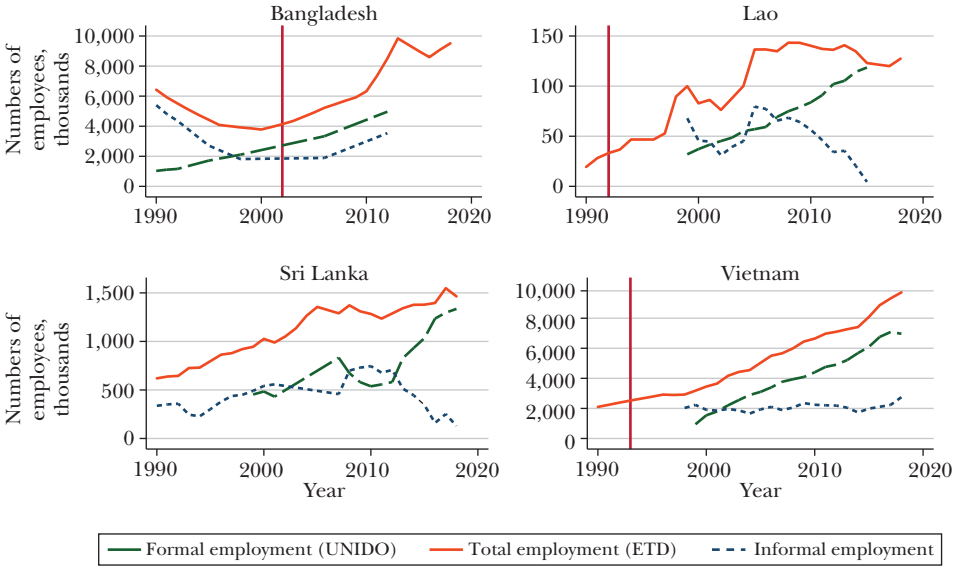

Source: ETD (de Vries et al. 2021), GGDC 10-sector (Timmer et al. 2015), UNIDO Indstat2.
Note: The vertical red line indicates the start of the country's growth acceleration. Informal employment is calculated as the difference between GGDC (total) and UNIDO (formal) employment pre-1990, and ETD (total) and UNIDO (formal) employment from 1990 onward. The Bangladesh series for formal employment post-2000 comes from two datapoints in 2006 and 2012.

Figure 3b shows starkly different patterns for four Asian late industrializers: Bangladesh, Lao People's Democratic Republic, Sri Lanka, and Vietnam. Again, the vertical lines represent the beginning of each countries' growth acceleration. Unlike the African cases, formal sector manufacturing employment in these four countries grows rapidly following the growth acceleration. Over these same periods and apart from Bangladesh, employment in small and informal firms remains relatively flat and even starts to decline slightly in Vietnam after 2005: for a description of the transition in Vietnam, see McCaig and Pavcnik (2017). Like the African countries, the Asian late industrializers were largely agrarian societies before the start of their growth accelerations. Although not shown, increases in formal manufacturing employment in these countries coincided with gradual declines in agricultural employment shares. These patterns combined with the evidence in Figure 2b are consistent with the idea that structural change in these countries was a result of positive supply shocks to the manufacturing sector.

Employment and Labor Productivity Growth in Formal Manufacturing Firms

The informal sector has apparently absorbed a large majority of the growth of manufacturing workers in African countries. Why are formal sector manufacturing

firms not creating more jobs in Africa? One possibility is that these firms are not very productive, making it unprofitable for them to expand. To explore this possibility, we take a closer look at the performance of the formal manufacturing sector following a two-pronged strategy. We begin by summarizing the recent evidence for Ethiopia and Tanzania documented in Diao et al. (2021). We then turn to UN Industrial Development Organization data to gauge whether the results for Ethiopia and Tanzania are indicative of what is happening in the other African countries for which we have data.

Diao et al. (2021) provide an in-depth analysis of the formal manufacturing sectors in Ethiopia and Tanzania. The core of their analysis rests on two newly created panels of manufacturing firms using formal firm census data, one for Tanzania covering 2008–2016 and one for Ethiopia covering 1996–2017. In both cases, the panel covers firms with ten employees or more. In both Tanzania and Ethiopia, large firms, exporters, and foreign firms all have significantly higher levels of labor productivity. These results are consistent with a very large theoretical and empirical literature on manufacturing firm performance in Africa and elsewhere: for example, see Bigsten and Söderbom (2006) for a summary of early research on the manufacturing sector in Africa based on the Regional Program for Enterprise Development and Harrison and Rodríguez-Clare (2010) for a summary of the literature on the impact of trade and foreign direct investment on firm performance.

However, the firm and sector-level results for Tanzania and Ethiopia suggest that the best-performing firms are not absorbing employment. In Tanzania, labor productivity growth in large firms is on the order of 8 percent per year, and 13 percent per year for large exporters. By contrast, labor productivity growth in firms that start with less than 50 employees ranges between negative 3 percent and zero. By contrast, employment growth in these small firms is as high as 13 percent while employment growth averages zero percent in firms with 50 or more employees. In Ethiopia, the picture is somewhat better although labor productivity growth is not as high as it is in Tanzania. Average sectoral employment growth in Ethiopia's large firms is around 6.5 percent per year while it is around 7.5 percent per year for firms with between 10 and 49 employees. (Although the employment growth rates may seem high in Ethiopia, it is important to remember that the country is starting from an extremely low base.) Labor productivity growth is around 5 percent per year for both small and large firms. To complete the analysis of Ethiopia, the authors examine a group of mechanized firms with less than ten employees using the Small-Scale Industries survey. Average annual employment growth among these firms is around 17 percent, while labor productivity growth is an average of zero, with considerable heterogeneity. The combined evidence for Ethiopia and Tanzania suggests that the economywide performance of employment and output in the manufacturing sector is not driven by large firms in the formal sector. Instead, it appears that the growing share of small less productive firms is dragging down the economywide productivity performance of the manufacturing sector.

To determine whether this interpretation applies to other African countries, we turn to the UN Industrial Development Organization data, which we use to compute

Table 1
Formal Manufacturing Sector Performance, 2000–2018

	Employment growth	*Real output per worker growth*	*Real value added per worker growth*	*Real exports growth*
Bangladesh[a]	0.066	0.042	0.000	0.100
Botswana	0.016	0.070	0.044	0.037
Cameroon[b]	0.038	0.082	0.119	0.089
Ethiopia[c]	0.065	0.047	0.041	0.093
Ghana	0.086	0.059	0.099	0.073
Kenya	0.019	0.065	0.067	0.056
Lao People's Dem Rep[d]	0.074	0.072	0.083	0.208
Lesotho[e]	0.065			0.031
Mauritius	–0.031	0.027	0.033	–0.012
Nigeria[f]	0.015			0.133
Senegal	0.025	0.021	0.023	0.030
South Africa	–0.005	0.018	0.014	0.025
Sri Lanka	0.049	0.024	0.019	0.024
United Republic of Tanzania	0.036	0.033	0.081	0.134
Vietnam	0.078	0.063	0.065	0.178

Notes: Employment, output per worker, and value added per worker series come from UNIDO INDSTAT2 database; the export series comes from the World Bank World Development Indicators (WDI); monetary values are deflated using the manufacturing Producer Price Index from WDI. The results for employment, output per worker, and value added per worker growth all cover the same years, while the results for export growth use any years for which data are available in the period. Unless otherwise indicated below, the countries all have 14–19 years of useable data post-2000. Detailed information on the data availability can be found in Appendix Table A1.

[a] Bangladesh only has two datapoints, in 2006 and 2012.

[b] Cameroon only has four datapoints, in 2000–2002 and 2008.

[c] We use census data from the LMSM survey for Ethiopia because the UNIDO data from 2014–2015 are estimates, and the census data extend past 2015 allowing us to include 2016–2017. The LMSM series match well to the UNIDO series pre-2010 and the LMSM is the source of the UNIDO data.

[d] Lao's export data is only available from 2010–2019.

[e] Lesotho only has employment data from 2001–2009, but not output or value-added data—we therefore only report employment and export (2000–2017) growth results.

[f] Nigeria does not have UNIDO INDSTAT2 data post-2000, so we estimated formal employment from a report released by Nigeria's National Bureau of Statistics that has employment statistics for 2010–2012 (NBS 2014); we extrapolated values up through 2017 using a linear trend. Because of the limited data availability in UNIDO, we only report employment and export (2001–2019) growth results.

growth in real output per worker, real value added per worker and real exports to gauge the health of the formal manufacturing sector. The results of this analysis are presented in Table 1. The data are organized into four columns: employment growth, growth in real output per worker, real value added per worker, and exports. The data coverage across the countries of Africa is admittedly sparse, but we do have information for the six countries with at least 100,000 workers in the formal manufacturing sector. We also include four Asian comparator countries: Bangladesh, Lao, Sri Lanka, and Vietnam.

In considering Table 1, first notice that growth in real value added per worker (or output) is positive in all the African countries. This is consistent with what we

reported for Ethiopia and Tanzania and suggests that the poor productivity performance of economywide manufacturing is not a result of subpar performance in the formal manufacturing sector. Second, apart from Ethiopia and Kenya, employment growth in formal manufacturing is considerably weaker in Africa than in Bangladesh, Lao, and Vietnam. This gap exists despite the fact the levels of manufacturing employment are considerably higher in the Asian comparator countries; for example, total formal manufacturing employment was around 300,000 in Ethiopia in 2017, while it was more than 6 million in Vietnam.

The upshot of these comparisons is that the formal manufacturing sectors in the African countries for which we have data appear to be performing reasonably well. This evidence is consistent with a large literature on the productivity of Africa's larger manufacturing firms.[2] But the formal manufacturing sector in Africa, unlike their counterparts in Asia, does not appear to be absorbing significant amounts of labor—at least not yet—except perhaps for Nigeria. Teal (2016) succinctly describes the issue in Ghana as follows: "It is the inability of larger firms, particularly those employing more than 100, to grow in numbers and employment that needs to be explained if the inability of Ghana to produce more productive jobs in its manufacturing sector is to be understood." As a step towards a deeper understanding of the issue, we use firm-level census data to compare the manufacturing sectors in Ethiopia and Vietnam.

Ethiopia versus Vietnam: A Case of the Missing Huge Firms

Ethiopia is the second most-populous country in Africa. Its government has pursued an aggressive industrialization strategy, which largely revolves around attracting investment in labor-intensive manufacturing for export. As of 2021, Ethiopia's strategy mirrors to a large degree the strategy pursued by Vietnam; as noted earlier, both countries have been heavily influenced by China's use of special economic zones.

In some ways, the manufacturing sectors of Ethiopia and Vietnam look similar. In both countries, annual average employment growth between 2000 and 2017 was between 7 and 8 percent. In both countries, employment growth was driven almost entirely by the entry of new firms, not the expansion of existing firms. Employment growth in both countries is more rapid in firms with between 10 and 49 employees at 7.5 percent in Ethiopia and 12.1 percent in Vietnam. Average annual labor productivity growth in Vietnam was 7.5 percent—higher but not that much higher than the 5.2 percent registered in Ethiopia.

There are high rates of entry and exit in both countries indicating a significant degree of dynamism. Capital-labor ratios are comparable in both, yet Ethiopia is much poorer than Vietnam—a point to which we will return. In both countries,

[2]For example, John Sutton in his Enterprise Maps argues that the reasonable performance of manufacturing firms serving domestic markets can be explained by the discipline imposed by having to compete with imports. *Made in Africa: Learning to Compete in Industry* (2012) demonstrates the relatively high productivity of large African firms compared to firms with less than 10 employees.

public-sector firms played a significant role in the early stages of industrialization. Public-sector firms are still involved in manufacturing in both countries today but by 2017, the employment share of public-sector firms was considerably lower at around 8 percent in Ethiopia and 5 percent in Vietnam (in 2017).

Of course, there are also some striking differences. The most notable is the sheer size of the formal manufacturing sector in Vietnam relative to Ethiopia. An 8 percent increase in formal sector manufacturing in Vietnam adds an additional 520,000 jobs while an 8 percent increase in formal sector manufacturing in Ethiopia adds around 28,000 jobs. If employment in Ethiopia's manufacturing sector continues to grow at 8 percent per year, it will take Ethiopia 38 years to catch up to the level of employment seen in Vietnam today. Part of this difference has to do with the fact that Vietnam had a much larger industrial base when it embarked on its strategy of export-led manufacturing development. In 1986, at the start of the "Doi Moi" reform period in Vietnam, the country already registered 2.5 million employees in the manufacturing sector. By contrast, when the Ethiopian government embarked on its reforms in the mid-2000s, the country had fewer than 100,000 workers in the formal manufacturing sector.

Like Ethiopia, the share of informal employment in Vietnam was very high at the onset of its reforms in 1986. Thirteen years after the start of reforms, the share of informal employment was still around 86 percent, although it was lower in the manufacturing sector. And between 1990 and 1995 the share of informal employment in Vietnam's manufacturing sector increased by 10 percentage points, peaking at around 58 percent in 1995 (General Statistics Office 2006). This growth in informal manufacturing sector employment is similar as what we observe in Africa today, albeit of a shorter duration. By 2009, the share of informal employment in Vietnamese manufacturing had fallen to 43 percent (McCaig and Pavcnik 2015). Today it stands at around 36 percent.

A striking difference between Ethiopia and Vietnam is the rapid expansion of foreign firms in Vietnam's manufacturing sector. In 1990, at the onset of Vietnam's reforms, there were less than 1,000 workers employed in foreign-owned firms. Between 1990 and 2000, employment in foreign-owned manufacturing enterprises grew at an annual average rate of 47.3 percent (General Statistics Office 2006). Employment growth in domestic private and state-owned enterprises paled in comparison at 3 percent and 2 percent respectively. Although employment growth in foreign-owned enterprises slowed down after 2000, it remains the dominant source of employment growth in Vietnam's manufacturing sector. Between 2000 and 2017, annual employment growth in foreign-owned manufacturing enterprises averaged close to 14 percent; employment growth in domestic private firms averaged 4 percent while it was −6 percent in state-owned enterprises. By 2017, the share of manufacturing employment in foreign-owned enterprises exceeded 60 percent in Vietnam.

By contrast, in 2017 the share of employment in foreign-owned enterprises in Ethiopia was less than 10 percent and the share of employment in domestic private firms exceeded 60 percent. However, prior to the onset of the pandemic, the landscape in Ethiopia was rapidly changing. In 2014, the Ethiopian Industrial Parks

Development Corporation (IPCD) was established to help promote exports and job creation primarily in the manufacturing sector. This first park, Eastern Industrial Zone, was established in 2012–2013, and as of 2020, it hosted 91 firms employing over 18,000 workers. Since then, 13 additional parks have been opened, which in total are home to 154 firms that employ about 93,000 workers. Indeed, 19 of these firms employ more than 1,000 workers.[3] Of these employees, 74 percent are engaged in apparel production. In absolute terms, the number of employees in the parks is still relatively small. However, since 2015, employment in industrial parks increased by 78 percent, comparable to what we saw in Vietnam over the period 1990–2000.

All but three of the new industrial parks are owned by the Ethiopian government: two of the other owners are Chinese and one is Bangladeshi. Almost all the firms in these parks are foreign-owned; 79 percent are owned by Asian investors, with China in the lead at 66 percent, 12.4 percent are owned or partially owned by Ethiopian investors, and 8.6 percent are owned by EU and UK investors.

Within Ethiopia's industrial parks, 83 percent of employment is in firms with more than 1,000 employees. Women account for 74 percent of total employment; firm managers tend to hire young female workers who have completed a 10th grade education. These women typically work eight hours a day for six days a week plus an average of two hours overtime each day bringing the average hours worked per week to 60. Ultimately, the sustainability of these parks depends in part on working conditions. Using a phone survey of firm managers, Meyer, Krkoska, and Maaskant (2021) find that the base salary in most of the firms in these industrial parks exceeds the cost of basic needs as measured by the local poverty line. When bonuses, overtime pay, incentive payments, and in-kind benefits are included, total compensation is roughly four times the cost of basic needs. However, they also find that 21 percent of firms in the industrial parks report paying a base wage below the local poverty line. Based on the information in the report, it is not clear how far below the poverty line these wages fall and what happens to total compensation for these workers when non-wage benefits are included.

A unique feature of many of the foreign firms in Vietnam is their sheer size. For example, in 2017, there were 756 foreign-owned firms in Vietnam with more than 1,000 employees; in Ethiopia, this number is only 21. Moreover, there were 125 foreign firms with more than 5,000 employees in Vietnam while there are no firms of this size in Ethiopia's manufacturing sector. This difference does not apply only to foreign firms, although the "huge" firms are more prevalent among foreign-owned enterprises and in the labor-intensive sectors of apparel and footwear. We call this the case of the missing huge firms.

A final point worth mentioning is the productivity growth of the very small firms in Vietnam. The census of establishments in Vietnam covers firms with less than ten

[3]A list of Ethiopia's industrial parks, with some basic data on when they started, number of firms and workers, and share of workers who are female and/or in the garment industry, is available in the online Appendix available with this article at the *JEP* website. Our evidence is consistent with that of Meyer, Krkoska, and Maaskant (2021).

employees. Productivity growth in these very small Vietnamese firms was around 9 percent per year between 2000 and 2017; the corresponding growth rate in Ethiopia is imprecisely estimated at zero percent per year. We do not have an explanation for this difference. The firms with less than ten employees in the Vietnamese enterprise survey are all registered; we don't know the registration status of the small-scale industries in Ethiopia, but in 2014, more than 85 percent of these firms had a license. In any case, given the prevalence of small manufacturing establishments in Africa, this difference seems worth investigating.

The evidence discussed so far suggests that employment growth in Africa's small and informal manufacturing firms is considerably more rapid than employment growth in manufacturing firms which employ ten or more workers. However, labor productivity growth in the larger formal firms is reasonably strong. This mixture of outcomes is not all bad news. Productivity growth in formal sector firms is indicative of an advancing technology frontier. The rapid growth in small manufacturing firms is indicative of entrepreneurial spirit; indeed, anyone who has spent time in an African country understands the incredible ingenuity of African entrepreneurs. Nevertheless, if African governments want to expand manufacturing exports and employment opportunities for those who prefer formal wage work, they will need to grow their formal manufacturing sectors. How might this happen? In the next section, we consider some of the opportunities.

Opportunities

African Continental Free Trade Area

The African Continental Free Trade Area was founded as a free trade area in 2018 with 54 of the 55 African Union nations as signatories (the exception being Eritrea). To date, 36 states have ratified the agreement, and trade under the agreement officially commenced at the start of 2021.[4] Its key functions include progressively eliminating tariffs on intra-African trade (with alternate timelines for implementation based on countries' income status), implementing rules-of-origin, monitoring and eliminating nontariff barriers, as well as establishing an online negotiating forum, a digital payment system, and the African Trade Observatory. Arguably, the largest potential gains of the African Continental Free Trade Area are dynamic and arise mainly from access to larger markets and economies of scale in production. Another less tangible but potentially important benefit of the agreement is political. Most of Africa's economies are relatively small; this limits their bargaining power relative to the rest of the world in international forums such as the World Trade Organization. Regional integration has the potential to change this dynamic.

[4] A useful source of background information on the African Continental Free Trade Area is the South-Africa based nonprofit Trade Law Center often called "tralac." For details on the agreements and membership behind the African Continental Free Trade Area, a starting point is https://www.tralac.org/resources/by-region/cfta.html (accessed June 1, 2021).

To what extent might the African Continental Free Trade Act catalyze employment and export growth within Africa? A comparison between Ethiopia and Tanzania is instructive (based on Diao and McMillan 2019). Both countries showed a steady upward trend in the value of exports in the two decades leading up to the pandemic. However, Ethiopia's exports go almost exclusively to countries outside of Africa, which is consistent with what we know about the Ethiopian governments' push to include Ethiopia in global value chains. By contrast, a large majority of Tanzania's manufacturing exports go to other countries in Africa. What is perhaps surprising is that Tanzania's export volume and growth from 1998 through 2017 are more than double that of Ethiopia. After all, the government of Ethiopia has aggressively incentivized manufacturing for export with its industrial parks and tax incentives while as far as we can tell, the Tanzanian government has been much more laissez-faire.

One reason for the differential export performance is that exports from Ethiopia and Tanzania are very different. The top 50 products exported from Ethiopia account for 65 percent of Ethiopia's manufacturing exports; 84 percent of the top 50 products are classified as textiles including leather and footwear. More than 85 percent of Tanzania's export products are resource-intensive, with 50 percent classified as agro-processed goods and another 35 percent classified as material-intensive products. The agro-processed goods consist of items like bottled juices, cooking oils, and packaged flour while the resource intensive products consist of items such as wood products and furniture, household articles made from plastic materials such as buckets, washbasins, chairs and clothing hangers, and construction materials such as cement, glass, and ceramic products. In sum, agro-processed and resource-intensive goods account for 68 percent of total manufacturing exports from Tanzania (Diao and McMillan 2019). Intra-African trade in manufactured exports, like that occurring in Tanzania, has also been documented elsewhere (Hallward-Driemeier and Nayyar 2017). In many ways, this trend bodes well for the African free trade act. African countries still import much of their food; the evidence from Tanzania suggests that some of this demand could be met by Africa-based agro-processors.

But while Tanzania's export performance is impressive, there has been little employment growth in Tanzania's formal manufacturing sector. The issue (to which we will return in the next section) is the high capital intensity of resource-based manufacturing. Nonetheless, agro-processing has the potential to create jobs and wealth indirectly for logistics and packaging companies, restaurants and hotels, agricultural input suppliers, and so on. For perspective, Sexton et al. (2015) estimate using input-output tables that California's food and beverage processing sector directly accounted for around $25 billion in value-added and 198,000 jobs in 2012. However, the indirect benefits associated with the food and beverage industry were far greater and include an additional $57 billion in value added and another 562,000 jobs. The extent to which these sorts of linkages can generate large scale job creation in Africa is an open question.

One potential challenge for the African Continental Free Trade Area involves the rule-of-origin provisions, which define the products that are eligible for preferential tariff treatment. To qualify, a product must be wholly obtained or substantially transformed within an African country that is a member of the agreement (Signé and Madden 2020). The goal of rules-of-origin is to prevent trade "deflection": that is, a situation where exports arrive from outside Africa's free trade area in any one country and are then re-exported to other African countries under the preferential rules. On the other side, Signé and Madden (2020) find that low intra-African trade volumes mean that the fixed cost of compliance with rules-of-origin may be burdensome for many traders. Moreover, underdeveloped African value chains make it difficult and costly for African exporters to source intermediate inputs from domestic or regional sources. For these reasons, overly strict rules-of-origin could mean that firms or products which need inputs from outside Africa are effectively excluded from the free trade agreement.

The In-Between Sector

We have documented the rapid growth in small and informal firms in the manufacturing sector of many African countries along with their relatively low levels of labor productivity. However, there is a great deal of heterogeneity in the productivity of these small firms. For example, Diao et al. (2018, 2020) show that in 2010, 15 percent of the small firms in Tanzania have labor productivity higher than economywide manufacturing labor productivity. Also, more than 50 percent of firm owners report that they would not leave their business for a full-time salaried job. While 15 percent may seem high in comparison to what others have found, it is only half the share of "gung-ho" entrepreneurs (30.4 percent) identified by Banerjee et al. (2019); these entrepreneurs are those whose businesses benefitted from access to microfinance in India. Following Lewis (1979), we describe this group of firms as "the in-between firms" to signal their status and performance as somewhere between formal (most productive) and informal (least productive) manufacturing firms.

How can policymakers support the productive small firms while not encouraging over-investment in unproductive activities? The importance of this question cannot be overstated. These small firms are considerably more likely to use labor-intensive technologies and will thus be an important source of employment in African countries (and elsewhere) for years to come.[5] One popular strategy for spurring growth in promising ventures in both the developed and developing world is business plan competitions, which seek to identify and encourage entrepreneurs with growth

[5]A related question which we do not take up here is the extent to which support for small firms has encouraged the proliferation of unproductive businesses (for example, Martin, Nataraj, and Harrison 2017). However, in the context of sub-Saharan Africa where wage employment is scarce, it is difficult to think of this as a misallocation of resources. Instead, we might think of indiscriminate programs targeted at small businesses as a kind of social safety net. A final consideration is the extent to which support in the form of financial resources displaces or substitutes for alternative sources of financing (Fafchamps et al. 2014).

aspirations by helping them to develop a detailed plan and then providing financing or in-kind benefits (such as training) to those deemed most likely to succeed.

YouWin! is an example of a successful, large-scale nationwide business plan competition initiated in Nigeria in 2011. The top-scoring plans overall and in each region won awards of roughly $50,000, and then out of the 1,900 plans that were semifinalists in the competition, 700 were selected at random as winners. An evaluation of this competition using the underlying random variation by McKenzie (2017) tracked winners over five years and shows that winning firms had higher survival, profits, sales, and employment. Winning also increased the likelihood that a firm has more than ten employees by 20 percent. For present purposes, what is important about the business competition is that manufacturing was the second most common sector for new firms, comprising 13 percent, and third most common for existing firms, comprising 14 percent of winners. The types of products being manufactured by firms in the competition are very heterogeneous, and include processed food products, books and media, metal products, chemicals and detergents, and a range of other items.

However, a follow-up paper by McKenzie and Sansone (2019) highlights the difficulties of picking the outright competition winners. They compare the relative performance of three approaches to predicting outcomes from the YouWin! competition: business plan scores from judges with business experience, simple ad hoc models used by researchers who study entrepreneurship, and machine learning approaches. The results are sobering: i) business plan scores from judges are uncorrelated with performance; ii) gender, age, ability and business sector do have some predictive power (education might have an effect also, but applicants were required to apply using the internet which likely screened out individuals with little education); iii) machine learning methods did not enhance predictive power; and iv) the overall predictive power of all approaches is very low, highlighting the fundamental difficulty of picking competition winners. This of course does not mean that the program had no impact; it just means that it could have had a similar or even stronger impact had it been possible to identify promising projects with greater accuracy. Business plan competitions appear to have been successful in several other African countries, although not on the same scale: for a description of results for Ethiopia, Tanzania, and Zambia, see Fafchamps and Quinn (2017). In fact, business plan competitions in other African countries are in the works.

While such competitions seem useful, there is no substitute for business-to-business linkages for raising productivity and employment in small firms. The "putting out system" or sub-contracting has been important since the first industrial revolution and is still common in many parts of the world including China and India; this small enterprise sector is often referred to as the "cottage industry." A detailed description of the way this worked in Taiwan several decades ago highlights the sophisticated nature of these enterprises in one village in rural Taiwan (Niehoff 1987). In that setting, all but one of the enterprises were self-financed with no formal guidance apart from production and export brokers; these brokers were identified as an important source of ideas regarding types of

factories and commodities suitable for household entrepreneurship. We have some evidence of linkages between domestic and foreign firms in Ethiopia (Abebe, McMillan, and Serafinelli forthcoming). However, survey data from Ethiopia indicate that less than 5 percent of large firms do any sub-contracting. Overall, we have very limited knowledge about linkages between small and large firms in Africa.

The African Center for Economic Transformation (ACET) is taking business-to-business linkages seriously. In early 2021, ACET launched the ACET Private Sector Development Program. The program has two symbiotic objectives. The first—at the macro level—is to promote evidence informed private-sector friendly public policies and regulations to strengthen the ecosystem for small- and medium-sized enterprises. The second—at the micro level—is an incubator program designed to integrate small- and medium-sized enterprises in the manufacturing space into local, regional and global value chains. The incubator phase of the program assists firms in all aspects of business from input sourcing to management training; ACET has partnered with firms such as Price Waterhouse Coopers and EVC Africa Ltd to provide this assistance (https://acetforafrica.org/psd/acet-business-transform/). ACET is currently piloting the incubator program with 10 businesses in Ghana with plans to expand (also to other countries) if the pilot is successful. The 10 firms range in size from about 6 to 46 employees and operate in agro-processing, cosmetics, construction, electric vehicles and plastic waste recycling (Brown and Odoom 2021).

The importance of small manufacturing firms in Africa cannot be overstated; they are considerably more likely to use labor-intensive technologies and will thus be an important source of employment in African countries for years to come. It therefore makes sense to devote more time and energy to understanding their dynamics, especially the prospects for outsourcing to these firms.

Our lack of information about small firms in developing economies is not unique to Africa (Li and Rama 2015). This is not surprising given the large share of informality among small firms, which makes it difficult to come up with adequate sampling frames. Nevertheless, it has not prevented generalizations about small firms based on tiny samples (for example, La Porta and Shleifer 2016). To see why this can be problematic, consider the work of Bassi et al. (2020) on three sectors of urban manufacturing in Uganda with a mix of "small" and large firms. The authors show that the way in which one defines the borders of a small firm matters a lot. For example, the active rental market in carpentry equipment allows firms which manufacture two-panel doors to achieve scale, collectively turning a seemingly unproductive one-person business into a productive one, where the productivity gains come through mechanization.

Value Addition in Natural Resources

Africa is home to immense reserves of minerals that could help to spur its industrialization. Indeed, the US economy became the world's largest extractor and exporter of natural resources at precisely the same time it became the largest

industrial power via heavy utilization of resources (Wright 1990). In the modern economy, natural resources are a key part of global value chains. For example, tantalum is used in the production of cellphones, DVD players, laptops, and gaming devices. By building backward and forward linkages and developing associated industries, African countries can tap into Africa's advantage as the dominant supplier of these resources.

Industrialization prospects depend critically on how natural resources are managed. Africa—with the exceptions of Botswana and South Africa—has had a bad record on this front. But over time, governance has improved and the domestic private sector is considerably more robust and mature. For example, the Dangote Group (a Nigerian multinational founded by Aliko Dangote) is setting up the largest oil refinery in Africa. This sort of private sector involvement can play a useful role in disciplining African governments and holding them accountable. Oil refining in Nigeria could be a significant boon to a continent that has mostly exported crude oil, and then imported refined oil.

Green Manufacturing

Green energy involves production and uses including both deploying renewable energy sources like compressed natural gas, wind, solar, and biomass, along with achieving higher energy efficiency in operations. An International Renewable Energy Agency (2020) report highlights Africa's substantial endowment in renewable resources like biomass, geothermal, hydropower, solar, and wind power. Africa's estimated solar power generation potential greatly exceeds that of other regions (Kabir et al. 2018; Schwerhoff and Sy 2020).

As countries around the world seek to reduce greenhouse gas emissions, Africa's comparative advantage in clean energy production could be a boon to both domestic and foreign manufacturing firms in terms of cost reductions and compliance with international climate agreements. A transition to green manufacturing could also make light manufacturing in Africa more cost competitive, thereby increasing employment. In the example of South Africa, Winkler and Black (2021) argue that part of South Africa's unemployment problem is a result of policies that encouraged mining and heavy industry. In those industries, emissions are especially high and employment low due to a reliance on coal-powered energy and heavily subsidized capital-intensive factories. They argue that an employment-intensive, low energy, light manufacturing industrial policy is more aligned with South Africa's real comparative advantage.

More broadly, low-cost renewable energy sources could have an important impact on electricity access in Africa. Low-cost electrification has the potential to raise the productivity of small and large manufacturing firms alike. This could have an especially large effect in rural areas where electrification rates average only 17 percent (Altenburg and Assmann 2017). Of course, encouraging green energy would not alleviate some of the problems associated with energy distribution and pricing (for example, see Burgess et al. 2020; Jack and Smith 2015).

Pharmaceuticals

Africa relies heavily on imports of health commodities, with most countries importing between 70 and 90 percent of pharmaceuticals consumed (Conway et al. 2020). The only African countries with relatively sizable pharmaceutical production industries are Kenya, Nigeria, and South Africa, and the entire continent accounts for just 3 percent of global medicine production. Furthermore, almost all manufacturing capacity in Africa produced generic medicines, for which firms import the active pharmaceutical ingredients (Banda, Wangwe, and Mackintosh 2016). In 2014, according to a 2021 report on "Vaccine manufacturing in Africa," funded by the UK government, Africa produces just 1 percent of the vaccines it administers, with the other 99 percent being imported.[6] Africa's pharmaceutical industry employed 250,000 workers and created $6.8 billion in gross value added, significantly lower than other global regions—for example, Latin America employed 466,000 workers and created $24.6 billion in gross value added in pharmaceuticals manufacturing (IFPMA 2017).

There has been increasing focus on this need in the last decade; for example, the African Union has a Pharmaceutical Manufacturing Plan for Africa, which aims for countries to develop their capacities in pharmaceuticals production, innovation, and research and development (Banda, Wangwe, and Mackintosh 2016). However, the COVID-19 pandemic has compounded the need for local pharmaceutical industries. According to the World Health Organization (2021), only Senegal and South Africa currently have the capacity to produce vaccines, although the Ethiopian government is currently building an industrial park dedicated to pharmaceuticals (https://www.ipdc.gov.et/service/parks/12) and Nigeria's relatively large pharmaceuticals industry has had the capacity to produce vaccines in the past.

A silver lining of the pandemic may be that it focuses the international community on assisting African governments in shoring up capabilities in the pharmaceuticals industry. There is some evidence that this may be happening. In January 2021, ten African countries along with China introduced a resolution at the World Health Organization calling for greater local production of medicines, which was co-sponsored by 100 countries including all 54 countries in Africa. The Director-General of the World Trade Organizations, Ngozi Okonjo-Iweala, is urging members to facilitate the transfer of technology for vaccine production to more countries. A spokesperson for the Switzerland-based public-private vaccine alliance GAVI has stated that the United States supports manufacturers to transfer not only intellectual property but also know-how in a bid to boost global production (Zarocostas 2021).

[6]For the presentation slides for this report, see https://www.dcvmn.org/IMG/pdf/20210316_vx_mf_africa_dcvmn_briefing_vpre-read.pdf, accessed June 1, 2021.

Challenges

Two of the common concerns raised about the prospects for manufacturing in African nations are the business climate and the risk of political instability. These issues are real, but we believe they are much less important than a more fundamental issue: manufacturing around the world shows a pattern of rising capital intensity, which raises questions about the prospects for manufacturing in countries with a comparative advantage in low-skilled labor.

Rising Capital Intensity of Manufacturing

Technological innovation in manufacturing in recent decades has favored capital over labor. At the same time, the spread of global value chains and increased openness to trade have had the effect of homogenizing technology around the world (Rodrik 2018). As a result, trade integration tends to reduce the employment intensity of manufacturing production in developing countries (Sen 2019). Thus, participation in global value chains tends to increase labor productivity, but not employment (Pahl and Timmer 2020). Indeed, Pahl et al. (2019) show that technological change in formal manufacturing has led to employment declines in Kenya, Senegal, and South Africa. This confluence of events makes it more difficult for low-income African countries to gain a foothold in formal manufacturing—even in their own domestic markets.

Using firm-level data from Ethiopia and Tanzania, Diao et al. (2021) show that the capital intensity of formal manufacturing in both countries far exceeds economywide capital intensity. This is especially true of the larger most productive firms, where capital intensity approaches (or exceeds) levels observed in the Czech Republic, a country with per capita income around 20 times higher. High levels of capital intensity (and possibly of skill intensity as well, although the study does not measure that) appear to be an important reason behind the poor employment performance of Africa's most productive manufacturing firms. This creates a conundrum for countries in Africa: competing with established producers on world markets is only possible by adopting technologies that make it harder to generate significant employment.

In Figure 4, we compare the evolution of the capital intensity of manufacturing in Ethiopia, Vietnam, and the United States. Overall manufacturing capital intensity in Ethiopia and Vietnam look similar, with comparable levels and growth trajectories in capital-labor ratios from 2006 to 2017. The key distinction is that Ethiopia's capital-labor ratio in manufacturing is four times the capital intensity for its economy as a whole capital (Diao et al. 2021); in Vietnam, the two were roughly equal. In US manufacturing, capital intensity is actually lower than economywide capital intensity. In short, the prevailing capital intensity of manufacturing does not play to Ethiopia's comparative advantage in relatively low cost and low skilled labor.

Might certain sectors of manufacturing be less capital-intensive? The mass produced ready-made garment industry is the least capital-intensive industry within manufacturing, and it still employs large numbers of workers across the

Figure 4
Capital Labor Ratios, Total Manufacturing and Apparel

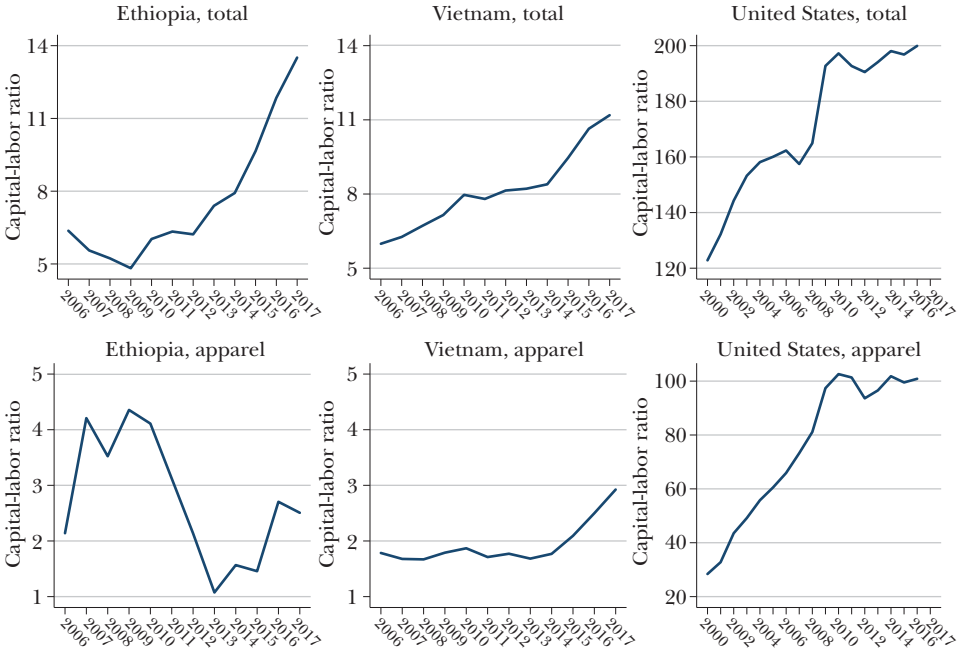

Note: Capital-labor ratios were measured in thousands of 2012 US dollars and are measured as annual averages weighted by sector employment shares. The top row is total manufacturing and the bottom row is apparel. For the Ethiopia apparel figure, we limit the sample to firms in the sector for at least five years from 2006 to 2017, to reduce the impact of entry and exit on the trend (there are only 43 firms in Ethiopia in the apparel sector per year on average, compared to 2,606 in Vietnam). These figures are produced using a cleaned panel of firms for Ethiopia (LMSM) and Vietnam (enterprise survey), and the United States data come from the NBER-CES Manufacturing Database (2012 NAICS version). NBER-CES data are organized at the NAICS 6-digit level so with Ethiopia and Vietnam we aggregate the firm-level data to the ISIC 4-digit level so that we can use similar methods. The 2011 data for Ethiopia is missing firms from apparel, so that value is replaced through interpolation.

globe.[7] In 2019, the ready-made garment industry in the Asia and the Pacific region employed an estimated 65 million garment sector workers or 75 percent of all ready-made garment workers worldwide (ILO 2020). More than half of these workers are in China and Vietnam, two countries where wages are rising. To get a sense for technological changes in the ready-made garment industry, Figure 4 also examines trends in capital-labor ratios in the apparel industry in Ethiopia, Vietnam, and the United States; we include the US as a benchmark for where the industry might be headed.

[7]Authors' calculations based on firm-level data for Ethiopia, Tanzania, Vietnam, the NBER productivity database and, for European countries, the EU KLEMS database (van Ark and Jäger 2017).

In 2017, capital-labor ratios in the apparel industries of Ethiopia and Vietnam are a fraction of what they are for total manufacturing. (Contrary to popular perception, this pattern does not hold for textiles or leather.) Moreover, in Ethiopia and Vietnam, capital-labor ratios in apparel are not rising the way they are in total manufacturing. There is an uptick in apparel manufacturing's capital intensity in Vietnam between 2015 and 2017, but the overall level remains at less than one-third of total manufacturing's capital intensity. In the US economy, although capital intensity in apparel far exceeds that in Ethiopia or Vietnam (as shown by the difference in the vertical axis), capital-labor ratios increased substantially in the apparel industry between 2000 and 2010 but then leveled off. While the apparel products produced in the United States are likely not comparable to those produced in Ethiopia and Vietnam, the leveling of capital intensity in the apparel industry is indicative of a stall in technological change in this industry. One reason for this appears to stem from the difficulty associated with mechanizing the sewing process.[8] At least for now, the ready-made garment industry appears to offer some opportunity for certain nations in Africa.

The Business Environment

Much has been made of the poor business environment in Africa and business environment does matter, of course. But as nations across Asia have shown, where there are profits to be made, businesses find a way to work around business environment problems. Similarly, despite the business environment in Africa, formal manufacturing firms have performed well in terms of productivity growth (Diao et al. 2021).

Indeed, measuring the business environment by the World Bank Doing Business index, many countries of Africa compare favorably to countries of Asia that have experienced rapid growth. In 2013, for example, Ghana ranked 27 countries ahead of Vietnam in the Doing Business indicators. According to these indicators, it was considerably easier to get credit in in Ghana than in Vietnam, paying taxes was less of a hassle, insolvency was much more quickly resolved, and access to electricity was less problematic. In terms of how well investors are protected, there was a 40-point difference between the two countries in favor of Ghana (McMillan, Rodrik, and Sepulveda 2017). A comparison between the rankings of countries in Africa and those of countries in Asia with established bases in manufacturing for the year 2019 offers several similar examples. Rwanda ranks 40 points ahead of Vietnam at 29, Mauritius and Kenya are also ranked ahead of Vietnam at 21 and 61 respectively. Seventeen African countries rank ahead of Cambodia. Bangladesh has five million garment workers (ILO 2020), but out of 48 countries in Africa only eight countries are ranked below Bangladesh and seven of these countries are at war. Nigeria is ranked 30 points ahead of Bangladesh.

[8] Interview with Willy Shih, Robert and Jane Cizik Professor of Management Practice, Harvard Business School, May 7, 2021.

An important aspect of doing business not covered by the World Bank's Doing Business index is the ease and cost of international travel. For example, Campante and Yanagizawa-Drott (2018) find that air links increase business links and that the movement of people fosters the movement of capital; the advent of just-in-time manufacturing has also raised the importance of short-term air shipping. In a novel study of traders in Lagos, Nigeria, Grant and Startz (2021) use Nigeria's 2016 ranking in Henley and Partners Visa Restrictions Index to motivate estimating the welfare gains associated with relaxing travel restrictions between China and Nigeria. A look at these rankings for 2021 reveals similar patterns to what we observed with the Doing Business Indicators rankings. Travel is less restrictive for Ethiopians and Nigerians than it is for Bangladeshi citizens, and 27 African countries are ranked ahead of Vietnam, which ranks 88th along with Chad and Mali.

Political Instability

Political instability, with its associated violence and uncertainty, disrupts markets and growth. For example, Ksoll, Macchiavello, and Morjaria (2016) estimate the impact of electoral violence on cut flower exporters in Kenya and find that this violence induced a large negative supply shock reducing exports by around 50 percent. Worker absence was responsible for much of the reduction; larger firms and firms with direct contractual relationships suffered somewhat smaller losses. Evidence from Ethiopia indicates significant declines in manufacturing activity associated with political unrest in 2015 and 2016 (Abreha, Maemir, and Sánchez-Martín 2021).

Prior to 2017, countries in sub-Saharan Africa had made significant progress toward democracy and political stability. On the Polity 2 scale—ranging from –10 (hereditary monarchy) to +10 (consolidated democracy)—the population weighted average for 46 African countries rose by 10 points from –5 to +5 between 1990 and 2016 (for details, see Figure A.1 in the online Appendix). From 1990 to 2008, the incidence of civil wars in Africa dropped from 18 to 8. But currently the situation seems to be taking a turn for the worse. This year alone, there were five coups in Africa, more than in any other year over the past two decades (as reported by Munshi and Schipani 2021) and the ongoing civil war in Ethiopia has shuttered factories in the Tigray region. These may be temporary setbacks on the road to progress. However, there is some evidence that climate change and religious extremism are increasing the incidence of conflict in Africa (McGuirk and Nunn 2020).

Conclusion

Manufacturing has an important role to play in the development of Africa. Some patterns are clear. Manufacturing employment has increased at a rapid pace since 2000. Labor productivity growth in Africa's large manufacturing firms appears healthy but employment growth in these firms has been disappointing. There is a concern that manufacturing is becoming more capital-intensive, and thus may not be

as powerful a direct creator of jobs as in the past. In the ready-made garment sector, there still seems to be an opportunity for considerable employment expansion, at least for now. However, the processing of natural resources or agricultural products which are both abundant in Africa has always been capital intensive. The evidence from Tanzania clearly shows that while this type of manufacturing can contribute significantly to value added and export growth, it does little for employment. But even if the capital intensity of manufacturing precludes huge employment gains, the indirect employment gains associated with manufacturing could be large. Moreover, the managerial and logistical capabilities of large international manufacturing firms that have operations in Africa could be transferred to other activities through worker turnover or informal networks (Newman et al. 2016; Abebe, McMillan, and Serafinelli forthcoming).

A pattern that stands out in the African context is the rapid growth of small and informal firms in the manufacturing sector. On the one hand, this is exemplary of Africa's entrepreneurial spirit. On the other hand, average productivity growth in these firms appears to be a drag on labor productivity growth in manufacturing. However, there is substantial productive heterogeneity among these small firms. Integrating some of the more productive small firms into domestic value chains could have large payoffs.

References

Abebe, Girum, Margaret S. McMillan, and Michel Serafinelli. Forthcoming. "Foreign Direct Investment and Knowledge Diffusion in Poor Locations: Evidence from Ethiopia." *Journal of Development Economics.*

Abreha, Kaleb G., Woubet Kassa, Emmanuel K.K. Lartey, Taye A. Mengistae, Solomon Owusu, and Albert G. Zeufack. 2021.*Industrialization in Africa: Seizing Opportunities in Global Value Chains.* Washington, DC: World Bank Group.

Abreha, K., H. Maemir, and M.E. Sánchez-Martín. 2021. "Productivity Dynamics and Drivers in the Ethiopian Manufacturing: Some New Evidence." Background paper for the upcoming World Bank Ethiopia Country Economic Memorandum.

Altenburg, Tim, Claudia Assmann, Dani Rodrik, Emilio Padilla, Stefan Ambec, Michela Esposito, Akexander Haider et al. 2017. *Green Industrial Policy: Concept, Policies, Country Experiences.* Geneva, Bonn: UN Environment, German Development Institute/Deutsches Institut für Entwicklungspolitk.

Banda, Geoffrey, Samuel Wangwe, and Maureen Mackintosh. 2016. "Making Medicines in Africa: An Historical Political Economy Overview." In *Making Medicines in Africa*, 7–24, edited by Maureen Mackintosh, Geoffrey Banda, Paula Tibandebage, and Watu Wamae. London: Palgrave Macmillan.

Banerjee, Abhijit, Emily Breza, Esther Duflo, and Cynthia Kinnan. 2019. "Can Microfinance Unlock a Poverty Trap for Some Entrepreneurs?" NBER Working Paper 26346.

Bart van Ark, and Kirsten Jäger. 2017. "Recent Trends in Europe's Output and Productivity Growth Performance at the Sector Level, 2002–2015." International Productivity Monitor 33 (Fall).

Bassi, Vittorio, Raffaela Muoio, Tommaso Porzio, Ritwika Sen, and Esau Tugume. 2020. Achieving Scale Collectively. https://ssrn.com/abstract=3674922.

Bigsten, Arne, and Måns Söderbom. 2006. "What Have We Learned from a Decade of Manufacturing Enterprise Surveys in Africa?" *World Bank Research Observer* 21 (2): 241–65.

Brown, Edward K., and Charles Odoom. 2021. Zoom Interview, December 15, 2021.

Burgess, Robin, Michael Greenstone, Nicholas Ryan, and Anant Sudarshan. 2020. "The Consequences of Treating Electricity as a Right." *Journal of Economic Perspectives* 34 (1): 145–69.

Campante, Filipe, and David Yanagizawa-Drott. 2018. "Long-Range Growth: Economic Development in the Global Network of Air Links." *Quarterly Journal of Economics* 133 (3): 1395–1458.

Central Statistical Agency 1996–2017. Medium and Large Scale Manufacturing Industries Survey, 1996–2017. Addis Ababa, Ethiopia.

Central Statistical Agency. 2014. "Nigerian Manufacturing Sector: Summary Report 2010-2012." https://www.nigerianstat.gov.ng/pdfuploads/Manufacturing%20sector%202010-2012.pdf.

Central Statistical Agency. 2019. "National survey of micro small & medium enterprises (MSMEs) 2017." https://smedan.gov.ng/downloads/.

Conway, Michael, Tania Holt, Adam Sabow, and Irene Yuan Sun. 2020. "Should Africa Make Its Own Drugs?" *McKinsey & Company,* January 10. https://www.mckinsey.com/industries/public-sector/our-insights/should-sub-saharan-africa-make-its-own-drugs.

De Vries, Gaaitzen, Linda Arfelt, Dorothea Drees, Mareike Godemann, Calumn Hamilton, Bente Jessen-Thiesen, Ahmet Ihsan Kaya, Hagen Kruse, Emmanual Mensah, and Pieter Woltjer. 2021. "UNU-WIDER Technical Note." The Economic Transformation Database (ETD): Content, Sources, and Methods. https://doi.org/10.35188/UNU-WIDER/WTN/2021-2 (accessed May 4, 2021).

Diao, Xinshen, Mia Ellis, Margaret S. McMillan, and Dani Rodrik. 2021. "Africa's Manufacturing Puzzle: Evidence from Tanzanian and Ethiopian Firms." NBER Working Paper 28344.

Diao, Xinshen, Kenneth Harttgen, and Margaret McMillan. 2017. "The Changing Structure of Africa's Economies." *World Bank Economic Review* 31 (2): 412–33.

Diao, Xinshen, Josaphat Kweka, and Margaret McMillan. 2018. "Small firms, structural change and labor productivity growth in Africa: Evidence from Tanzania." World Development 105 (2018): 400–415.

Diao, Xinshen, Josaphat Kweka, Margaret McMillan, and Zara Qureshi. 2020. "Economic Transformation in Africa from the Bottom Up: New Evidence from Tanzania." *World Bank Economic Review* 34 (S1): S58–62.

Diao, Xinshen, and McMillan, M., 2019. "Manufacturing in Ethiopia and Tanzania: Challenges and Opportunities." In *The Annual Proceedings of the Wealth and Well-Being of Nations, 2019.* Beloit College Press.

Diao, Xinshen, Margaret McMillan, and Dani Rodrik. 2019. "The Recent Growth Boom in Developing Economies: A Structural-Change Perspective." In *The Palgrave Handbook of Development Economics*, 281–334, edited by Machiko Nissanke and José Antonio Ocampo. Cham: Palgrave Macmillan.

Fafchamps, Marcel, David McKenzie, Simon Quinn, and Christopher Woodruff. 2014. "Microenterprise Growth and the Flypaper Effect: Evidence from a Randomized Experiment in Ghana." *Journal of Development Economics* 106: 211–26.

Fafchamps, Marcel, and Simon Quinn. 2017. "Aspire." *Journal of Development Studies* 53 (10): 1615–33.

General Statistics Office. 2006. *Vietnamese Industry in 20 Years of Renovation and Development.* Ha Noi: General Statistics Office.

Grant, Matthew, and Meredith Startz. 2021. "Cutting Out the Middleman: The Structure of Chains of Intermediation." https://drive.google.com/file/d/1euKQt5Vwlw56J7cXnZvl622ASkGOJnM3/view.

Hallward-Driemeier, Mary, and Gaurav Nayyar. 2017. *Trouble in the Making? The Future of Manufacturing-Led Development.* Washington, DC: World Bank Publications.

Harrison, A., and Rodríguez-Clare, A. 2010. "Trade, Foreign Investment, and Industrial Policy for Developing Countries. In *Handbook of Development Economics*, 4039–4214. Amsterdam: Elsevier.

International Federation of Pharmaceutical Manufacturers and Associations (IFPMA). 2017. *The Pharmaceutical Industry and Global Health: Facts and Figures 2017.* https://www.ifpma.org/wp-content/uploads/2017/02/IFPMA-Facts-And-Figures-2017.pdf.

International Labor Organization. 2020. *The Supply Chain Ripple Effect: How COVID-19 Is Affecting Garment Workers and Factories in Asia and the Pacific.* International Labor Organization.

International Renewable Energy Agency. 2020. *Scaling up Renewable Energy Deployment in Africa.* International Renewable Energy Agency.

Jack, B. Kelsey, and Grant Smith. 2015. "Pay as You Go: Prepaid Metering and Electricity Expenditures in South Africa." *American Economic Review* 105 (5): 237–41.

Kabir, Ehsanul, Pawan Kumar, Sandeep Kumar, Adedeji A. Adelodun, and Ki-Hyun Kim. 2018. "Solar Energy: Potential and Future Prospects." *Renewable and Sustainable Energy Reviews* 82 (1): 894–900.

Kruse, Hagen, Emmanuel Mensah, Kunal Sen, and Gaaitzen de Vries. 2021. "A Manufacturing Renaissance? Industrialization Trends in the Developing World." UNU-WIDER Working Paper 2021/28.

Ksoll, Christopher, Rocco Macchiavello, and Ameet Morjaria. 2016. "Guns and Roses: Flower Exports and Electoral Violence in Kenya." *Global Poverty Research Lab Working Paper* 17–102.

La Porta, Rafael, and Andrei Shleifer. 2016. "The Unofficial Economy in Africa." In *African Successes, Volume 1: Government and Institutions*, 261–306. Chicago: University of Chicago Press.

Lewis, W. Arthur. 1979. "The Dual Economy Revisited." *Manchester School 47* (3): 211–29.

Li, Yue, and Martín Rama. 2015. "Firm Dynamics, Productivity Growth, and Job Creation in Developing Countries: The Role of Micro- and Small Enterprises." *World Bank Research Observer* 30 (1): 3–38.

Martin, Leslie A., Shanthi Nataraj, and Anne E. Harrison. 2017. "In with the Big, Out with the Small: Removing Small-Scale Reservations in India." *American Economic Review 107* (2): 354–86.

McCaig, Brian, and Margaret S. McMillan. 2020. "Trade Liberalisation and Labour Market Adjustment in Botswana." *Journal of African Economies* 29 (3): 236–70.

McCaig, Brian, and Nina Pavcnik. 2015. "Informal Employment in a Growing and Globalizing Low-Income Country." *American Economic Review* 105 (5): 545–50.

McCaig, Brian, and Nina Pavcnik. 2017. "Moving out of Agriculture: Structural Change in Vietnam." In *Structural Change, Fundamentals, and Growth: A Framework and Case Studies*, Chapter 2, edited by Margaret McMillan, Dani Rodrik, and Claudia Sepúlveda. Washington, DC: International Food Policy Research Institute.

McGuirk, Eoin F., and Nathan Nunn. 2020. "Nomadic Pastoralism, Climate Change, and Conflict in Africa." NBER Working Paper 28243.

McKenzie, David. 2017. "Identifying and Spurring High-Growth Entrepreneurship: Experimental Evidence from a Business Plan Competition." *American Economic Review* 107 (8): 2278–2307.

McKenzie, David, and Dario Sansone. 2019. "Predicting Entrepreneurial Success Is Hard: Evidence from a Business Plan Competition in Nigeria." *Journal of Development Economics* 141: 102369.

McMillan, Margaret, Dani Rodrik, and Claudia Sepulveda. 2017. "Structural Change, Fundamentals and Growth: A Framework and Case Studies." NBER Working Paper 23378.

Meyer, Christian Johannes, Eduard Krkoska, and Koen Maaskant. 2021. *Wages and Compensation in Ethiopia's Industrial Parks: Evidence from a Firm Survey*. Washington, DC: The World Bank Group.

Munshi, Neil, and Andres Schipani. 2021. "Failure of Democracy: Why Are Coups on the Rise in Africa." *Financial Times*, November 14. https://www.ft.com/content/a669c8e3-a744-445c-b613-9ff83059c90c/.

National Bureau of Statistics. 2014. *Nigerian Manufacturing Sector: Summary Report 2010-2012*. Nigeria: National Bureau of Statistics.

Newman, Carol, John Page, John Rand, Abebe Shimeles, Måns Söderbom, and Finn Tarp. 2016. *Made in Africa: Learning to Compete in Industry*. Washington, DC: Brookings Institution Press.

Nguimkeu, Pierre, and Albert Zeufack. 2019. "Manufacturing in Structural Change in Africa." World Bank Policy Research Working Paper 8992.

Niehoff, Justin D. 1987. "The Villager as Industrialist: Ideologies of Household Manufacturing in Rural Taiwan." *Modern China* 13 (3): 278–309.

Pahl, Stefan, Marcel P. Timmer, Reitze Gouma, and Pieter Woltjer. 2019. "Jobs in Global Value Chains: New Evidence for Four African Countries in International Perspective." World Bank Policy Research Working Paper 8953.

Pahl, Stefan, and Marcel P. Timmer. 2020. "Do Global Value Chains Enhance Economic Upgrading? A Long View." *Journal of Development Studies 56* (9): 1683–1705.

Robinson, James, 2015. "Natural Resources and Industrialization: Implications for Africa." Unpublished.

Rodrik, Dani. 2018. "New Technologies, Global Value Chains, and Developing Economies." NBER Working Paper 25164.

Schwerhoff, Gregor, and Mouhamadou Sy. 2020. "Where the Sun Shines." *Finance & Development* 57 (001).

Sen, Kunal 2019. "What Explains the Job Creating Potential of Industrialisation in the Developing World?" *Journal of Development Studies* 55 (7): 1565–83.

Sexton, Richard J., Josué Medellín-Azuara, and Tina L. Saitone. 2015. "The Economic Impact of Food and Beverage Processing in California." *Agricultural and Resource Economics Update* 18 (4).

Signé, Landry. 2018. *The Potential of Manufacturing and Industrialization in Africa: Trends, Opportunities,*

and Strategies. Washington, DC: Africa Growth Initiative at Brookings.

Signé, Landry, and Payce Madden. 2020. "Considerations for Rules of Origin under the African Continental Free Trade Area." *Journal of African Trade* https://dx.doi.org/10.2991/jat.k.201205.001.

Sutton, John. 2012. *Competing in Capabilities: The Globalization Process*. Oxford: Oxford University Press.

Teal, F., 2016. "Firm Size, Employment and Value-Added in African Manufacturing Firms: Why Ghana Needs Its 1 Per Cent." Centre for the Study of African Economies, University of Oxford Working Paper 2016-07.

Timmer, M., de Vries, G.J., and De Vries, K. 2015. "Patterns of Structural Change in Developing Countries." In *Routledge Handbook of Industry and Development*, 79–97, edited by John Weiss and Michael Tribe. Routledge.

UNIDO. 2020. INDSTAT 2 Industrial Statistics Database at 2-digit level of ISIC Revision 3. Vienna. Available from http://stat.unido.org. (accessed October 1, 2019).

Winkler, Harald, and Anthony Black. 2021. "Creating Employment and Reducing Emissions: Options for South Africa." SARCHl Industrial Development Working Paper 2021-06.

World Health Organization. 2021. "What Is Africa's Vaccine Production Capacity?" March 19, *WHO Africa*. https://www.afro.who.int/news/what-africas-vaccine-production-capacity.

Wright, Gavin. 1990. "The Origins of American Industrial Success, 1879–1940." *American Economic Review* 80 (4): 651–68.

Zarocostas, John. 2021. "What Next for a COVID-19 Intellectual Property Waiver?" *Lancet* 397 (10288): 1871–72.

Journal of Economic Perspectives—Volume 36, Number 1—Winter 2022—Pages 33–56

Agricultural Technology in Africa

Tavneet Suri and Christopher Udry

Economic development typically involves a structural transformation in which a dominantly agrarian economy moves towards being more manufacturing- and services-based. However, economies of countries across Africa have an especially high reliance on agriculture, both for output and jobs, compared with the rest of the world. The high shares of agriculture in GDP and employment in Africa largely reflect the low level of GDP per capita on the continent (Herrendorf et al. 2014). Figure 1 shows that agriculture is almost 20 percent of GDP in Africa, compared with a world average of about 5 percent. Moreover, the share of agriculture in GDP of the African region has remained stable over the last 50 years, whereas the share for other regions that started high in 1970—South East Asia and South Asia—has fallen a lot. Panel B shows how agricultural shares of employment have declined across regions of the world in the last 30 years. Africa now has the highest share of employment in agriculture at about 50 percent, given the declines in the South Asia region, while the world average of employment in agriculture is closer to 30 percent.

There is also wide variation in these shares across regions within Africa, the highest being East and West Africa and the lowest South Africa. Panel C of Figure 1 shows that agriculture's share of GDP has been falling in East Africa, where it has historically been highest, but not in other regions. Panel D shows the wide disparity

■ *Tavneet Suri is the Louis E. Seley Professor of Applied Economics, MIT Sloan School of Management, Cambridge, Massachusetts. Christopher Udry is Robert E. and Emily King Professor of Economics, Northwestern University, Evanston, Illinois. Their email addresses are tavneet@mit.edu and christopher.udry@northwestern.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.36.1.33.

Figure 1
Agriculture's Contribution to GDP and Jobs

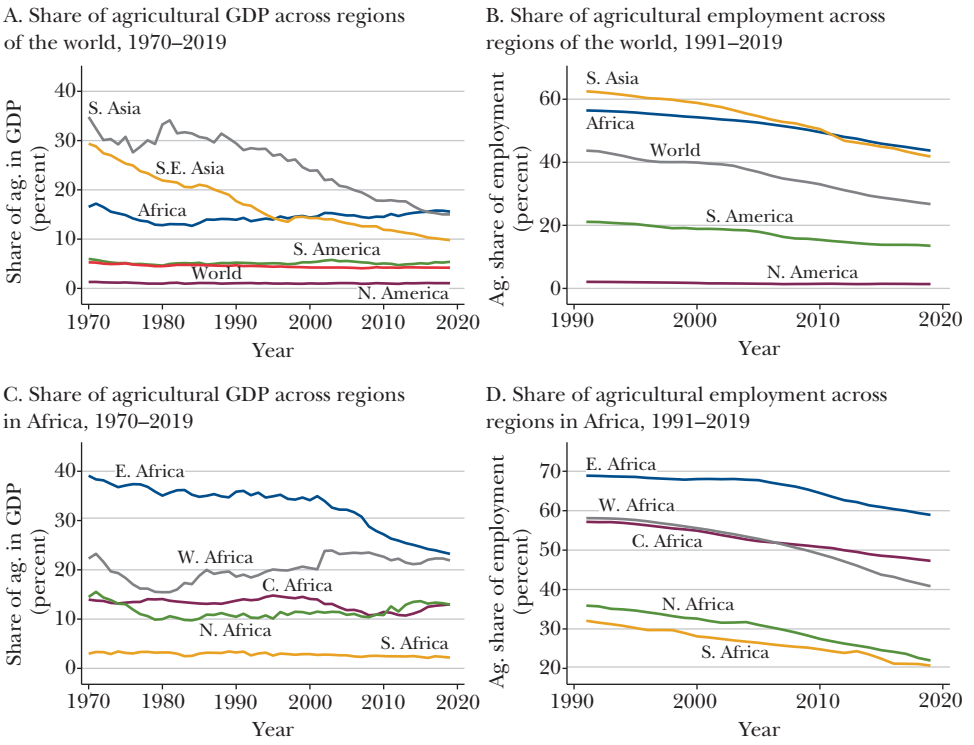

Source: Panels A and C are based on data from the United Nations, Food and Agriculture Organization (FAOSTAT) at https://www.fao.org/faostat/en/#home. Panels B and D are based on data from the World Development Indicators from the World Bank at https://databank.worldbank.org/source/world-development-indicators. The countries that are included in Africa and regions of Africa are described for each of these two main data sources in the online Appendix, available with this article at the *JEP* website.

in the share of the workforce in agriculture across regions of Africa, although the share is dropping everywhere.

A first step towards structural transformation happens as the agricultural sector evolves from smallholder farmers growing mainly food crops (cereals) for self-consumption to larger scale farmers growing food crops primarily for sale. At present, about 80 percent of African farmers are smallholders with under two hectares of land, who together account for 40 percent of cultivated area (Lowder et al. 2016), although farm sizes do seem to have been on the rise recently in some African countries (Jayne et al. 2016). Increasing agricultural productivity through improved technology is key to this process of agricultural and structural transformation (Bustos et al. 2016, 2020; Dercon and Gollin 2014; Gollin et al. 2021). Examples of specific technological changes that improve labor productivity in agriculture that might be part of this structural transformation would include mechanization of

farm activities including land preparation and transportation, and the use of labor-saving agrochemicals like pesticide. There are clearly documented causal links from increased agricultural productivity to reduced poverty (for a good review, see de Janvry and Sadoulet 2010) and improved child nutrition (for example, Glennerster and Suri 2018).

There are many historical examples of the flexibility and openness to innovation of farmers across Africa: centuries-old examples like the introduction of maize, cassava, and sweet potatoes to Africa as part of the Columbian exchange; decades-old examples like the transformation of the economy of Ghana with the introduction of cocoa (Hill 1963); and more recent examples like the emergence of commercial flowers for export from Kenya and Ethiopia. But overall, these changes have not been sufficient to generate sustained productivity growth across the board in agriculture in Africa, and productivity growth in African agriculture has been slow relative to that in the rest of the world. Across the world, value added per worker in agriculture is lower than it is in the rest of the economy, but the gap is larger in Africa than it is elsewhere (Gollin et al. 2014).

In this paper, we start with some background about agricultural productivity and technology adoption in Africa, highlighting how it has lagged. We then discuss what may explain these lags and what we know about each of these explanations. We discuss the importance of heterogeneity, along a variety of dimensions, in understanding the patterns of agricultural productivity and technology use in Africa. Public and private investment in new agricultural technology in Africa has been extremely limited, so there is no large set of profitable technologies waiting on the shelf to be adopted by Africa's farmers. However, there are viable directions for policy and an important set of unanswered research questions.

The What: Trends in Agricultural Technology Use and Productivity in Africa

We first describe some basic patterns of agricultural productivity and innovation in Africa.[1] Figure 2 shows agricultural yields for cereals (as measured in value/hectare) in Africa versus the world over the last 60 years. Both the level and growth rate of yields in Africa lag other regions. Of course, looking at Africa as a whole masks considerable heterogeneity. For example, in more detailed breakdowns from the same data, cereal yields in the South Africa region have climbed substantially to $900/hectare in the last decade or so, while yields in countries of West Africa and

[1] Much of the aggregate data comes from the Food and Agriculture Organization of the United Nations, which in turn largely relies on ministerial or national statistical office sources. The World Bank's Living Standards Measurement Study—Integrated Surveys on Agriculture program (Christiaensen and Demery 2017) has provided essential new data on agriculture in Africa, complementing both these official sources and the many smaller-scale researcher-led surveys.

Figure 2
Cereal Yields by Region of the World, 1961–2018
(in thousands of US dollars per hectare)

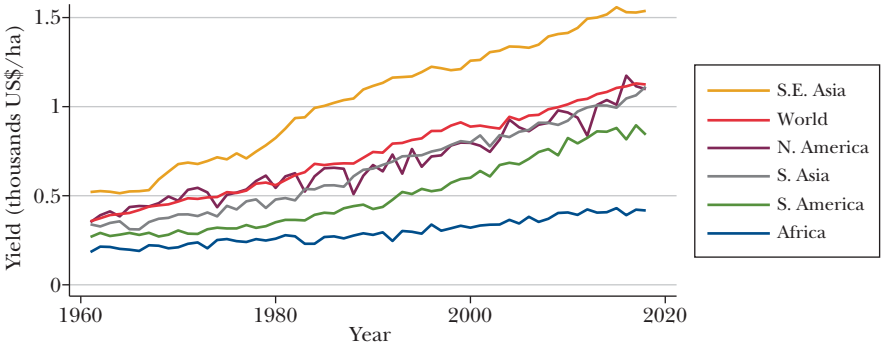

Source: United Nations, Food and Agriculture Organization (FAOSTAT) at https://www.fao.org/faostat/en/#home.
Note: Includes barley, buckwheat, canary seed, cereals not identified separately (canagua/coaihua, quihuicha/Inca wheat, adlay/Job's tears, wild rice, other minor unclassified locally relevant cereals), fonio, maize, millet, oats, rice (paddy), rye, sorghum, triticale, and wheat.

North Africa rose more modestly to about $500/hectare, and yields in Central and East Africa have risen only slightly to about $300/hectare.

Underlying these yield trends are, of course, changes in farm size, though there is much less variation across low- and middle-income countries on the average farm size. Figure 3 below shows some of this variation, highlighting how the average farm size now is likely not that different in Africa compared to South Asia (though both of these regions have dramatically smaller farm sizes than in Latin America and the high-income countries). It is worth noting, however, that data on farm size is less consistently collected and so may be noisier (Lowder et al. 2016).

These yield gaps across the world and across Africa are largely a consequence of the dramatically different technologies being used. Technologies in agriculture involve the biological processes of plant and animal growth, coupled with the physical actions that create the uniform conditions that distinguish a farm from natural growth. Agricultural technology is embedded in seeds, breeding stock, irrigation and other water management methods, chemical inputs such as fertilizer and pesticides, agronomic practices such as fallowing patterns and plant spacing, and equipment like hand tools, tractors, or pumps. Large variation across regions is apparent along many of these dimensions. In Figure 4, we show fertilizer use comparing Africa to regions of the world as a whole, and comparing across regions within Africa. Most other regions in the world mostly caught up to North American levels of fertilization at least two decades ago. Although there have been substantial increases in fertilizer use in South and North Africa, gains in other areas of Africa appear quite modest.

Figure 3
Average Farm Size by Region of the World, 1960–2000

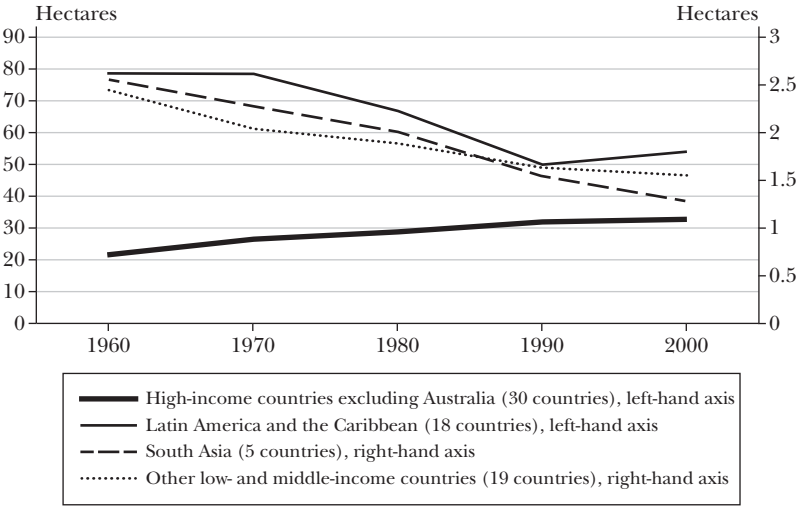

Source: Lowder et al. (2016)

Figure 4
Fertilizer Use, 1961–2018
(kilograms/hectare)

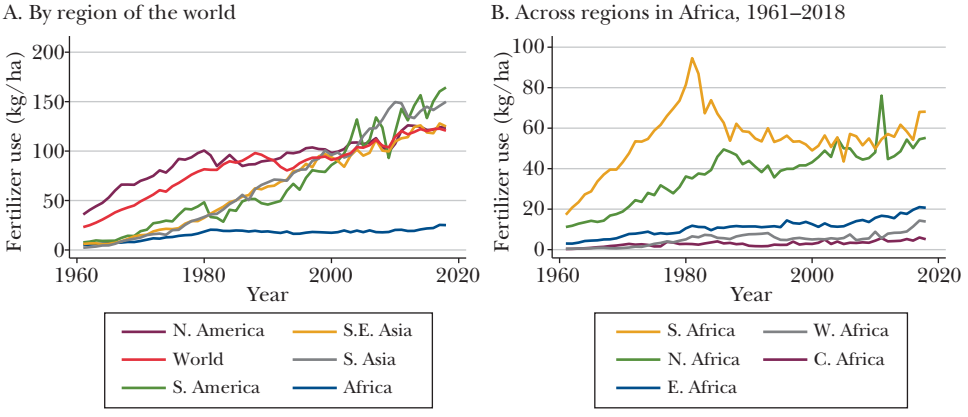

Source: United Nations, Food and Agriculture Organization (FAOSTAT) at https://www.fao.org/faostat/en/#home.

Figure 5
Trends in Irrigation, 1960–2020

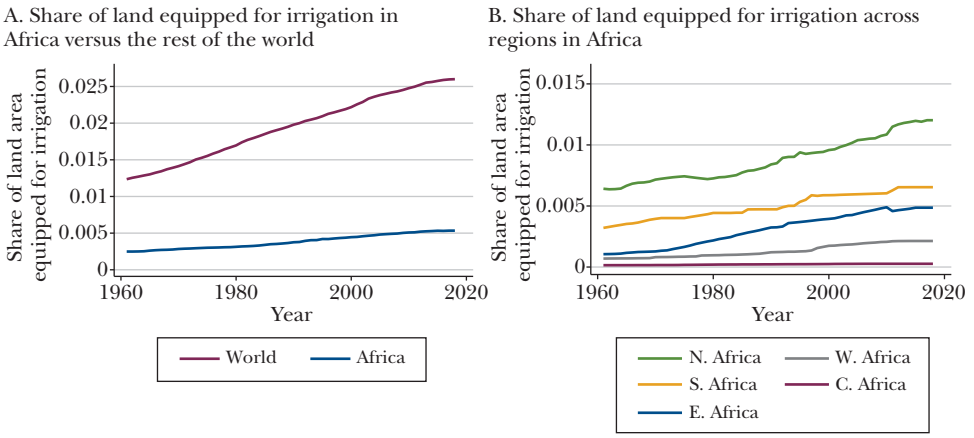

Source: United Nations, Food and Agriculture Organization (FAOSTAT) at https://www.fao.org/faostat/en/#home.

In the case of irrigation as well, Africa lags quite far behind, with little sign of catch up. As shown in Panel A of Figure 5, the share of agricultural land area equipped for irrigation remains at approximately 0.005 for the continent, compared to over 0.025 for the world. Again, there is substantial heterogeneity across regions of Africa, with irrigation concentrated in North Africa and in South Africa (Panel B, Figure 5). Almost all farmers in West, Central, and Eastern Africa rely exclusively on rainfall for their crops.

Fertilizer and irrigation are two key agricultural technologies. As we will discuss later in the paper, these inputs also have the property of reducing heterogeneity in characteristics of land: for example, irrigation projects, to some extent, level out natural differences in supply of water to farmland. In turn, the extent of heterogeneity has implications for how widely other technologies like seed varieties and cultivation practices can be broadly shared.

The Why: Reasons for Africa's Technological Stagnation in Agriculture

Our aim here is to provide a framework for understanding the range of constraints that account for the stagnation of technological progress in agriculture in Africa. We highlight some of the more recent studies in African contexts that help us understand why technologies that appear to improve productivity are not developed or adopted. A key source of evidence on the importance of these constraints is provided by the 26 randomized controlled trials across 10

African countries supported by the Agricultural Technology Adoption Initiative (Bridle et al. 2019). Broader reviews can also be found in de Janvry et al. (2017), Binswanger-Mkhize and McCalla (2010), and Suri and Udry (2021)—the last review specifically for the case of Africa. Ultimately, none of these constraints taken alone seems able to explain the low technology adoption (and hence low yields) in Africa.

Summary: No Single Binding Constraint

No single constraint explains the low productivity in African agriculture; instead, different combinations of constraints seem to bind for different farmers. As a result, packages of interventions may end up being the most useful approach to adoption of new technologies and improving agricultural productivity. Currently, we do not know much about effective packages of interventions (we return to this in more detail below). Examples of packages include a program in Kenya, organized by a nonprofit called Drumnet, that combined training, agricultural credit, improved access to saving, input supply, and marketing assistance and showed that this led to increased adoption of a novel export crop (and the related inputs) (Ashraf et al. 2009). Also in Kenya, the One Acre Fund's program of group lending, crop insurance, regular training, input supply, and market facilitation support generated adoption of improved practices, increased use of fertilizer, and increased yield and farm profits (Deutschmann et al. 2019). A multi-faceted "economic inclusion" program in Niger with training, large cash grants, and a set of psycho-social interventions generated substantial increases in improved livestock and use of fertilizer and phytosanitary inputs (Bossuroy et al. 2021).

However, such multifaceted programs have not always been successful. The Drumnet export-promotion program in Kenya, for example, collapsed after a year when farmers were unable to meet EU production requirements. Some earlier programs of this type were costly failures, like the Integrated Rural Development Programs implemented in the 1970s and 1980s (Chambers 2013), or the minimal effects of some community-driven rural development programs (Appiah et al. 2020). We lack a clear understanding of what elements of these programs are essential and in which environments.

Farmers do show an ability to overcome these constraints when the technology is sufficiently productive (that is, if it increases yields by enough) and hence actually profitable (Suri 2011). There are many examples of technology being adopted in specific contexts, such as cocoa in Ghana and Cote d'Ivoire, flowers in Kenya and Ethiopia, and improved cassava in Nigeria. But this notion of "sufficiently productive and profitable" is limited in geographic, economic, and social extent by the heterogeneity that we discuss in the next section.

Below, we summarize some of the main constraints that have been studied in the literature and some overall findings on these constraints. We do not detail findings in specific papers—instead, we draw on the extensive review in Suri and Udry (2021) and refer readers there for a more detailed, paper-by-paper literature review.

Credit, Liquidity, and Savings Constraints

If farmers face binding credit constraints, or barriers to saving, they may not invest in new agricultural technologies. Much of the work on this topic has focused on adoption of fertilizer and/or improved seeds (which are often complementary inputs). Many countries in sub-Saharan Africa have hosted experiments with input subsidy schemes; good reviews of experiences with fertilizer subsidies can be found in Druilhe and Barreiro-Hurlé (2012) and Smale and Thériault (2018). These papers do find increases in adoption when inputs are subsidized, though the increases are small relative to the gap documented above.

Another approach is to use credit or savings to improve the quality of output, using a range of policy tools such as the provision of improved storage technologies and encouraging the use of technologies to combat carcinogenic aflatoxins that can grow on a fungus in maize. Finally, one could provide credit or savings products directly to farmers, including commitment savings technologies, which have been shown, in some cases, to improve technology use or labor (again, for detailed descriptions of the related underlying literature, see Suri and Udry 2021).

Overall, constraints on borrowing and saving seem to affect production decisions by many farmers in Africa, but only modestly. Alleviating potential credit constraints increases the use of agricultural technologies like improved seed and fertilizer, but ultimately by modest amounts and only in some contexts.

Insurance Constraints

Farmers who are not able to insure risk may bias their decisions towards low-risk, low-return technologies. This bias may be exacerbated if they are poor, in which case downside risk may be health- and life-threatening.

Financial innovation may mitigate the consequences of risk. For example, rainfall-based crop insurance might address weather risk. However, the studies in J-PAL (2016) across a number of countries find that take-up of rainfall insurance at actuarially fair rates is very low (for example, Carter et al. 2017), though high when subsidies are substantial. When the price risk of crops is insured directly, there are also impacts on the adoption of agricultural technologies.

There is evidence that when risks are insured, adoption of technologies like fertilizer and new seed varieties can improve, although important challenges remain in designing these insurance mechanisms. Improved information to serve as the basis for such insurance contracts may become available as the cost of extremely high-resolution satellite images starts to fall and remote sensing-based crop yield measurement improves (Benami et al. 2021; Lobell et al. 2020). Such data could help: it would get better (and independently verified) estimates of the impacts of any rainfall event (using data from rainfall stations still leaves a lot of uninsured risk for farmers as the number of rainfall stations is low and therefore does not offer complete coverage) and hence provide a much cheaper alternative to verifying the impacts of any rainfall event via in-person visits.

Finally, if farmers truly worry about downside risk given they are close to the consumption floor, guaranteeing them an income floor may also change their

investments in technology and their productivity in fundamental ways. For example, Banerjee et al. (2021) find that a universal basic income in Kenya raises agricultural investments, both in terms of agricultural assets (including livestock) as well as the use of agricultural technologies.

Information Constraints

Farmers in Africa may not know about new technologies or how to use them effectively. Providing information to farmers directly, as well as through social networks, may address this problem. There are a wide variety of studies of various extension programs, public and private, but most of these studies do not find transformational effects on technology adoption (for overviews, see Caldwell et al. 2019; Bridle et al. 2019; Suri and Udry 2021). These modest impacts are found irrespective of the mode of delivering the information (for example, using direct training programs as in traditional extension or via text messaging or video messaging). While a number of studies show that farmers learn about new technologies through social networks, this new knowledge is unable to close much of the adoption gap.

Overall, extension and information programs seem to have limited effects on technology adoption, except perhaps in settings where the technologies or the crops are truly new.

High Transaction Costs and Infrastructure

Farmers across Africa face high transaction costs across multiple dimensions: both input and output markets are hard to access, with poor supply chain investments and infrastructure for many agricultural products. Search costs are high and there may often be many layers of intermediaries between sellers of the agricultural good and the final buyer (Startz 2020; Grant and Startz 2021). The traders who play an important intermediary role in farmers' access to output markets may have market power (Bergquist and Dinerstein 2020; Casaburi et al. 2013; Newman et al. 2018). Markets are often not well-integrated, though this is slowly changing in Africa with growth in cell phone adoption and use. Sometimes, farmers cannot easily sell their output or easily access inputs. The profits from new technology use may be low given this current state of infrastructure.

The literature on market access is growing, including studies that highlight the importance of inadequate supply chains for input and output markets as barriers to adopting new technology. African agriculture often seems to lack markets for quality—that is, the lack of different prices for varying quality of crops. Providing quality incentives seems to be important in driving farmers' investments in quality, but we still have a lot to learn here. There is ample evidence that farmers lose significant quantities of harvested crops to post-harvest and storage losses. At the farmer level, these can amount to between 4 percent and 8 percent of output (Sheahan and Barrett 2017). There is a small (but growing) literature on the role of technologies that avoid such losses, showing that these may improve technology use and productivity. However, again, there is room to reinforce some of these early studies.

Digital financial tools may also hold some promise to improve the agricultural value chain. Studies on these aspects have only just started, so there is still a lot to learn. An example is using technology to build better trading markets, such as the Ethiopian Commodity Exchange, a digital crop market in Uganda called Kudu (Newman et al. 2018), and M-Shamba in Kenya. A similar example in the area of crop insurance is the Kilimo Salama rainfall insurance program in Kenya, which relies on mobile money for payments (Greatrex et al. 2015).

Imperfect Labor Markets

Given the low population density in many African economies, farmers in rural labor markets may find it hard to hire the right types of labor at the right times. There is not much research about the constraints that poorly functioning labor markets play in agricultural productivity in Africa, in part because it is hard to design interventions in labor markets that seem likely to improve labor allocation and also to be implementable from a policy perspective. However, some studies have shown that large-scale investments (such as irrigation and mechanization) are not fully exploited due to poorly functioning labor markets and the inability to hire enough labor when needed.

Imperfect Land Markets

Imperfect land markets and poorly defined property rights for farmers may explain some of the low level of investments in agricultural technologies, as farmers may not feel confident in their access to the future returns on these investments. A small literature addresses how improvements in property rights may affect the use of agricultural technology. Across the literature in this area, it does seem that improvements in property rights generally increase investment, but there are exceptions.

Large-scale commercial agriculture is rare in sub-Saharan Africa, at least in part due to the barriers in gaining access to land (Collier and Dercon 2014). Land market imperfections inhibit the growth of particularly productive farms (Restuccia and Santaeulalia-Llopis 2017), preclude farmers from capturing potential returns to scale associated with mechanization (Foster and Rosenzweig 2017), and slow the development of robust value chains that can support technological innovation (Barrett et al. forthcoming).[2] However, the same customary land tenure systems that inhibit the commercialization of land provide guarantees of some access to land for most people in rural Africa and play a central role in informal social protection. Transformation of land tenure systems to facilitate agricultural innovation requires attention to these broader issues.

[2] Gollin and Udry (2021) argue that Restuccia and Santaeulalia-Llopis (2017) overestimate the potential gains to reallocating land to more productive farmers, but do not dispute the existence of misallocation due to land market imperfections.

The Where and When: The Broader Role of Heterogeneity

Many agricultural technologies exhibit an extreme sensitivity to local circumstances: nutrients, moisture, temperatures, and solar energy are all required in appropriate proportions and timing for crops and animal husbandry, and farmers face different threats from diseases and pests (Evenson and Westphal 1995; Moscona and Sastry 2021). Overlaid on these natural conditions is heterogeneity in infrastructure and market access and hence in the prices of inputs and outputs, again both over time and space. Such heterogeneity is almost the very nature of agriculture. Even in agricultural regions outside Africa with well-developed agricultural technology, we see fine geographical variation in optimal technology choices: in the American state of Illinois, for example, optimal maize varieties differ at the scale of a few dozen miles (for example, see https://burrusseed.com/product-selection-guide/).

We believe that this heterogeneity is a key to understanding technological stagnation in African agriculture more broadly. For Kenya, Suri (2011) shows large heterogeneity in both gross and net returns to hybrid seed and fertilizer. Indeed, the extent of heterogeneity appears to be far larger in sub-Saharan Africa than in temperate regions. Claassen and Just (2011) study the variance of log yield across farms in the United States: they find that the 95th percentile of corn yield is 190 percent larger than the 5th percentile yield. For comparison, the 95–5 ratio for Uganda is 9,304 percent and for Tanzania 2,558 percent (Gollin and Udry 2021). A substantial share of this variation is due to heterogeneity in soil, moisture, temperature, and other dimensions of growing conditions over space. This exceptional degree of heterogeneity is at least in part a consequence of the eighteenth- to twentieth-century history of low population density on most of the continent (itself caused in part by the slave trades and associated violence and political disruptions, as discussed in Manning 2014). The low population density makes infrastructure costs high, so that farmers rely on rainfall for irrigation and on periodic fallowing (rather than fertilizer) to maintain soil fertility (Boserup 1965). Agricultural production therefore is more tightly connected to local conditions: rainfall realizations, plot-level availability of soil nutrients, the existence (or not) of a nearby road and/or market. The implications of such heterogeneity can be far-reaching, not just in explaining the lack of adoption, but also in underscoring the challenges of developing new agricultural technologies and designing policy that ultimately improves yields.

Heterogeneity in Soil and Land Quality

The nutrient composition, physical properties, and biochemistry of the soil affects the profitability of any given seed, fertilizer or other chemical input, or other agronomic practice. Sanchez (2019) provides an extensive review of soil properties and soil management across the world, highlighting the extreme heterogeneity in soil quality and the importance of the interaction of nutrients and water availability. For example, the maize yield response to fertilizer is strongly dependent on sandiness and soil carbon content (Burke et al. 2020). These land characteristics vary in

important ways across communities, at scales as small as hundreds of meters: Hengl et al. (2021) show some maps of data collected by the International Soil Reference and Information Centre (ISRIC), displaying differences in the soil acidity at a resolution of 30 meters. There is even substantial variation in soil characteristics *within* farms, and farmers adjust production decisions to accommodate that variation (for example, Tjernström et al. 2015). Soil quality is not fixed, but evolves over time in response to farmer actions, weather shocks, and environmental influences (like seed dispersal and the movement of pests).

Given large variations in soil quality across even small geographic areas, optimal technology choices will vary between plots or even on a specific plot over time. For example, Harou et al. (2020) show that in Morogoro Rural, Tanzania, within-village variation accounts for almost one-third of overall variation across the district in several key soil nutrients, and thus recommended amounts and types of fertilizer would vary similarly within villages. Farmers who were provided with plot-specific recommendations for appropriate fertilizer use (along with vouchers for reduced cost access to inputs) were more likely to apply the recommended fertilizer, and increased yields by over 150 percent relative to the control group.

Heterogeneity in Weather

The return to investment in cultivation in sub-Saharan Africa (and South Asia) can vary widely, depending upon the local weather (Rosenzweig and Udry 2020). As a recent example, McCullough et al. (2020) use an experimental crop trial meta-database to show that the profitability of adopting fertilizer varies strongly with weather outcomes. This temporal heterogeneity makes the adoption decisions of farmers much more difficult; it also demands a longer-term perspective for those who seek to understand these adoption decisions.

There has been some investment in crop varieties that are more resistant to poor weather realizations, and in particular to varieties that may but outperform the alternatives during poor weather realizations. This is an exciting area of varietal improvement, though such varieties are still few and far between. Examples of such varieties are Swarna Sub-1 flood resistant rice in India (Emerick et al. 2016) and drought resistant maize (Boucher et al. 2019).

Heterogeneity in Access to Markets

Scarce and low-quality road infrastructure, together with imperfect competition in transportation, generate large variation across space in both input and output prices. Atkin and Donaldson (2015) demonstrate that in Nigeria and Ethiopia, the effect of distance on prices of traded goods is four or five times the effect in the United States. Direct trucking costs in Africa are similarly much larger than those in developed countries (Teravaninthorn and Raballand 2009). Median trade costs in Africa are about five times higher than everywhere else in the world (Porteous 2019).

These high transportation costs interact with a spatially dispersed population of farmers to produce markets for agricultural output and inputs that are far from

a perfectly competitive benchmark (Bergquist and Dinerstein 2020). The traders to whom farmers sell (often at farm gate) are not competitive, further affecting output prices (for example, Casaburi et al. 2013). The consequent price mark-ups also vary across space. Some countries have a policy of uniform national prices of some inputs; such policies are problematic given high transportation costs, because it leads to spatial variation in access to these inputs and in the returns to and adoption of these inputs (as Suri 2011 shows for Kenya).

These high, heterogenous, and variable trade costs interact with the adoption of new technologies in specific ways: for example, subsidies for fertilizer only raise farmer incomes when trade costs are low (Porteous 2020).

Pervasive Heterogeneity and Technology Adoption

This pervasive heterogeneity presents challenges, both for African farmers attempting to make optimal technology decisions and for researchers seeking to understand farmer technology choices. A technology that is profitable on one set of farms in a particular year may fail on other nearby farms or in other seasons.

Agricultural technologies vary in the degree and nature of their sensitivity to local circumstances. The profitability of adopting mechanical field preparation in savannah regions (that is, a grassland with numerous but widely spaced trees), for example, depends on factor prices and topography, but not as strongly on local variations in soil characteristics. The profitability of adopting improved storage technologies depends on prices and local insect populations, but not as much on topography. Innovations tied more closely to the biological and physical characteristics of a plot, like chemical inputs and seed varieties, exhibit the most heterogeneity in returns (for example, Harou et al. 2017).

Farmers presented with an opportunity to adopt a new agricultural technology must translate information they receive regarding the performance of the technology in specific circumstances to its likely performance on their own farms. Understanding how farmers make this translation remains incomplete. Information from extension agents and expert advice can be effective, but some farmers place more weight in learning from others who are more like themselves (for example, Munshi 2004; BenYishay and Mobarak 2019).

Heterogeneity has implications for the optimal supply of agricultural innovations as well, because the circumstantial sensitivity of many innovations limits the extent of their potential impact. This reduces the incentive for private research. In addition, any cost-benefit analysis for public support of research into new agricultural technologies needs to take the more restricted impact into account.

The How: Where to from Here?

Given what we know to date about the stagnation of both agricultural technology use and agricultural productivity in Africa, what policy actions might be appropriate?

Investments in Technology for Agriculture

There is a need for more and longer-term investments in research and development for agricultural technologies themselves in Africa. Spending per farmer on R&D for agriculture in Africa is two orders of magnitude lower than in developed countries. The United Nations Food and Agriculture Organization, based on data from ASTI (IFPRI), estimates that spending per farmer on R&D across Africa has been about $50 (in US dollars at purchasing power parity) or less since 1990.[3] For comparison, R&D spending per farmer in Brazil was about $1200 in 2000 and has more than doubled since then.

If one looks across regions of Africa, one can see substantial increases in South Africa and East Africa, where research and development spending per farmer has recently climbed substantially to nearly $250 and almost $150, respectively. But these changes make the average R&D spending per farmer for Central, North, and West Africa, hovering at less than $25, look even lower. Countries like Uganda, Zimbabwe, Niger, Burundi, and Ethiopia had growth in public agricultural R&D spending at an annual rates of 6 percent or more from 2000 to2014, according to data from CGIAR ASTI. Conversely, countries like Togo, Gabon, and Guinea averaged annual declines of 4 percent or more in public agricultural R&D spending over that time frame.

In short, the lack of productive new agricultural technologies ready to be adopted in Africa is no mystery; it's the result of low levels of agricultural research and development investment in the past.[4] In addition, volatility in already low levels of investments make the problem worse. Rawat (2020) shows for public investment in R&D, the volatility in these expenditures is highest in sub-Saharan Africa (and lowest in South Asia).

Private research and development is not filling the gap. Although private R&D in agriculture has grown tremendously over the last 20-odd years and it is expected to play a stronger role in the future (Fuglie 2016), it is still only one-quarter of overall agricultural R&D. Moreover, almost all private agricultural R&D is in (and for) the developed world, as illustrated in Figure 6 below (Heisey and Fuglie 2018).

Agricultural research and development is a process that builds on itself. For example, in South Asia after the introduction of the initial "green revolution" varieties of rice, Evenson and Gollin (2003) show that there was a sustained research and development effort leading to multiple generations of new varieties gradually

[3]This estimate for Africa includes only a sub-sample of countries where there is consistent data during the last three decades: Benin, Botswana, Burkina Faso, Burundi, Republic of Congo, Côte d'Ivoire, Ethiopia, Gabon, The Gambia, Ghana, Guinea, Kenya, Madagascar, Malawi, Mali, Mauritius, Niger, Nigeria, Senegal, South Africa, Sudan, Togo, and Zambia.

[4]Perhaps surprisingly, the share of government expenditure that goes to agriculture in Africa is similar to other world regions, according to FAOSTAT data. The expenditure on agriculture includes agricultural research and extension services, input subsidies, irrigation, marketing, and rural road infrastructure. Across Africa, about 5 percent of all government expenditures go to agriculture, whereas for the world as a whole, as well as comparable regions like South East Asia, only about 3.5 percent of government spending goes to agriculture as defined here.

Figure 6
Public and Private Agricultural Research and Development Spending by Region, 2011 *(millions of US dollars, converted at purchasing power parity exchange rates)*

						EMEA		*APAC*	
	High-income countries	*Other countries*	*World*	*North America*	*Latin America*	*High-income countries*	*Other countries*	*High-income countries*	*Other countries*
Public agricultural R&D	18,212	24,116	42,328	5,501	6,901	7,879	2,999	4,832	14,215
Private agricultural R&D, allocated to country of company incorporation	12,326	1,933	14,260	6,458	36	5,023	110	845	1,788
Private agricultural R&D, allocated to location of company product sales	9,510	4,750	14,260	5,106	1,968	3,696		3,489	

Source: Heisey and Fuglie (2018).
Note: EMEA refers to Europe, Middle East, and Africa; APAC refers to the Asia Pacific region.

increased yields, improved agronomic performance and nutritional content, and increased resistance to pests. In India, it is common to have 20–40 new varieties of rice released each year since 1970, along with 10–20 new varieties of both maize and wheat each year, as reported by the Indian Council of Agricultural Research (ICAR)–Indian Institute of Maize Research (IIMR). In addition, the numbers of new varieties introduced has risen in recent years (as shown in SeedNet India Portal). In contrast, in Kenya where maize is the main staple, it was common to have five or fewer new varieties introduced annually from 1970 to 2000, although since then the number of new varieties of maize being introduced has risen to Indian levels, according to data from the Kenya Plant Health Inspectorate Service (KEPHIS). Of course, what matters is not just the number of varieties but the quality or yield improvement provided by these varieties. As Karanja (1996) highlights for the earlier research and development efforts in Kenya, research yields were exhibiting a "plateau effect," with newly released varieties in 1989 having smaller yield advantages over their predecessors than previously released ones.

Broad or Customizable Technology?

The substantial heterogeneity experienced by most African farmers can be addressed in two ways. One approach is that a combination of irrigation, permanent cultivation, and terrain engineering can reduce local differences. The other is to seek new technologies that are more customizable or that can be profitable across a wider range of circumstances.

In developed countries, there is much more customization of new agricultural technologies to very local circumstances than there is in Africa. More than 60 years ago, Griliches (1957) was showing how new technologies were adapted to local conditions in US agriculture. It has long been common for farmers in developed country agriculture who are separated by only a few dozen miles and facing almost

uniform input and output prices to prefer different seed varieties for major grains. Hurley et al. (2004) show that within single farm plots in the American state of Minnesota, optimal fertilizer doses vary can by a factor of two. Recently, this customization of technology to (relatively) small variations in growing conditions, even within a single plot, has intensified in high-income countries with the growth of precision agriculture (for example, Stoorvogel et al. 2015; Schimmelpfennig 2016; North Dakota State University 2021).

Agricultural research efforts in most of Africa have been less successful in adapting technology to local variation. In most of Africa, only extremely limited customization of modern varieties of seed is available. In agronomic trials in northern Ghana, for example, the best performing maize seed was a variety developed for South African conditions (van Asselt et al. 2018). Fertilizer recommendations to African farmers are often uniform over large areas of highly variable soils (Michelson et al. 2021).

Developing, testing, and adopting new technologies in this environment poses significant challenges. Customization requires much better feedback from a much larger and more diverse set of farmers. One promising direction is a greatly expanded use of farmer participatory trials to map heterogeneity and optimize recommendations. A related approach would use the rapid expansion of information network availability across Africa to carry out on-farm trials of new technologies at a much larger scale than has been possible in the past (Newman et al. 2012). However, we first need to build up the base of relevant local information on what technologies are viable for testing in what local areas of each country.

One obstacle to this approach is the low number of agricultural research stations and the historically low numbers of staff and high turnover at these stations (Lipton 1988) across countries in Africa. The United States has 607 research stations (Pearson and Atucha 2015); according to data on farmers from USDA (2020), this is about 134 research stations per 100,000 farmers. For comparison, we looked at data on the number of research stations in some African countries from Beye (2002), along with data on farmers from about 2002 from the Agricultural Science and Technology Indicators (ASTI) published by CGIAR. In Ghana there are only 14 agricultural research stations (0.28 per 100,000 farmers); in Malawi also only 154 (0.34 per 100,000 farmers); in Mali, seven (1.02 per 100,000 farmers); in Madagascar, 24 (0.42 per 100,000 farmers); in Kenya, 25 (0.22 per 100,000 farmers), in Cameroon, 30 (0.79 per 100,000 farmers); and in Senegal, also 30 (0.94 per 100,000 farmers). In Africa, the tests by research stations of new varieties simply cannot provide the necessary local information on returns to varieties that farmers need, which in turn can help to explain the large differences in returns between station trials and on-farm trials that are often observed (Laajaj et al. 2020).

We also do not know enough about what seed varieties are actually used by African farmers. Recent work on DNA fingerprinting shows that there are discrepancies between self-reports of what farmers say they are using and the DNA results (Poets et al. 2020), though much more research remains to be done here. Similarly, the labs doing quality testing in these countries may need more investment.

As one example, there are disputes over the quality of inputs in different countries; Bold et al. (2017) and Ashour et al. (2019) argue that there is significant adulteration of fertilizer and herbicides, respectively, in Uganda, but Sanabria et al. (2013) and Michelson et al. (2021) provide evidence that fertilizer is of good quality across multiple countries. Clearly, we need a significant investment push to generate fine-scale information on what technologies are used and the returns to those technologies.

In that vein, large-scale participatory trials of new technologies could provide a foundation for a broader process of integrating "citizen science" into agricultural research and development and agriculture extension activities. As discussed earlier, the same information technology that would permit the communication of trial protocols and results between scientists and participants is already being used to provide inexpensive, timely information and advice to farmers. It could also provide a forum for generating new ideas for trials and for crowd-sourcing farmer input into research investment decisions (for example, Cole and Fernando 2021).

Cheaper Inputs

Many African farmers face extremely high fertilizer prices. Over 60 percent of the fertilizer in Africa was imported in 2016 (Africa Fertilizer Market Development Report 2018) and this, along with extremely high transport costs and low population density on the continent, implies many farmers often face extremely high fertilizer prices, which makes fertilizer unprofitable for them as a technology. For example, fertilizer prices in some African countries (Cedrez et al. 2020) are well over double those in the United States, as shown in Figure 7. Policy can try to target these costs directly in a multitude of ways. First, investing in infrastructure that will lower transport costs not only lowers the prices of inputs but also potentially improves market access for outputs, as lower costs of search and trading lead to improved market integration (Casaburi et al. 2013; Newman et al. 2018). Another alternative would be to encourage more local production of fertilizer or more geographically local trade in fertilizer. As of January 2021, there were only 135 fertilizer plants across all of sub-Saharan Africa (excluding South Africa): 17 are manufacturing plants and 101 are processing plants (IFDC 2021). In some countries, the price of fertilizer may be high because a small number of importers dominate the market, making it not very competitive. In these cases, some combination of short-term assistance to entrants or pro-competitive policy interventions may be needed.

Policy can additionally try to target these costs directly in a multitude of ways. One obvious approach would be to subsidize fertilizer, but it is not clear that such subsidies are beneficial in the long term or are a sustainable policy tool to increase adoption and yields. Long-term subsidies on nitrogen-based fertilizers have sometimes led to systematic (and potentially serious) overuse and hence longer term soil fertility issues, as Kishore et al. (2021) document in Bangladesh, India, Nepal, and Sri Lanka.

Supporting other inputs may be useful as well. The earlier discussion noted that although rainfall insurance affects the use of technology, farmers are still largely

Figure 7
Fertilizer Prices across Countries in Africa and the United States (2016)
(in US dollars/kilogram)

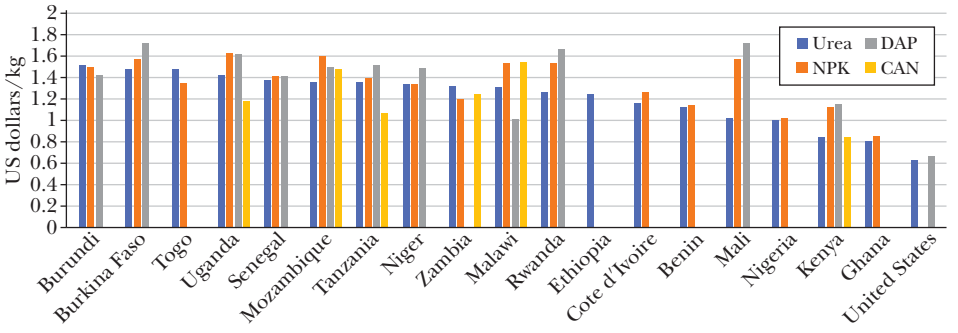

Source: Cedrez et al. (2020) for Africa; USDA Economic Research Service (2019) for the United States.

not willing to pay for it at actuarially fair prices (J-PAL 2016). Subsidies for such insurance could become one of the basic social welfare tools for these economies. Investments in improved weather forecasting could also reduce weather risks more directly.

Conclusion: Pressing Questions

When it comes to improving agricultural technology and productivity in Africa, we have learned much about what doesn't work, which is valuable. However, it also makes apparent how much we do not know. Here are some pressing questions.

First, we do not know enough about how to provide incentives to either the public or the private sector for the development of new agricultural technologies that are locally customized or how to provide incentives for experimentation with these technologies. Rao et al. (2019) estimate the returns to agricultural research and development and show that they remain as high as ever.

Second, are improvements in agricultural technology and productivity the most useful way for raising the standard of living and creating a path out of poverty, or should the focus be on investments in the non-agricultural sector? There is some work showing tight connections between these two sectors (for example, Haggblade et al. 2010; Gollin et al. 2021). However, it is still an open question as to whether the most productive policies would seek to change the returns to non-agricultural investments, and in this way to draw labor out of agriculture and facilitate structural transformation, or whether direct investment in increasing agricultural productivity should be the focus.

Third, can the integration of rural and urban markets in Africa provide better incentives to farmers? A lot of urban food production comes from imports, so there may be a role for the demand side and better market integration in driving technology adoption to replace these imports with locally produced goods (for a review, see de Janvry and Sadoulet 2020). Creating market incentives that remunerate quality, especially for high value crops, may be one step towards sparking this demand side (Bernard et al. 2017).

Fourth, there is very little irrigated farmland across Africa. Many farmers use small-scale pump and hand irrigation along streams for market gardens, especially in peri-urban areas of Africa. More than half of irrigated land in Liberia, Sierra Leone, Ghana, and Nigeria is smallholder informal irrigation, not reliant on fixed infrastructure (Drechsel et al. 2006). But the aggregate area covered by this smallholder irrigation remains tiny. There is a dearth of quantitative information on the availability and characteristics of groundwater across sub-Saharan Africa, but it appears that the potential for high water yield boreholes for irrigation is not widespread, so there may be important geological constraints (Xu et al. 2019). Population density may also be a major constraint to large infrastructure investments required for surface irrigation. If these constraints are binding in certain areas, is there a way to scale down large-scale infrastructure investments? An example of what we have in mind would be the small-scale wet coffee mills that TechnoServe helped cooperatives build across East Africa (IPE 2017).

Fifth, we still know little about the role of the state in agriculture. In particular, there is not much work on crony capitalism in agriculture, or the political economy around how policy priorities or large infrastructure investments are decided specifically in agriculture.

Sixth, there is more to learn about some of the constraints we highlighted above, especially when it comes to labor, land markets, and the environment. A number of multifaceted interventions have proven to be cost effective in some settings: for example, One Acre Fund in Kenya and the Niger Economic Inclusion Program, but much remains to be learned in identifying which constraints affect which farmers.

Finally, there are many studies about how climate change is likely to affect agricultural output in Africa and around the world. But we have relatively little understanding of how farmers and entire agricultural systems might adapt to these changes.

■ *We would like to thank Preksha Jain and Mansa Saxena for fantastic research assistance, and Lorenzo Casaburi, Alain de Janvry, Douglas Gollin, and Jeremy Magruder for valuable comments.*

References

Appiah, Ernest, Kate Baldwin, Dean S. Karlan, and Christopher Udry. 2020. "How Political Insiders Lose Out When International Aid Underperforms: Evidence from a Participatory Development Experiment in Ghana." Unpublished.

Ashour, Maha, Daniel Orth Gilligan, Jessica Blumer Hoel, and Naureen Iqbal Karachiwalla. 2019. "Do Beliefs About Herbicide Quality Correspond with Actual Quality in Local Markets? Evidence from Uganda." *Journal of Development Studies* 55 (6): 1285–1306.

Ashraf, Nava, Xavier Giné, and Dean Karlan. 2009. "Finding Missing Markets (And a Disturbing Epilogue): Evidence from an Export Crop Adoption and Marketing Intervention in Kenya." *American Journal of Agricultural Economics* 91 (4): 973–90.

Atkin, David, and Dave Donaldson. 2015. "Who's Getting Globalized? The Size and Implications of Intra-National Trade Costs." NBER Working Paper 21439.

Banerjee, Abhijit, Michael Faye, Alan Krueger, Paul Niehaus, and Tavneet Suri. 2021. "Effects of a Universal Basic Income in Kenya." Unpublished.

Barrett, Christopher B., Thomas Reardon, Johan Swinnen, and David Zilberman. Forthcoming. "Agri-Food Value Chain Revolutions in Low-and Middle-Income Countries." *Journal of Economic Literature.*

Benami, Elinor, Zhenong Jin, Michael R. Carter, Aniruddha Ghosh, Robert J. Hijmans, Andrew Hobbs, Benson Kenduiywo, and David B. Lobell. 2021. "Uniting Remote Sensing, Crop Modelling and Economics for Agricultural Risk Management." *Nature Reviews Earth and Environment* 2 (2): 140–59.

BenYishay, Ariel, and A. Mushfiq Mobarak. 2019. "Social learning and Incentives for Experimentation and Communication." *Review of Economic Studies* 86 (3): 976–1009.

Berthold, Herrendorf, Richard Rogerson, and Ákos Valentinyi. 2014. "Growth and Structural Transformation." In *Handbook of Economic Growth,* Vol 2B, edited by Philippe Aghion and Steven N. Durlauf, 855–941. Amsterdam: Elsevier.

Bergquist, Lauren Falcao, and Michael Dinerstein. 2020. "Competition and Entry in Agricultural Markets: Experimental Evidence from Kenya." *American Economic Review* 110 (12): 3705–47.

Bernard, Tanguy, Alain De Janvry, Samba Mbaye, and Elisabeth Sadoulet. 2017. "Expected Product Market Reforms and Technology Adoption by Senegalese Onion Producers." *American Journal of Agricultural Economics* 99 (4): 1096–1115.

Beye, Gora. 2002. *Impact of Foreign Assistance on Institutional Development of National Agricultural Research Systems in Sub-Saharan Africa.* Washington, DC: Food & Agriculture Organization of the United Nations.

Binswanger-Mkhize, Hans, and Alex F. McCalla. 2010. "The Changing Context and Prospects for Agricultural and Rural Development in Africa." In *Handbook of Agricultural Economics,* Vol 4, edited by Prabhu Pingali and Robert Evenson, 3571–3712. Amsterdam: Elsevier.

Bold, Tessa, Kayuki C. Kaizzi, Jakob Svensson, and David Yanagizawa-Drott. 2017. "Lemon Technologies and Adoption: Measurement, Theory and Evidence from Agricultural Markets in Uganda." *Quarterly Journal of Economics* 132 (3): 1055–1100.

Boserup, Ester. 1965. *The Conditions of Agricultural Growth: The Economics of Agrarian Change under Population Pressure.* London: Earthscan.

Bossuroy, Thomas, Markus Goldstein, Dean Karlan, Harounan Kazianga, William Pariente, Patrick Premand, Catherine Thomas, Christopher Udry, Julia Vaillant, and Kelsey Wright. 2021. "Pathways out of Extreme Poverty: Tackling Psychosocial and Capital Constraints with a Multi-faceted Social Protection Program in Niger." World Bank Policy Research Working Paper 9562.

Boucher, Steve, Michael Carter, Olaf Erenstein, Travis Lybbert, Jonathan Malacarne, Paswel Marenya, and Laura Paul. 2019. *Bundling Innovative Risk Management Technologies to Accelerate Agricultural Growth and Improve Nutrition.* Washington, DC: Feed the Future Innovation Lab for Markets, Risk & Resilience.

Bridle, Leah, Jeremy Magruder, Craig McIntosh, and Tavneet Suri. 2019. "Experimental Insights on the Constraints to Agricultural Technology Adoption." Unpublished.

Burke, William J., Sieglinde S. Snapp, and Thom S. Jayne. 2020. "An In-Depth Examination of Maize Yield Response to Fertilizer in Central Malawi Reveals Low Profits and Too Many Weeds." *Agricultural Economics* 51 (6): 923–40.

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

Bustos, Paula, Gabriel Garber, and Jacopo Ponticelli. 2020. "Capital Accumulation and Structural Transformation." *Quarterly Journal of Economics* 135 (2): 1037–94.

Caldwell, Richard, Rachel Lambert, Jeremy Magruder, Craig McIntosh, and Tavneet Suri. 2019. *Improving Agricultural Extension and Information Services in the Developing World*. Durham, UK: VoxDev.

Carter, Michael, Alain de Janvry, Elisabeth Sadoulet, and Alexandros Sarris. 2017. "Index Insurance for Developing Country Agriculture: A Reassessment." *Annual Review of Resource Economics* 9 (1): 421–38.

Casaburi, Lorenzo, Rachel Glennerster, and Tavneet Suri. 2013. "Rural Roads and Intermediated Trade: Regression Discontinuity Evidence from Sierra Leone." SSRN Working Paper 2161643.

Cedrez, Camila B., Jordan Chamberlin, Zhe Guo, and Robert J. Hijmans. 2020. "Spatial Variation in Fertilizer Prices in Sub-Saharan Africa." *PloS one* 15 (1): e0227764.

CGIAR ASTI. 2021. ASTI. https://www.asti.cgiar.org/data (accessed June 21, 2021).

Chambers, Robert. 2013. *Rural Development: Putting the Last First*. New York: Routledge.

Christiaensen, Luc, and Lionel Demery. 2017. *Agriculture in Africa: Telling Facts from Myths*. Washington, D.C.: World Bank Group.

Claassen, Roger, and Richard E. Just. 2011. "Heterogeneity and Distributional Form of Farm-Level Yields." *American Journal of Agricultural Economics* 93 (1): 144–60.

Cole, Shawn A., and A. Nilesh Fernando. 2021. "'Mobile'izing Agricultural Advice: Technology Adoption, Diffusion and Sustainability." *Economic Journal* 131 (633): 192–219.

Collier, Paul, and Stefan Dercon. 2014 "African Agriculture in 50 Years: Smallholders in a Rapidly Changing World?" *World Development* 63 (1): 92–101.

De Janvry, Alain, and Elisabeth Sadoulet. 2010. "Agricultural Growth and Poverty Reduction: Additional Evidence." *World Bank Research Observer* 25 (1): 1–20.

De Janvry, Alain, and Elisabeth Sadoulet. 2020. "The Puzzle of Lagging Sub-Saharan Africa Agriculture: Toward a Theory of Connectedness." Unpublished.

De Janvry, Alain, Elisabeth Sadoulet, and Tavneet Suri. 2017. "Field Experiments in Developing Country Agriculture." In *Handbook of Economic Field Experiments*, Vol 2, edited by Abhijit Vinayak Banerjee and Esther Duflo, 427–66. North-Holland: Elsevier.

Dercon, Stefan, and Douglas Gollin. 2014. "Agriculture in African Development: Theories and Strategies." *Annual Review of Resource Economics* 6 (1): 471–92.

Deutschmann, Joshua W., Maya Duru, Kim Siegal, and Emilia Tjernström. 2019. "Can Smallholder Extension Transform African Agriculture?" NBER Working Paper 26054.

Drechsel, Pay, Sophie Graefe, Moise Sonou, and Olufunke O. Cofie. 2006. *Informal Irrigation in Urban West Africa: An Overview*. Colombo, Sri Lanka: International Water Management Institute.

Druilhe, Zoé, and Jesús Barreiro-Hurlé. 2012. "Fertilizer Subsidies in sub-Saharan Africa." ESA Working Paper 12-04.

Emerick, Kyle, Alain de Janvry, Elisabeth Sadoulet, and Manzoor H. Dar. 2016. "Technological Innovations, Downside Risk, and the Modernization of Agriculture." *American Economic Review* 106 (6): 1537–61.

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

Evenson, Robert E., and Larry E. Westphal. 1995. "Technological Change and Technology Strategy." In *Handbook of Development Economics,* Vol 3A, edited by Jere Behrman and T.N. Srinivasan, 2209–99. North Holland: Elsevier.

Foster, Andrew D., and Mark R. Rosenzweig. 2017. "Are There Too Many Farms in the World? Labor-Market Transaction Costs, Machine Capacities and Optimal Farm Size." NBER Working Paper 23909.

Fuglie, Keith. 2016. "The Growing Role of the Private Sector in Agricultural Research and Development World-Wide." *Global Food Security* 10 (1): 29–38.

Glennerster, Rachel, and Tavneet Suri. 2018. "Agricultural Technology and Nutrition in Sierra Leone." Unpublished.

Gollin, Douglas, Casper Worm Hansen, and Asger Wingender. 2021. "Two Blades of Grass: The Impact of the Green Revolution." *Journal of Political Economy* 129 (8): 2344–84.

Gollin, Douglas, David Lagakos, and Michael E. Waugh. 2014. "Agricultural Productivity Differences Across Countries." *American Economic Review* 104 (5): 165–70.

Gollin, Douglas, and Christopher Udry. 2021. "Heterogeneity, Measurement Error, and Misallocation: Evidence from African Agriculture." *Journal of Political Economy* 129 (1): 1–80.

Grant, Matthew, and Meredith Startz. 2021. "Cutting Out the Middleman: The Structure of Chains of Intermediation." Unpublished.

Greatrex, Helen, James Hansen, Samantha Garvin, Rahel Diro, Sari Blakely, Margot Le Guen, Kolli Rao, and Daniel Osgood. 2015. *Scaling Up Index Insurance for Smallholder Farmers: Recent Evidence and Insights*. Copenhagen, Denmark: CGIAR Research Program on Climate Change, Agriculture and Food Security (CCAFS).

Griliches, Zvi. 1957. "Hybrid Corn: An Exploration in the Economics of Technological Change." *Econometrica* 25 (1): 501–22.

Haggblade, Steven, Peter Hazell, and Thomas Reardon. 2010. "The Rural Non-Farm Economy: Prospects for Growth and Poverty Reduction." *World Development* 38 (10): 1429–41.

Harou, Aurélie P., Yanyan Liu, Christopher B. Barrett, and Liangzhi You. 2017. "Variable Returns to Fertiliser Use and the Geography of Poverty: Experimental and Simulation Evidence from Malawi." *Journal of African Economies* 26 (3): 342–71.

Harou, Aurélie P., Malgosia Madajewicz, Hope Michelson, Cheryl A. Palm, Nyambilila A. Amuri, Christopher Magomba, Johnson M. Semoka, Kevin Tschirhart, and Ray Weil. 2020. "Increasing Maize Yields with Soil Testing and Subsides in Tanzania." Unpublished.

Heisey, Paul W., and Keith O. Fuglie. 2018. *Agricultural Research Investment and Policy Reform in High-Income Countries*. Washington, DC: United States Department of Agriculture.

Hengl, Tomislav, Matthew A.E. Miller, Josip Križan, Keith D. Shepherd, Andrew Sila, Milan Kilibarda, Ognjen Antonijević et al. 2021. "African Soil Properties and Nutrients Mapped at 30 M Spatial Resolution Using Two-Scale Ensemble Machine Learning." *Scientific Reports* 11 (1): 1–18.

Hill, Polly. 1963. *Migrant Cocoa-Farmers of Southern Ghana: A Study in Rural Capitalism*. London: Cambridge University Press.

Hurley, Terrance M., Gary L. Malzer, and Bernard Kilian. 2004. "Estimating Site-Specific Nitrogen Crop Response Functions." *Agronomy Journal* 96 (5): 1331–43.

IFDC. 2021. *Register of Fertilizer: Manufacturing and Processing Facilities in Sub-Saharan Africa (5th Edition)*. Muscle Shoals, Alabama: IFDC.

Indian Council of Agricultural Research (ICAR) and Indian Institute of Maize Research (IIMR). 2019. *List of Cultivars of Maize Developed by AICRPM and Released for Cultivation in India Since Inception of Project*. Punjab, India: Indian Institute of Maize Research.

IPE Triple Line. 2017. *Evaluation of the TechnoServe East Africa Coffee Initiative*. London: IPE Triple Line.

Jayne, Thomas S., Jordan Chamberlin, Lulama Traub, Nicholas Sitko, Milu Muyanga, Felix K. Yeboah, Ward Anseeuw et al. 2016. "Africa's Changing Farm Size Distribution Patterns: The Rise of Medium-Scale Farms." *Agricultural Economics* 47 (S1): 197–214.

J-PAL, CEGA, and ATAI Policy Bulletin. 2016. *Make it Rain*. Cambridge, MA: Abdul Latif Jameel Poverty Action Lab, Center for Effective Global Action, and Agricultural Technology Adoption Initiative.

Karanja, Daniel David. 1996. "An Economic and Institutional Analysis of Maize Research in Kenya." Food Security International Development Working Paper 54693.

Kenya Plant Health Inspectorate Service (KEPHIS). 2020. "National Crop Variety List." https://kephis.org/images/pdf-files/UPDATED%202020%20August%20NATIONAL%20VARIETY%20LIST1.pdf. (accessed June 21, 2021).

Kishore, Avinash, Muzna Alvi, and Timothy J. Krupnik. 2021. "Development of Balanced Nutrient Management Innovations in South Asia: Perspectives from Bangladesh, India, Nepal, and Sri Lanka." *Global Food Security* 28: 100464.

Laajaj, Rachid, Karen Macours, Cargele Masso, Moses Thuita, and Bernard Vanlauwe. 2020. "Reconciling Yield Gains in Agronomic Trials with Returns Under African Smallholder Conditions." *Scientific Reports* 10 (1): 14286.

Lipton, Michael. 1988. "The Place of Agricultural Research in the Development of sub-Saharan Africa." *World Development* 16 (10): 1231–57.

Lobell, David B., George Azzari, Marshall Burke, Sydney Gourlay, Zhenong Jin, Talip Kilic, and Siobhan Murray. 2020. "Eyes in the Sky, Boots on the Ground: Assessing Satellite-and Ground-Based Approaches to Crop Yield Measurement and Analysis." *American Journal of Agricultural Economics* 102 (1): 202–19.

Lowder, Sarah K., Jakob Skoet, and Terri Raney. 2016. "The Number, Size, and Distribution of Farms, Smallholder Farms, and Family Farms Worldwide." *World Development* 87 (1): 16–29.

Manning, Patrick. 2014. "African Population, 1650-2000: Comparisons and Implications of New Estimates" in *Africa's Development in Historical Perspective*, edited by Emmanuel Akyeampong, Robert Bates, Nathan Nunn and James Robinson, 131–52. Cambridge: Cambridge University.

McCullough, Ellen B., Julianne D. Quinn, Andrew M. Simons, Leslie J. Verteramo, and Joshua Woodard. 2020. "Modeling Climate Smart Soil Health Investments in Sub-Saharan Africa." Unpublished.

Michelson, Hope, Anna Fairbairn, Brenna Ellison, Annemie Maertens, and Victor Manyong. 2021. "Misperceived Quality: Fertilizer in Tanzania." *Journal of Development Economics* 148 (1): 102579.

Moscona, Jacob, and Karthik Sastry. 2021. "Inappropriate Technology: Evidence from Global Agriculture." SSRN Working Paper 3886019.

Munshi, Kaivan. 2004. "Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution." *Journal of Development Economics* 73 (1): 185–213.

Newman, Greg, Andrea Wiggins, Alycia Crall, Eric Graham, Sarah Newman, and Kevin Crowston. 2012. "The Future of Citizen Science: Emerging Technologies and Shifting Paradigms." *Frontiers in Ecology and the Environment* 10 (6): 298–304.

Newman, Neil, Lauren Falcao Bergquist, Nicole Immorlica, Kevin Leyton-Brown, Brendan Lucier, Craig McIntosh, John Quinn, and Richard Ssekibuule. 2018. "Designing and Evolving an Electronic Agricultural Marketplace in Uganda." In *COMPASS '18: Proceedigs of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies,* ed. Ellen Zegura, 1–11. New York, NY: Association for Computing Machinery.

Pearson, Calvin H., and Amaya Atucha. 2015. "Agricultural Experiment Stations and Branch Stations in the United States." *Natural Sciences Education* 44 (1): 1–5.

Poets, Ana, Kevin Silverstein, Philip Pardey, Sarah Hearne, and James Stevenson. 2020. *DNA Fingerprinting for Crop Varietal Identification: Fit-for-Purpose Protocols, their Costs and Analytical Implications.* Rome, Italy: CGIAR Standing Panel on Impact Assessment (SPIA).

Porteous, Obie. 2019. "High Trade Costs and Their Consequences: An Estimated Dynamic Model of African Agricultural Storage and Trade." *American Economic Journal: Applied Economics* 11 (4): 327–66.

Porteous, Obie. 2020. "Trade and Agricultural Technology Adoption: Evidence from Africa." *Journal of Development Economics* 144: 102440.

Rao, Xudong, Terrance M. Hurley, and Philip G. Pardey. 2019. "Are Agricultural R&D Returns Declining and Development Dependent?" *World Development* 122 (1): 27–37.

Rawat, Stuti. 2020. "Global Volatility of Public Agricultural R&D Expenditure." In *Advances in Food Security and Sustainability,* Vol 5, edited by Marc J. Cohen, 119–43. Cambridge, MA: Academic Press.

Restuccia, Diego, and Raul Santaeulalia-Llopis. 2017. "Land Misallocation and Productivity." NBER Working Paper 23128.

Rosenzweig, Mark R., and Christopher Udry. 2020. "External Validity in a Stochastic World: Evidence from Low-Income Countries." *Review of Economic Studies* 87 (1): 343–81.

Sanabria, Joaquin, Georges Dimithe, and Emmanuel K.M. Alognikou. 2013. *The Quality of Fertilizer Traded in West Africa: Evidence for Stronger Control.* Muscle Shoals, AL: International Fertilizer Development Center.

Sanchez, Pedro A. 2019. *Properties and Management of Soils in the Tropics.* London: Cambridge University Press.

Schimmelpfennig, David. 2016. *Farm Profits and Adoption of Precision Agriculture.* Washington, DC: United States Department of Agriculture.

SeedNet India Portal. 2021. "Seed Division: Gazette Notification S.O. 500(E)." https://seednet.gov.in/CMS/SeedGo/85th%20Gazzette%20Notification.pdf. (accessed June 2, 2021).

Sheahan, Megan, and Christopher B. Barrett. 2017. "Food Loss and Waste in Sub-Saharan Africa." *Food Policy* 70 (1): 1–12.

Smale, Melinda, and Véronique Thériault. 2018. "A Cross-Country Summary of Fertilizer Subsidy Programs in sub-Saharan Africa." Feed the Future Innovation Lab for Food Security Policy Research Paper 169.

Startz, Meredith. 2021. "The Value of Face-to-Face: Search and Contracting Problems in Nigerian Trade." Unpublished.

Stoorvogel, Jetse J., Lammert Kooistra, and Johan Bouma. 2015. "Managing Soil Variability at Different Spatial Scales as a Basis for Precision Agriculture." In *Soil-Specific Farming: Precision Agriculture,* edited by Rattan Lal and B. A. Stewart, 37–72. Boca Raton, FL: CRC Press.

Suri, Tavneet. 2011. "Selection and Comparative Advantage in Technology Adoption." *Econometrica* 79 (1): 159–209.

Suri, Tavneet, and Christopher Udry. 2021. African Agricultural Technology. https://voxdev.org/voxdevlit.

Teravaninthorn, Supee, and Gaël Raballand. 2009. *Transport Prices and Costs in Africa: A Review of the International Corridors*. Washington, DC: World Bank.

Tjernström, Emilia, Michael Carter, and Travis Lybbert. 2015. "The Dirt on Dirt: Soil Characteristics and Variable Fertilizer Returns in Kenyan Maize Systems." Unpublished.

The Food and Agricultural Organization of the United Nations (FAOSTAT). https://www.fao.org/faostat/en/#data. (accessed June 21, 2021).

U.S. Department of Agriculture. 2019. "Fertilizer Use and Price." Economics Research Service. https://www.ers.usda.gov/data-products/fertilizer-use-and-price.aspx. (accessed June 2, 2021).

U.S. Department of Agriculture. 2020. "Occupational Employment and Wage Statistics." Economic Research Service. https://www.bls.gov/oes/current/oes_stru.htm#45-0000. (accessed June 21, 2021).

Van Asselt, Joanna, Federica Di Battista, Shashi L. Kolavalli, and Christopher R. Udry. 2018. "Agronomic Performance of Open Pollinated and Hybrid Maize Varieties: Results from On-Farm Trials in Northern Ghana." Unpublished.

Xu, Yongxin, Paul Seward, Cheikh Gaye, Lixiang Lin, and Daniel O. Olago. 2019. "Preface: Groundwater in sub-Saharan Africa." *Hydrogeology Journal* 27 (3): 815–22.

Journal of Economic Perspectives—Volume 36, Number 1—Winter 2022—Pages 57–80

Time Use and Gender in Africa in Times of Structural Transformation

Taryn Dinkelman and L. Rachel Ngai

Changes in time allocation are central to the process of economic development. As an economy develops, a substantial portion of activities that were once conducted within the household—for example, cooking, producing clothing and household goods, providing child- and eldercare—are outsourced to the market. This shift in activity from home to market has implications for how women and men spend their time. When unpaid household work can be outsourced to the paid market sector, this tends to generate additional compensated female jobs, improving gender equality in labor market outcomes and in the home. In addition, when individuals can choose to allocate their time between home and market based on their talents, the economy as a whole benefits from reduced labor misallocation.

The process of structural transformation, in which production shifts across the agricultural, manufacturing, and services sectors, is linked to the emergence of new types of labor market opportunities for women. Typically, as an economy grows and the service sector expands, home production can be more easily outsourced to the market via a process of marketization. Two market jobs are often created: one for

■ *Taryn Dinkelman is Associate Professor of Economics, University of Notre Dame, Notre Dame, Indiana. L. Rachel Ngai is Associate Professor of Economics, London School of Economics, London, United Kingdom. Dinkelman is a Faculty Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts and a BREAD Research Fellow; Ngai is a Member, Centre for Macroeconomics, London, United Kingdom; both authors are Research Associates, Centre for Economic Policy Research, London, United Kingdom.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.36.1.57.

the woman who is able to reduce her time in home production and another in the service sector, which itself tends to be intensive in female labor.

In this paper, we examine these links between time allocation, structural transformation, and marketization of home services in Africa using both macro and micro data. It is impossible to think about women's market work in a growing economy without thinking about women's work in the home and how the latter constrains or facilitates the former. Our aim is not to speculate on potential sources of economic growth in African countries or the triggers for structural transformation, but rather to draw attention to the understudied links between home and market work in Africa and to discuss potential barriers to the smooth transition of time out of the home and into market work when the economy grows. A key point we emphasize is that tracking changes in the allocation of time across the development process will require more time use surveys collected from a wider range of African countries and for more years in each country.

We start by defining home production and provide a macro-level picture of how female market work has shifted across sectors in Africa since the 1970s. Over the last 50 years, the share of female workers in agriculture on the continent has shrunk, while female employment in services has grown.

Next, we draw out two patterns about the nature of women's work in Africa. First, women in North African countries have especially low labor market participation rates and very low market hours. Second, in much of sub-Saharan Africa, high female labor force participation coexists with low average market hours and high home production hours.[1] The majority of women's market work in Africa is unpaid, occurring on family farms and in own-account or family firms. Family farms and family firms allow women to combine home and market work at the same location; this pattern of female employment was also evident in the mid-19th century United States (Ruggles 2015).

Because so many hours are devoted to home production each day among African women, we investigate the nature of this work. A key contribution of this paper is that we assemble microdata from time-use surveys across four African countries at different levels of development to show that the allocation of home production hours across different activities closely resembles historical time use patterns in the United States in the 1920s and 1960s.

With these key features in hand, we discuss some of the constraints and opportunities involved in moving towards a more efficient and equitable allocation of labor time in Africa. We highlight two key frictions that could slow down the marketization of home production in Africa: 1) much of sub-Saharan Africa lacks infrastructure to manage necessary household work, including physical infrastructure for household utilities and transportation to the market, and child- and family-care services; and 2) in North African countries in particular, social norms about women's market work restrict the reallocation of female time across home and market.

[1] Heath and Jayachandran (2018) provide a recent review of female labor force participation patterns across the developing world, updating the work of Mammen and Paxson (2000) in this journal.

Although most high-income economies have already passed through the stages of structural transformation, many African countries are in the early stages of this process. There is still a large knowledge gap regarding women's time use and time allocation on the continent. Observing how women in Africa allocate time between home and market provides an opportunity to deepen our understanding of the process of structural transformation. Moreover, because many African countries are still in the early stages of structural transformation, understanding and addressing barriers to women's time reallocation from home to market across these countries is a first-order problem with implications for aggregate output and welfare.

Home Production and Structural Transformation

Time in home production is defined as the time spent on the production of goods and services for own (not market) use. Reid (1934, p. 11) provides a still-relevant definition:

> It consists of those unpaid activities which are carried on, by and for the members, which activities might be replaced by market goods, or paid services, if circumstances such as income, market conditions, and personal inclinations permit the service being delegated to someone outside the household group.

As Reid (1934, p. 47) observed, many home activities in now-developed economies have been transferred to the market through the process of *marketization*:

> As time went on, one form of production after another, spinning, weaving, sewing, tailoring, baking, butchering, soap-making, candle-making, brewing, pre-serving, laundering, dyeing, gardening, care of poultry, and other tasks have wholly or in part been transferred to commercial production. In addition, childcare, education, and the care of the sick are now to a large extent carried on by paid workers.

Typical modern-day examples of these home production activities are cooking, cleaning, laundry, and child- and eldercare. In many African countries, collecting wood and water for use in the household is also part of home production.[2]

The marketization process that Reid described is closely linked to structural transformation of a developing economy, namely the decline in agriculture, a hump-shaped rise and fall in manufacturing employment, and rising services along the development path. Such patterns have been well-documented both historically for

[2]Across African countries, there is ambiguity about whether wood and water collection should appear in the System of National Accounts (SNA) or not. For the purposes of this article, we count time spent in these activities as part of home production.

developed countries and over large cross-sections of advanced countries (Kuznets 1966; Maddison 1980; Herrendorf, Rogerson, and Valentinyi 2013). Lebergott (1993) documented this process of marketization with consumption expenditure data for the United States. Using harmonized cross-country time use data for recent decades, Freeman and Schettkat (2005) and Burda, Hamermesh, and Weil (2013) documented the importance of marketization of home production in understanding market hours across rich countries.

There is some debate in the literature about whether Africa's structural transformation follows typical patterns.[3] Figure 1 shows how female employment in African countries has shifted across sectors since the 1970s. Each panel shows the sectoral employment shares for women in eleven African countries (each country time series shown in different colored markers) at different levels of real GDP per capita between 1970 and 2010. Data are annual, employment shares are computed over all women in the labor market, and workers include paid and unpaid workers (self-employed and family workers). The patterns in Figure 1 indicate that in countries and years in which income is higher, female employment shares in agriculture are lower, while employment in manufacturing (to a lesser extent) and in services (to a greater extent) is higher. Manufacturing jobs account for only a small share of female employment on the continent: de Vries et al. (2021) note that under 10 percent of all female jobs in Africa are in manufacturing.[4]

The structural transformation in female jobs occurring in many African countries involves a shift from agriculture directly into the service sector. This is similar to the structural shifts in female employment in historical US data. For example, assembling data from 1860 to 2000, Ngai et al. (2021) document that during this 140-year period, jobs for American women moved out of the agricultural sector and into the service sector, and female employment shares in manufacturing remained low throughout.

How is the shift of female employment towards services during the structural transformation linked to home production? Historically, in high-income countries, the expansion of service sectors has enabled households to outsource home production via the process of marketization (Ngai and Pissarides 2008; Rogerson 2008). Using 136 time-use surveys from 43 countries over the period 1960–2014, Bridgman, Duernecker, and Herrendorf (2018) show that growth in market services across the globe is associated with a decline in home production hours. Moreover, growth in the service sector historically generated jobs intensive in both high-skilled and low-skilled female labor. These changes had important implications for the economic

[3] Rodrik (2016) observes that today's developing countries see their manufacturing sector shrink at lower levels of GDP per capita than in developed countries historically, a feature termed "premature de-industrialization." However, using more recent data, de Vries et al. (2021) show that patterns of deindustrialization observed in Africa in the 1990s seem to have reversed since the late 2000s.

[4] Male employment exhibits similar patterns across income levels and sectors (see online Data Appendix, available with this articles at the *JEP* website, for Figure 5.1), with one exception: men's employment share in manufacturing exhibits a clearer rise when agriculture declines at low levels of real GDP per capita.

Figure 1

Female Employment Shares by Sector and Level of Development in Sub-Saharan Africa

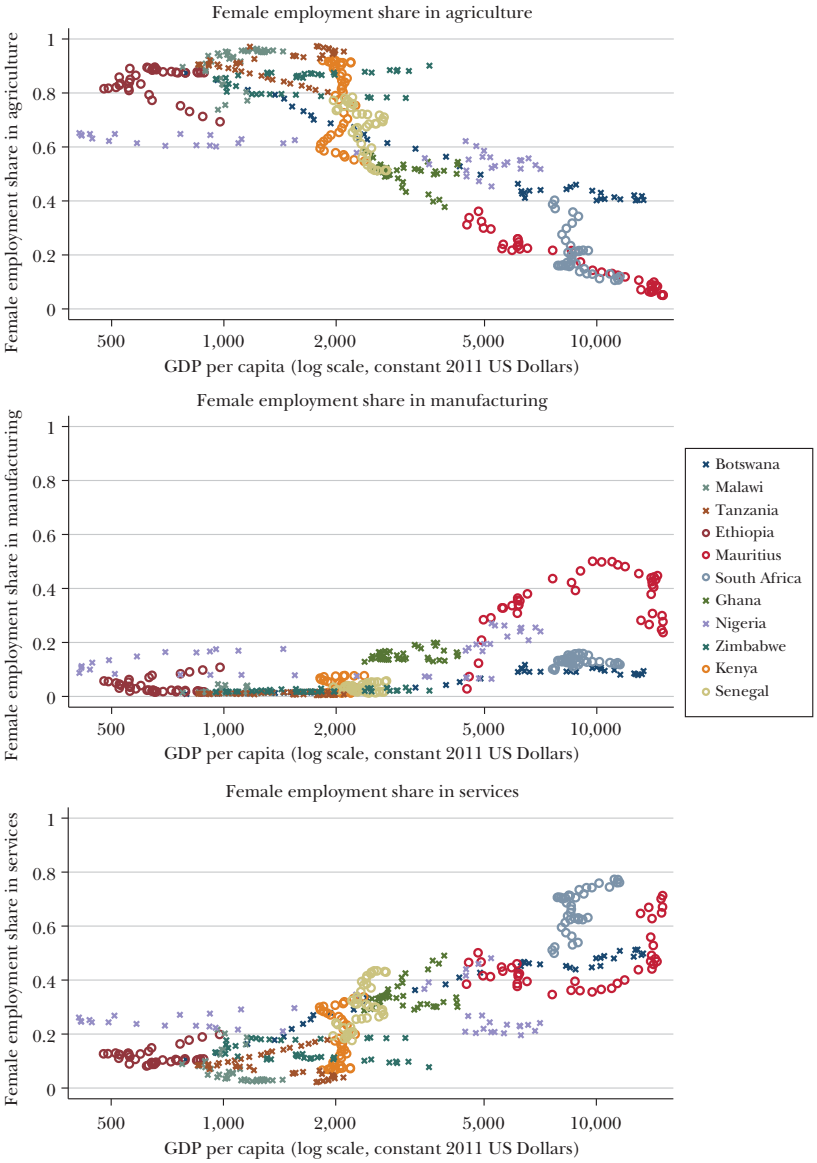

Source: Annual data on female employment shares by sector for eleven African countries from the Africa Sector Database collated by the Groningen Growth and Development Centre. Real GDP per capita (2011 international dollars) from Penn World Tables v9.1.

Note: The eleven countries are Botswana, Ethiopia, Ghana, Kenya, Malawi, Mauritius, Nigeria, Senegal, South Africa, Tanzania, and Zambia. De Vries, Timmer, de Vries (2015) describe the construction of the Africa Sector Database. Gender-specific sectoral employment shares are constructed for each country and year between 1970 and 2010 and include all paid and unpaid workers (self-employed and family workers) enumerated in population Census data and national Labor Force Surveys. See online Data Appendix for more details on figure construction.

role of women.[5] For women, the shift in time use between home and market is as important as the shift across market sectors during the structural transformation. In addition, at the macroeconomic level, misallocation of female time and talent across the home and market sectors can reduce aggregate output (Hsieh et al. 2019; Lee 2020).

How African Women Allocate Time between Home and Market

Participation in market work in Africa is linked to time spent in home production. Figure 2 reports patterns of female labor market participation and time use across market and home and how each varies with GDP in Africa. Figure 2a illustrates how the female labor force participation rate among African countries varies with GDP. We use data on labor force participation from the International Labour Organization (ILO); Gaddis et al. (2020) discuss some of the key challenges to measuring female labor force participation in sub-Saharan Africa.

Although Figure 2 only captures cross-country variation, the variation is consistent with the declining part of the familiar U-shape in female labor force participation with development—that is, labor force participation of women (especially married women) first falls and then rises with per capita GDP (Sinha 1965; Boserup 1970; Durand 1975; Goldin 1986). None of the countries in Figure 2a have reached high-income country status yet, and there is considerable variation in female labor force participation even within income bands. For example, South Africa and Algeria each have relatively higher levels of per capita GDP, but whereas South Africa's female labor force participation rate looks more like other middle- and high-income countries, Algeria's female labor force participation rate is below 20 percent, the lowest in the world. The lack of upturn in the female labor force participation rate that we see in Figure 2a may be related to some of the barriers to women's market work that we discuss below.

Labor force participation does not clearly reveal the extent of female time spent in market activities. Figure 2b illustrates how many hours women spend in market work (green triangles) and household work (orange circles) in five African countries at different levels of GDP per capita. The data show average weekly hours of market and home work across all women, including those who do and do not work in the market (or home) at all. To create this figure, we used the average weekly hours for women computed by Bridgman, Duernecker, and Herrendorf (2018) who in turn compute these averages using microdata from each African country (we use the aggregated data from this paper because not all of the microdata is publicly available).

Comparing 2a with 2b shows that high female labor force participation often coexists with low market hours and high home hours. For example, Uganda,

[5] For more discussion, see the pioneering work of Mincer (1962) on the role of home production for understanding female labor supply. Recent work in quantifying its importance includes Ngai and Petrongolo (2017) and Rendall (2018).

Figure 2

Extensive versus Intensive Margin Labor Market Activity for African Women

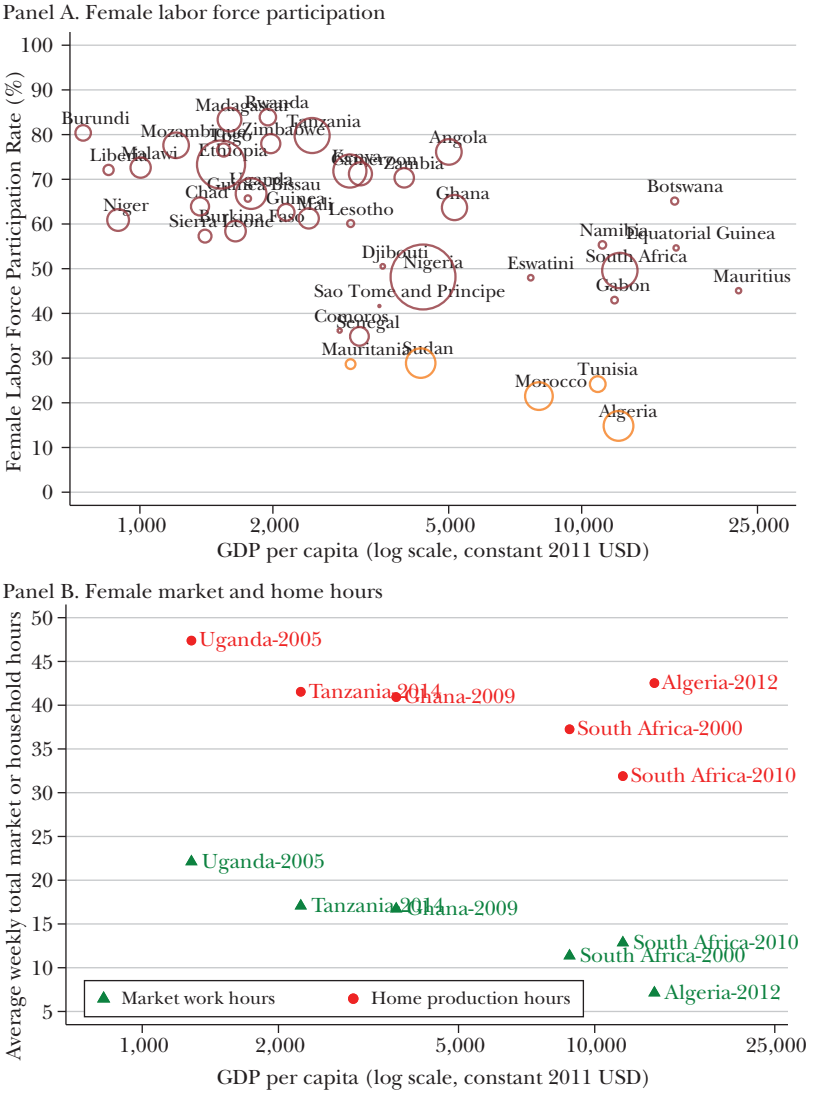

Source: Female labor force participation in panel A is from the International Labour Organization. Average weekly market and home hours for women are from Bridgman, Duernecker, and Herrendorf (2018) who use time use data to compute average hours for all women. Real GDP per capita (2011 international dollars) is from Penn World Tables v9.1.

Note: In panel A, circle sizes represent population, orange circles denote North African countries and red circles denote sub-Saharan African countries. In panel B, average market and home hours are reported for women aged 15 and older. We plot real GDP per capita on a log scale and label *x*-axes with the corresponding income per capita.

Tanzania, and Ghana all have female labor force participation rates in excess of 60 percent. Yet the average woman in these countries works fewer than 25 hours per week in the market, and at least 32 to 48 hours per week in home production.[6] Market work is far from being a full-time job for women in these countries, and home production is at least the equivalent of a part-time job. On the other end of the spectrum, all of the countries with low female labor force participation (below 40 percent) are in North Africa. Figure 2b shows that for one North African country (Algeria), women work less than 10 hours per week and report 40 hours of home time per week.

Figure 2b also shows that home production hours fall as income per capita rises (from $1,785 in Uganda to $12,200 in South Africa). Using a broader cross section of countries across the world, Bridgman, Duernecker, and Herrendorf (2018) show that female home hours decline to 25-30 hours per week once countries reach high-income status. Data from the American Time Use Surveys indicate that in 1965, the average woman in the United States worked 40.6 hours per week in the home; this fell to 28.7 hours per week in 2010 (for a detailed analysis of modern trends in US home production time, see Aguiar and Hurst 2007). Home production hours among African women in the 2000s resemble average hours in home production among American women in the mid-1960s. Most African countries likely have a long way to go before income growth leads to meaningful reductions in home production time. Time use surveys collected for a wider range of African countries and for more years in each country would be immensely useful for tracking these changes in home production hours as countries change and grow.

Together, Figures 2a and 2b illustrate substantial heterogeneity in female labor force participation across regions and show that many women tend to work a little in the market and a lot in home production. We next discuss two questions. Is there something about the nature of work in African countries that allows women to combine home and market activities? Do constraints to women's market work differ across countries where extensive margin participation is very low (like Algeria) compared with those countries where women participate at high rates but with few hours of market work (like Uganda)?

Combining Home and Market Work: The Role of Family Farms and Family Firms

In much of Africa, the large majority of women's time in market work is spent in unpaid family labor or own-account work. Most of these types of jobs—for example, growing vegetables on the family farm for market, producing meals for a family food stand, serving customers in a family business—are not remunerated. Yet these jobs offer flexibility of location and hours. Such jobs typically take place in the family home or close to it, or on the family farm where childcare, eldercare, and other types of home production must also take place. This flexibility makes

[6] As in other countries, women in Africa spend more time in home production than men do (data not shown). Among the countries in Figure 2b, women spend two to four times as many hours in home production as men do.

Figure 3
Composition of Women's Market Work by Region over Time

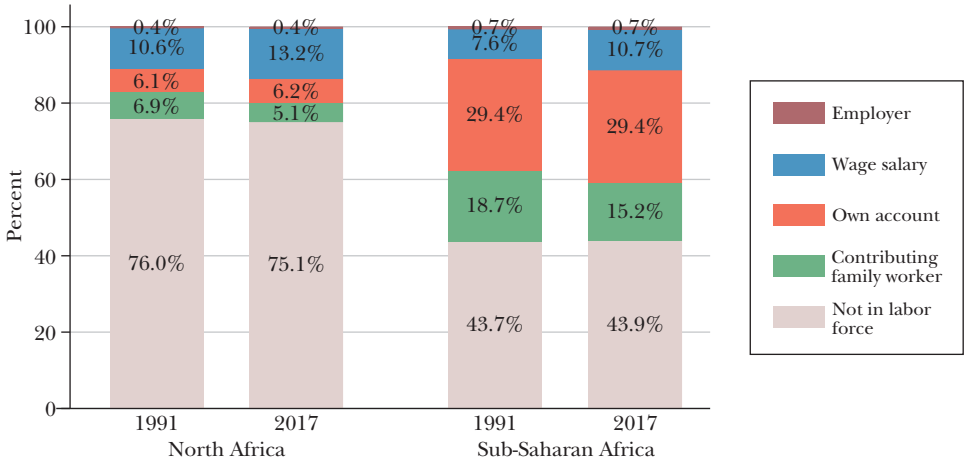

Source: Data on sector of market work and market participation are from the World Bank's World Development Indicators. See online Data Appendix for details on each data series.

it possible to shift time between market and home production at low cost, and obviates the need for women to specialize in either the home sector or the market sector.

In Figure 3, we show two snapshots of the profile of market work for women in 1991 and 2017, separately for North Africa and sub-Saharan Africa. The categories shown in this graph are exhaustive: a woman is either not in the labor force (rose-colored area), or she is part of the labor force and working as an unpaid contributing family worker (green), an own-account or self-employed worker (orange), a wage-worker (navy), or an employer (maroon).

Aggregating across countries, the female labor force participation rate in sub-Saharan Africa is around 56 percent. The vast majority of women doing any market work in sub-Saharan Africa are still own-account workers or contributing family workers, and not wage earners. The figure paints a generally stagnant picture of extensive margin participation and of the types of jobs held by women; over time, there has been a small increase in the share of women doing wage work and a small decline in women working as unpaid contributing family workers. This change is more notable in middle-income African countries, not shown separately here.

The employment pattern in low- and middle-income countries in Africa is consistent with what we see historically in developed countries. For example, Ngai, Olivetti, and Petrongolo (2021) show the role that work on family farms and within-family firms plays in accounting for changes in female labor force participation over time in the United States. As the structural transformation in the West occurred and family farms were consolidated into larger farms, home-based market work

disappeared, and with it, the ability of women to easily combine home and market production. The current types of market work done by women in sub-Saharan Africa suggest that any forces that shift workers off of family farms or lead to consolidation of these smaller farms may also have large impacts on women's time allocation across market and home. If marketization of home services is slow to materialize, moving away from family farms and firms could push women to withdraw from market work and reallocate hours towards home production.

A notable aspect of Figure 3 is the stark difference in employment patterns across North African countries (including Algeria, Egypt, Libya, Morocco, Sudan, Tunisia, Western Sahara) and the rest of the continent. The rate of female labor force participation for North African countries is around 25 percent, with the majority of this work being in wage employment. Family farms and firms do not play a large role in providing market opportunities for women in this region. Differences in social norms about acceptable market work for women may account for these large differences in women's time allocation in North Africa versus other countries. We return to this topic in the final section.

Service-Sector Jobs and Female Market Work in Africa

If structural transformation in Africa is to affect the way that women choose to allocate time across home and market, we need to know whether service sector jobs in Africa produce home production substitutes and generate female-intensive jobs for women across the skills distribution. Evidence on the growth of female employment in these home-production substitute sectors (as a subset of the service sector) exists for developed countries (for example, Addati et al 2018). These topics are an open area for research in African countries.

One way to measure the marketization of home production is to look at the share of service sector jobs that provide home-substitutable goods and services, like childcare workers, healthcare workers, cleaners, hotel workers, and restaurants. To illustrate how we might do this, we collect officially published statistics from South Africa, Kenya, Ghana, and Ethiopia in 2015 to capture female market work by sector. Figure 4 uses these data to graph the share of all female jobs (in blue) and the share of all female service sector jobs (in red) that could be regarded as home substitute activities. For comparison, we also plot female employment in home-substitutable sectors for the United States using similar data. Because the official statistics are aggregated at a fairly high level, we can only implement a broad definition of home-substitutable sectors, including jobs in education, health and social work, arts/entertainment, domestic work for private households, and all other services (including personal services like food and accommodation and miscellaneous repairs).

In all the countries shown here, women's market jobs that substitute for home production represent at least half of all of female jobs and make up the vast majority of service-sector work among women. In Ethiopia, Ghana, and South Africa, a larger share of all female jobs in the service sector are in home substitute work, relative to the United States and Kenya.

Figure 4
Share of Female Market Jobs Producing Home Production Substitutes

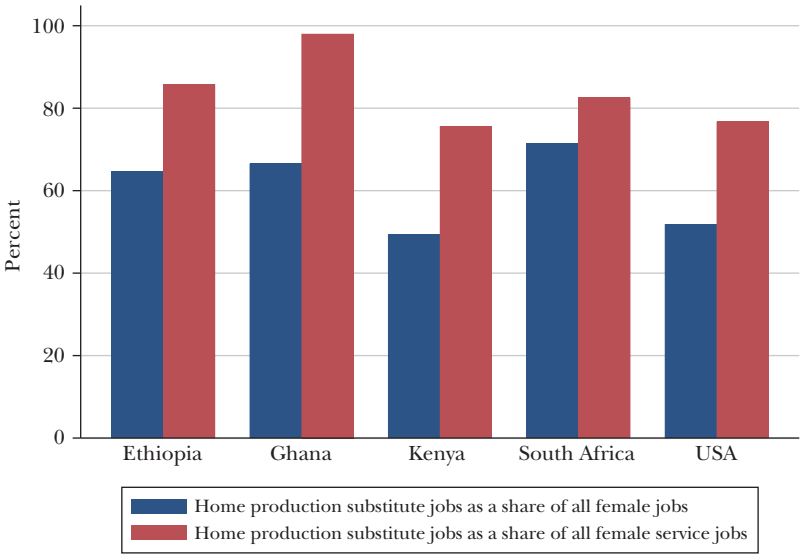

Source: We compiled data from official statistics of each country that were in turn each based on national household or labor force survey data. See online Data Appendix for details of how we classified each type of job.
Note: Share of female employment (total, or in services) that could be broadly classified as home production substitutes. We define home-substitutable jobs to include education, health and social work, arts/entertainment, domestic work for private households, and all other services (including personal services like food and accommodation, and miscellaneous repairs), as well as wholesale and retail trade jobs. Service sector jobs excluded from home-substitutable jobs include finance, insurance, and real estate; administrative and support services; information and communication; and professional activities. Public administration and defense are excluded from the denominator of both measures in Figure 4.

At least in this handful of African countries, the marketization of some types of home production seems well underway. These jobs provide opportunities for women to choose to reallocate time away from home production and towards market work. A natural and useful extension of this exercise would be to expand the sample to a larger number of African countries and to consider how the share of marketized home production substitute jobs has changed over time. An alternative way to calculate the extent of marketization, used in Ngai and Pissarides (2011), would be to begin with time-use surveys and make two calculations: the share of total market hours worked in jobs that produce home production substitutes and the share of total hours spent producing these services within the household (household hours). With multiple time-use surveys within a country, one could track marketization by observing how the location of time required for producing household services shifts out of the household and into the market.

Time in Home Production in Africa

The evidence presented earlier in Figure 2b shows that in a subset of African countries, the average woman spends 32–47 hours each week in home production and only 7–17 hours per week in market work. To understand low market hours in these countries, it is important to have a better understanding of what is happening within the home. We next take a closer look at home production activities in a subset of countries for which we have the microdata from national time use surveys or from household surveys with time use modules. Only a handful of African countries (listed in the online Data Appendix) have ever fielded such time use surveys that include data to estimate the extent of home production, and only since the 1990s. We use micro data from five of these surveys: two waves for South Africa and one each for Morocco, Ghana, and Sierra Leone. In the online Data Appendix, we discuss time use surveys in more detail and how we use these data to construct weekly hours in home production.

The four African countries on which we focus here occupy three different geographic regions: South Africa in the south, Morocco in the north, and Ghana and Sierra Leone in the west. They also represent very different stages of economic development as measured by real GDP per capita (shown in Table 1). All of the African countries we examine, except for South Africa in 2010, are poorer in the 2000s than the United States was in the 1920s. We compare the time-use patterns in home production in these countries with time-use patterns from the United States.

We focus on a sample of housewives in each country to understand the full extent of home production activities among those most responsible for these activities and to compare current female time use in Africa to historical patterns among housewives in the historical United States. Historical data on the United States from the 1920s are taken from a US Department of Agriculture report on surveys of American farmer housewives (Pidgeon 1937). Among the surveys that we use for the other countries, only Morocco contains a self-reported housewife category. To construct a comparable sample of housewives for surveys other than the US 1920s survey, we restrict samples to married women aged 15 to 59 who report not being in school or paid work in the last week and who report zero hours in market work and in education in the prior day.[7] The share of all women defined as housewives under this definition ranges from 12 to 18 percent in Sierra Leone, South Africa (2010), and the United States (2010), up to 44–47 percent in Morocco and Ghana, and 100 percent in the 1920s US data. This sample is clearly not representative of all women, but it lets us make some interesting historical

[7]Notice that the sample of US housewives in this section differs from the sample of all women used in Figure 2b. The zero hours of work or school restriction comes from the time use part of the survey and are available for all countries. The status of not being a paid worker or a student is a self-reported variable from a typical individual survey module; but these variables are not available for all countries. See online Data Appendix for details.

Table 1
Weekly Hours in Home Production among Housewives

	United States 1920s	*United States 1965*	*United States 2010*	*South Africa 2000*	*South Africa 2010*	*Morocco 2011*	*Ghana 2009*	*Sierra Leone 2003*
GDP per capita[a]	*$7,134*	*$18,130*	*$41,376*	*$5,873*	*$7,509*	*$3,621*	*$1,953*	*$641*
Panel A. Weekly hours								
Total hours	51.3	53.3	45.7	48.5	45.7	45.7	45.8	49.2
Cooking	25.1	11.5	7.0	16.5	17.0	23.6	22.7	9.8
(% of total)	(49%)	(22%)	(15%)	(34%)	(37%)	(52%)	(50%)	(20%)
Collecting firewood, water	1.5	0.0	0.0	1.9	1.1	0.5	1.7	1.9
	(3%)	(0%)	(0%)	(4%)	(2%)	(1%)	(4%)	(5%)
Cleaning	7.9	14.4	8.9	13.1	11.9	6.5	2.6	4.9
	(15%)	(27%)	(20%)	(27%)	(26%)	(14%)	(6%)	(10%)
Laundry	11.5	7.0	3.45	6.4	5.4	4.7	2.2	0.9
	(22%)	(13%)	(8%)	(13%)	(12%)	(10%)	(5%)	(2%)
Care of children, adults	3.6	10.0	15.7	8.1	7.2	7.4	9.5	18.4
	(7%)	(19%)	(34%)	(17%)	(16%)	(16%)	(21%)	(37%)
Household management	1.7	10.4	10.6	2.6	3.1	3.1	7.1	13.2
	(3%)	(20%)	(23%)	(5%)	(7%)	(7%)	(15%)	(27%)
Panel B. Sample features								
N	619	377	987	1,900	2,581	3,354	1,715	718
HH size	4.3	3.9	3.9	5.0	4.7	5.0	4.9	7.1
Housewives[b]	n/a	37%	18%	32%	16%	44%	47%	12%
Married women[c]	n/a	79%	54%	45%	38%	69%	56%	63%

Source: See online Data appendix for details of data sources.
Note: [a]GDP per capita is measured in 2011 International Dollars using the Penn World Tables v7.1. GDP per capita in the US in the 1920s is from Maddison's data reported in Herrendorf, Rogerson, and Valentinyi (2013). See text and Data appendix for definition of housewife. [b]Indicates the percentage of women in the overall sample defined as housewives; [c]indicates percentage of married women in the overall sample. 1920s US data are from a survey of farm housewives. Variable definitions: Cooking (food preparation, clean up, fetching wood and water); collecting firewood and water; cleaning (care of house, gardens); laundry (mending, laundry, making clothes); care (of children and adults in the household); household management (buying food, shopping, home management, travel for home management, other).

comparisons, it captures an important share of women in most countries, and it gives us a detailed picture of all of the activities that home production involves as well as their time intensity.

Table 1 shows average weekly hours spent on six broad categories of home production: Cooking/food preparation, collecting water and firewood, cleaning/care of household and gardens, laundry and repair of clothing, child and adult care, and general household management (which includes purchasing food, travel related to home production, and other tasks). The percent of total home production hours spent in each activity is in parentheses.

Three key facts emerge from this table, hinting at the possible barriers that African women may face in reallocating time out of home production and towards the market.[8]

Fact 1: Modern US and African housewives work very similar home hours.

African housewives do not work significantly more hours per week in the home than do American housewives. In fact, in some countries, African housewives work fewer hours. In the United States in 2010, housewives spent on average 45.7 hours per week in home production: about the same amount of time spent by Moroccan (45.7 hours), Ghanaian (45.8 hours), and (2010) South African (45.7) housewives. Only Sierra Leonean and South African housewives in 2000 report more hours in home production than American housewives in 2010.

One way to interpret the similarity in total home hours across African and US households in 2010 is that African women are combining the same total hours with less household capital and more household members, and hence are producing less total output of household services. African households certainly lack capital. Among the households in which these housewives live, only 4.3 percent (in Ghana) to 58 percent (in Morocco) have indoor plumbing. While most Moroccan households have access to electricity, only 12 percent of Sierra Leonean households do. Stove and refrigerator ownership is also low. One in 20 households in Ghana and 7 in 100 households in Sierra Leone have a stove; 1 in 5 households in Ghana and 1 in 20 households in Sierra Leone have a fridge. In contrast, by the mid-1960s, electricity and indoor plumbing were universal in the United States and over 90 percent of households owned a fridge and stove.[9]

Apart from these differences in access to infrastructure between American and African households, there is a far larger share of housewives in African households. Table 1 shows that only 18 percent of American women are housewives, compared with 16–32 percent in South Africa, and 44–47 percent in Morocco and Ghana (Sierra Leone only has 12 percent of women classified as housewives). The share of housewives in South Africa, Ghana, and Morocco in recent times is closer to the US share of housewives in the mid-1960s. This echoes the finding from Figure 2b, that weekly home time for the average African woman in our sample is closer to average home time among the average US woman in the mid-1960s.

The high share of African housewives working many hours in the home implies that there are potentially large time savings to be had among African housewives who transition out of home production and towards market work. Whether women are able to shift their hours away from home production services may depend on social norms around married women working in the market. Dividing the share of

[8]Table A.1 in the online Data Appendix shows the same breakdown of work hours in home production activities for all women (not just housewives) aged 15–59 in each country.

[9]Access to utilities and appliances are from the same African time use surveys and using data from Greenwood et al (2005), Vidart (2021), and the World Bank's World Development Indicators for the United States,

housewives by the share of married women in panel B of Table 1 for each country, we see that over 56 percent of married women in South Africa (2000), Morocco, Ghana, and Sierra Leone are housewives compared with the 46 percent of married US women who were housewives in the mid-1960s. In the final part of this paper, we discuss how social norms in Africa about married women's market work may constrain time reallocation across home and market.

Fact 2: Composition of home hours differs between modern US and African housewives.

For most African housewives, the bulk of their time is spent cooking, cleaning, and doing laundry. In South Africa, Ghana, and Morocco, cooking absorbs between one-third to just over one-half of all home production hours. Cleaning takes another 6–27 percent of home production time, while laundry takes 5–13 percent. In contrast, child- and elder-care take at most 21 percent of hours, and higher-skilled household management takes at most 15 percent of home production time.

Sierra Leone is a bit different from the other African countries. Although it is the poorest country by far, the composition of hours in home production reflects much higher loads of care work and household management, relative to other activities. One reason for this could be that households are much larger in Sierra Leone. The average household size is 7.1 in Sierra Leone, while the maximum household size in other countries is never larger than five.

The composition of home hours among modern US housewives is the exact reverse of South African, Ghanaian, and Moroccan housewives. Over half of home production hours in 2010 are spent in home management and in care work, with only 15 percent of hours spent cooking, 20 percent cleaning, and 8 percent doing the dreaded laundry. The first three columns of Table 1 demonstrate well-known historical facts about shifts in American time use throughout the 20th century. Housewife home production hours have fallen by about 11 percent percent between the 1920s and 2010—and Ramey (2009) notes even larger declines among all women. Cooking and laundry hours have shrunk dramatically, while time spent in family care activities and household management have more than quadrupled (Vanek 1973). Some activities have also disappeared: In the 1920s, American housewives spent 1.5 hours per week tending to the wood and water needs of their families, close to the 0.5–1.9 hours African housewives currently spend in these activities.

These historical shifts in the composition of home time in the United States reflect a combination of time-saving household innovations and the marketization of certain household activities that are amenable to mechanization: a structural transformation in the home. It is clear from the comparisons in Table 1 that African households have not yet experienced this kind of structural transformation in the home.

Fact 3: African housewives cook, clean, and do laundry more like US housewives of the 1920s and 1960s.

Housewives in South Africa, Ghana, Morocco spend their time much more along the lines of American women in the 1920s and the 1960s. In these three

African countries, cooking, cleaning, and laundry are all highly time-intensive home production activities.

Across all surveys (excepting Sierra Leone and 2010 United States), cooking takes up the lion's share of time spent in home production. In South Africa, Ghana, and Morocco, housewives spend between 16.5 and 23.6 hours per week cooking for their families. South African housewives spend more time cooking each week than American housewives did in the 1960s; Ghanaian and Moroccan women look more like American housewives of the 1920s, spending upwards of 20 hours per week cooking. Cleaning time in the middle- and lower-middle income African households also resembles historical US patterns. Cleaning took around 8 hours per week for US farm wives in the 1920s, while South African and Moroccan housewives spend 6.5–13.1 hours per week in house cleaning. Similarly, laundry is as time-intensive in modern South Africa and Morocco as it was between 1965 and 2010 in the United States.

The time intensity of cooking, cleaning, and doing laundry may be linked to the lack of marketized services—restaurants and grocery stories, laundromats and house cleaners—that can substitute for home production. It may also be linked to the lack of basic infrastructure and appliances; for example, without refrigerators for safe food storage, cooking must happen each day from scratch, and the lack of water and washing machines for laundry necessitates time-consuming cleaning of clothing. Investigating the availability of specific kinds of home production services and the access to basic infrastructure and appliances in African households would be a useful first step in exploring constraints to the reallocation of time from home to market. We take up these issues in the final section.

Reallocating Time: Constraints and Opportunities for Women in Africa

Structural transformation in market work has been on the rise in Africa since the late 2000s and is predicted to continue (de Vries et al. 2021). At the same time, total births per woman in sub-Saharan African fell from 5.75 to 4.62 in the two decades after 2000, while literacy among 15–24 year-old females in sub-Saharan Africa rose from 46 percent to almost 60 percent in the same period (according to World Bank Development Indicators data). These broad changes are likely to alter both the demand for home production services by households and the willingness of women to allocate time to home production relative to time in the market.

What are some of the constraints and opportunities for women reallocating unpaid time from their own homes into paid market work should they choose to do so? In this section, we discuss two broad frictions that could be important impediments to the marketization of home production as an economy undergoes structural transformation. First, we consider gender-specific social norms about appropriate market work and legal barriers restricting where and how women can spend their time. Second, we consider differences in household technologies or

care-work infrastructure (like public transportation and childcare) that could be a drag on reallocating time during any structural transformation of the economy.

Social Norms, Legal Barriers, and Female Time Allocation

Social norms about the appropriate nature of women's work, and legal barriers to women's work, restrict women's ability to work outside the home, to work in gender-mixed environments, or even to move about freely in the market place. For example, Jayachandran (2021) discusses how social norms in Islamic and Hindu societies affect what women can do in the market. A first-order effect of gender-specific social norms is on the extent to which women participate in market work at all. As one example of research in this area, Majbouri (2020) shows that female labor force participation is not sensitive to the fertility in the North African countries of Tunisia, Morocco, and Egypt, using twin births and sibling sex composition as instruments to estimate effects of having additional children. The implication is that social norms constrain female labor force participation more than high fertility.

Economists have only recently started to examine policies and test interventions that might relieve norms-based limitations on women's work. This work tends either to examine ways to boost women's formal work in the labor market within the constraints imposed by prevailing social norms, or to challenge social norms about appropriate women's work directly. Little of the prevailing evidence comes from African settings, although some studies in other parts of the world yield insights relevant for parts of Africa.

Social norms that restrict female mobility outside of the home may in part protect women from potentially dangerous public settings. For example, in many urban areas of developing countries, women are routinely harassed in public and on public transportation, and taking private taxis may be no safer. To allow women more freedom of mobility without changing the norm directly, some countries have started to experiment with providing gender-segregated transportation. In North Africa, the metro system in Cairo, Egypt, provides women-only public transit facilities; other urban areas in Latin America and in East Asia have adopted similar policies. These policies profoundly affect the experience of female mobility. In Mexico City, access to female-only subway transportation substantially reduces harassment of female commuters (Aguilar, Gutiérrez, and Villagrán 2021) during hours when this transportation was available. In Rio de Janeiro, Kondylis et al. (2020) show that one-quarter of women in their study were willing to pay a 20 percent premium to use the female-only bus transit. What remains is to show whether access to safer, female-only public transportation affects women's decisions to work in the market, at the margin.

Alternatively, although it may be difficult to change norms about how men and women work together, policy may have a role to play in incentivizing firms to pay the fixed costs of making workplaces more suitable for female workers. A recent example from Saudi Arabia is illustrative. Miller, Peck, and Seflek (forthcoming) note that three-quarters of all private sector firms in Saudi Arabia employed no women in 2009. The authors hypothesize that because firms were legally obligated

to provide separate physical and social spaces for male and female workers, the costs of integrating the labor force were too high. Yet, once firms were incentivized to hire Saudi nationals—itself a gender-neutral policy that had the unintended effect of increasing the relative demand for Saudi women—private firms paid these fixed costs and started hiring women. The share of women in the formal workforce tripled in four years (Miller, Peck, and Seflek forthcoming). Here, government policy offered a way to coordinate on a new equilibrium, opening up space for women's market work to expand without changing social norms directly.

Can social norms—and particularly men's beliefs—about acceptable women's work be changed directly? In rural India, Bernhardt et al. (2018) show that the perceived social cost of women's work falls on men, and that husbands' opposition to female labor is associated with their wives' lower take-up of employment. However, correcting perceived social costs of having a wife work outside the home can cause norms to shift quickly. In Saudi Arabia, Bursztyn et al. (2020) find that while the majority of married men in their sample support women working outside the home, these men underestimate the support of other men for this idea. The large gaps between (men's) privately held opinions and (men's) publicly accepted norms about women working outside the home generates a friction that keeps women at home. The authors experimentally and randomly correct male beliefs about what other men believe about the place of women outside of the home. They show that this belief correction increases married men's willingness to help their wives search for jobs and increases the chances that these women take up temporary jobs outside of the home. This evidence of belief updating connects to a larger literature about how information provision and exposure to role models can reduce some of the norm-based barriers to women's work (Jensen 2012; Jensen and Oster 2009).

Rather than challenging social norms about women in the market workplace, an alternative would be to challenge the social norm regarding who works in home production by seeking to change the acceptability of this work among men. Work-in-progress by McGavock et al. (2020) tests this idea. In an ongoing experiment that modifies the intervention in Bursztyn et al. (2020), the authors collect first- and second-order beliefs from married men in rural Ethiopia about the acceptability of men participating in home production activities that are usually performed by women (like collecting wood and water). Initial evidence supports the findings from Saudi Arabia: men overestimate the social stigma associated with their own participation in home production tasks. Follow-up work should provide additional insights into whether randomly correcting men's beliefs about social sanctions can change male or female time in home production.

Finally, formal legal barriers on the demand-side of the labor market also prevent women from fully participating in market work, or simply reduce women's market wages, making it less likely that women would choose to move hours of work into the market. Hyland, Djankov, and Goldberg (2020) document that women in North African and Middle Eastern countries have about half the legal rights of men. They show that gender inequality in the eyes of the law predicts worse labor

market outcomes for women: higher gender wage gaps and a lower female labor force participation rate.

In short, social norms are an important constraint on how women are able to allocate time between home and market, and likely play a large role in some parts of Africa, particularly in North Africa. A long line of research has documented changes in social norms about women's work in historical developed countries and in some developing countries (a non-exhaustive list of examples might begin with Fernández and Fogli 2009; Doepke, Tertilt, and Voena 2012; Alesina, Gutiérrez, and Soto Villagrán 2013). In the US historical experience, the social norms around and stigma attached to women working in dirty manufacturing jobs shifted dramatically after the 1950s and the rise of cleaner service sector jobs (Goldin 1995). Along with policy steps discussed in this section, there is hope that social norms regarding women's work outside of home production may respond to changing economic conditions through the structural transformation in Africa.

Physical and Care-Work Infrastructure and Female Time Allocation

The constraints that women face in allocating time between home and market may differ when female labor force participation is already high. As discussed earlier, this situation is widespread in sub-Saharan Africa, where most women combine low market hours and heavy time demands from home production. In this section, we address two questions: Could technological changes in home production (including easier access to clean water and fuel) save time worked in the home, allowing more hours to shift towards the market? If work shifts away from family farms towards the market place, how can countries prevent an exit from the labor force for those women who would like to continue working but have family responsibilities?

Economic theory is ambiguous on whether or how providing basic utilities to households will affect time spent in home or in market work: it depends on the relative productivity of labor in the home and market work and on the demand (and the income elasticity) for home production services. In developed countries, there has been an ongoing debate about whether access to durables that improved productivity of labor in the home increased women's ability to do market work (as in Greenwood, Seshadri, and Yorukoglu 2005), or led to increased demands for home production done by women (as in Reid 1934; Vanek 1973; Mokyr 2000).

The more recent empirical evidence is also mixed. Randomized trials in urban Morocco (Devoto et al. 2012) and rural Kenya (Kremer et al. 2011) show that improved access to clean water reduced time spent in home work, but none of the released time translated into more female market work in either country. In contrast, Meeks (2017) shows that new water infrastructure in rural Kyrgyzstan substantially reduces time spent in home production and shifts this time into the market and leisure. Dinkelman (2011) shows that grid electrification in South Africa changed the nature of home production in rural areas, with newly electrified homes relying more on electrical appliances for lighting, cooking, and heating. New access to electrification also increased net female employment. In that paper, evidence pointed

towards a net increase in labor supply, driving increased female market work. Yet in other settings like Kenya and India, researchers have found modest positive to zero impacts of electricity on female market work (for a review of recent work on effects of household electrification in this journal, see Lee, Miguel, and Wolfram 2020).

Electricity might directly increase the demand for female labor over generations by changing the types of market work women can do and incentivizing educational attainment. Vidart (2020) shows that US electrification in the late 1800s and early 1900s raised the return to hiring skilled workers and drew women out of the home. As these new jobs raised the returns to education, women responded by investing in more schooling. This link between infrastructure and investment in human capital, and then to employment in higher-skilled sectors—part of the structural transformation—is an under-explored area in African settings.

In the African context, releasing female time into the market may require simultaneously lifting a number of home production constraints. For example, access to efficient cookstoves alone may not reduce female time in home production if there is no way to store prepared food (refrigeration), or if someone still needs to take care of children. This point naturally leads to a consideration of an important complementary constraint on how women can use their time: the need for safe, quality childcare or eldercare.

The lack of access to marketized home production services, in particular childcare, likely constrains market hours and type of work chosen by women. Without childcare infrastructure, working when children are young requires market work that is amenable to interruption, has extreme flexibility, and is easy to access from a home base. Ironically, these characteristics almost perfectly describe jobs on a family farm and self-employment in home-based businesses in the informal sector—the kinds of jobs that women in Africa are *already* likely to occupy.

Does a lack of childcare partly account for the high prevalence of self-employment among women in developing countries? Does it also help us understand why there is so much variation in how women respond to changes in infrastructure and electricity? Heath (2017) highlights that the answer may depend on broader family structure. Using panel data from urban Ghana, she shows that with the arrival of a child, the women who drop out of the market workforce tend to have no other family members who can help with childcare, while those who remain and work more hours rely on older female relatives or elder children for childcare. Her paper implies that childcare constraints influence time allocation among a cross-generational array of women: mothers and their sisters, grandmothers, and older daughters.

Although self-employment offers great flexibility in hours or work, it does not always offer great flexibility in location. Delecourt and Fitzpatrick (2021) show that while almost 40 percent of self-employed females in Uganda bring their babies to work, no self-employed men report doing so. Women who bring their infants to work earn lower profits than men and then childless women. Their evidence suggests that women with young children are more likely to be stocked out of products that

require them to travel for replacement—hence, the lack of childcare constraints affects investment decisions at the level of the small firm.

Might childcare subsidies help to promote women's market work in developing countries? In urban Kenya where a private market for childcare exists, Clark et al. (2019) use a randomized controlled trial to show that subsidizing access to private childcare raises female employment by 8.5 percentage points. The subsidy both raised the labor force participation of married women and nudged unmarried mothers to take up more regular and inflexible hours jobs. Ongoing work in rural Burkina Faso by the World Bank (2021) tests whether providing mobile childcare facilities (tents and childcare workers) for women in public works programs affects labor market participation of women. In this case, because the work itself—digging ditches and clearing public lands—moves around, so does the childcare. More evidence would be extremely useful in shedding light on how households cope with the arrival of young children, and which women are most constrained on the intensive margin of market work by the requirements of childcare.

Simon Kuznets (1966, p. 157) noted that in order for economies to grow, they must undergo a process of structural transformation: "A high rate of modern economic growth is attainable only if the required marked shifts in industrial structure are not too impeded by resistance—of labour and of capital, of people and their resources in the old and accustomed grooves." Understanding the factors that keep African women's time in the "old and accustomed grooves" of home production or unpaid family work is an exciting new area for research, which would benefit from greater application of time use surveys on the continent.

■ *This work was partly supported by the Structural Transformation and Economic Growth (STEG) program at CEPR and by the Kellogg Institute for International Studies at the University of Notre Dame. Research assistance was provided by Regina Mannino, Brigid Meisenbacher, and Vivek Moorthy. We thank Gaaitzen de Vries for assistance with data from the GGDC Africa Sector Database.*

References

Addati, Laura, Umberto Cattaneo, Vale Esquivel, and Isabel Valarino. 2018. *Care Work and Care Jobs: For the Future of Decent Work*. Geneva: International Labour Organization.

Aguiar, Mark, and Erik Hurst. 2007. "Measuring Trends in Leisure: The Allocation of Time over Five Decades." *Quarterly Journal of Economics* 122 (3): 969–1006.

Aguilar, Arturo, Emilio Gutiérrez, and Paula Soto Villagrán. 2021. "Benefits and Unintended Consequences of Gender Segregation in Public Transportation: Evidence from Mexico City's Subway

System." *Economic Development and Cultural Change*. 69 (4).

Alesina, Alberto, Paola Giuliano, and Nathan Nunn. 2013. "On the Origins of Gender Roles: Women and the Plough." *Quarterly Journal of Economics* 128 (2): 469–530.

Bernhardt, Arielle, Erica Field, Rohini Pande, Natalia Rigol, Simone Schaner, and Charity Troyer-Moore. 2018. "Male Social Status and Women's Work." *AEA Papers and Proceedings* 108: 363–67.

Boserup, Ester. 1970. *Woman's Role in Economic Development*. London: Earthscan.

Bridgman, Benjamin, Georg Duernecker, and Berthold Herrendorf. 2018. "Structural Transformation, Marketization, and Household Production around the World." *Journal of Development Economics* 133: 102–26.

Burda, Michael, Daniel S. Hamermesh, and Philippe Weil. 2013. "Total Work and Gender: Facts and Possible Explanations." *Journal of Population Economics* 26: 239–61.

Bursztyn, Leonardo, Alessandra L. Gonz´alez, and David Yanagizawa-Drott. 2020. "Misperceived Social Norms: Women Working Outside the Home in Saudi Arabia." *American Economic Review* 110 (10): 2997–3029.

Clark, Shelley, Caroline W. Kabiru, Sonia Laszlo, and Stella Muthuri. 2019. "The Impact of Childcare on Poor Urban Women's Economic Empowerment in Africa." *Demography* 56 (4): 1247–72.

de Vries, Gaaitzen, Linda Arfelt, Dorothea Drees, Mareike Godemann, Calumn Hamilton, Bente Jessen-Thiesen, Ahmet Ihsan Kaya, Hagen Kruse, Emmanuel Mensah, and Pieter Woltjer. 2021. The Economic Transformation Database (ETD): content, sources, and methods. WIDER Technical Note 2/2021. https://doi.org/10.35188/UNU-WIDER/WTN/2021-2.

de Vries, Gaaitzen, Marcel Timmer, and Klaas de Vries. 2015. "Structural Transformation in Africa: Static Gains, Dynamic Losses." *Journal of Development Studies* 51 (6).

Delecourt, Solène, and Anne Fitzpatrick. 2021. "Childcare Matters: Female Business Owners and the Baby-Profit Gap." *Management Science* 67 (7): 3985–4642.

Devoto, Florencia, Esther Duflo, Pascaline Dupas, William Parienté, and Vincent Pons. 2012. "Happiness on Tap: Piped Water Adoption in Urban Morocco." *American Economic Journal: Economic Policy* 4 (4): 68–99.

Dinkelman, Taryn. 2011. "The Effects of Rural Electrification on Employment: New Evidence from South Africa." *American Economic Review* 101 (7): 3078–3108.

Doepke, Matthias, Michèle Tertilt, and Alessandra Voena. 2012. "The Economics and Politics of Women's Rights." *Annual Review of Economics* 4: 339–72.

Durand, J. 1975. *The Labor Force in Economic Development: A Comparison of International Census Data, 1946–66*. Princeton: Princeton University Press.

Fernández, Raquel, and Alessandra Fogli. 2009. "Culture: An Empirical Investigation of Beliefs, Work, and Fertility." *American Economic Journal: Macroeconomics* 1 (1): 146–77.

Freeman, Richard B., and Ronald Schettkat. 2005. "Marketization of Household Production and the EU–US Gap in Work." *Economic Policy* 20 (41): 6–50.

Gaddis, Isis, Gbemisola Oseni, Amparo Palacios-Lopez, and Janneke Pieters. 2020. "Who Is Employed? Evidence from Sub-Saharan Africa on Redefining Employment." World Bank Policy Research Working Paper 9370.

Goldin, Claudia. 1986. "The Female Labor Force and American Economic Growth." In *Long-term Factors in American Economic Growth*, Vol. 1, edited by Stanley L. Engerman and Robert E. Gallman, 557–604. Chicago: University of Chicago Press.

Goldin, Claudia. 1995. "The U-Shaped Female Labor Force Function in Economic Development and Economic History." In *Investment in Women's Human Capital*, edited by T. Paul Schultz, 61–90. Chicago: University of Chicago Press.

Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu. 2005. "Engines of Liberation." *Review of Economic Studies* 72 (1): 109–133.

Heath, Rachel. 2017. "Fertility at Work: Children and Women's Labor Market Outcomes in Urban Ghana." *Journal of Development Economics* 126: 190–214.

Heath, R., and S. Jayachandran. 2018. "The Causes and Consequences of Increased Female Education and Labor Force Participation in Developing Countries." In *Oxford Handbook on the Economics of Women*, edited by Susan L. Averett, Laura M. Argys, and Saul D. Hoffman. New York: Oxford University Press.

Herrendorf, B., R. Rogerson, and Ákos Valentinyi. 2013. "Growth and Structural Transformation." In *Handbook of Economic Growth*, edited by Philippe Aghion and Steven N. Durlauf, 855–941. Amsterdam: Elsevier.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow. 2019. "The Allocation of Talent and U.S. Economic Growth." *Econometrica* 87 (5): 1439–74.

Hyland, Marie, Simeon Djankov, and Pinelopi Koujianou Goldberg. 2020. "Gendered Laws and Women in the Workforce." *American Economic Review: Insights* 2 (4): 475–90.

Jayachandran, Seema. 2021. "Social Norms as a Barrier to Women's Employment in Developing Countries." *IMF Review* 69 (3): 576–95.

Jensen, Robert. 2012. "Do Labor Market Opportunities Affect Young Women's Work and Family Decisions? Experimental Evidence from India." *Quarterly Journal of Economics* 127 (2): 753–92.

Jensen, Robert, and Emily Oster. 2009. "The Power of TV: Cable Television and Women's Status in India." *Quarterly Journal of Economics* 124 (3): 1057–94.

Kondylis, Florence, Arianna Legovini, Kate Vyborny, Astrid Maria Theresia Zwager, and Luiza Cardoso de Andrade. 2020. "Demand for Safe Spaces: Avoiding Harassment and Stigma." World Bank Policy Research Working Paper 9269.

Kremer, Michael, Jessica Leino, Edward Miguel, and Alix Peterson Zwane. 2011. "Spring Cleaning: Rural Water Impacts, Valuation, and Property Rights Institutions." *Quarterly Journal of Economics* 126 (1): 145–205.

Kuznets, Simon. 1966. *Modern Economic Growth, Rate, Structure and Spread*. New Haven: Yale University Press.

Lebergott, Stanley. 1993. *Pursuing Happiness*. Princeton, NJ: Princeton University Press.

Lee, Kenneth, Edward Miguel, and Catherine Wolfram. 2020. "Does Household Electrification Supercharge Economic Development?" *Journal of Economic Perspectives* 34 (1): 122–44.

Lee, Jay Euijung. 2020. "Marriage and Misallocation: Evidence from 70 Years of U.S. History." PhD diss. London School of Economics.

Maddison, A. 1980. "Economic Growth and Structural Change in the Advanced Countries." In *Western Economies in Transition: Structural Change and Adjustment Policies in Industrial Countries*, edited by Irving Leveson and Jimmy W. Wheeler, 41–65. London: Croom Helm.

Majbouri, Mahdi. 2020. "Fertility and the Puzzle of Female Employment in the Middle East and North Africa." *Economics of Transition and Institutional Change* 28 (2): 225–44.

Mammen, Kristin, and Christina Paxson. 2000. "Women's Work and Economic Development." *Journal of Economic Perspectives* 14 (4): 141–64.

McGavock, T., E. McCullough, N.P. Magnan, and T. Assefa. 2020. "Stifled by Stigma? Experimental Effects of Updating Husbands' Beliefs on Participation in Women's Household Work." https://drive.google.com/file/d/10Rhv_GbMTthsTUYeGFJkKBtnxJ7_aaM2/view.

Meeks, Robyn. 2017. "Water Works: The Economic Impact of Water Infrastructure." *Journal of Human Resources* 52 (4): 1119–53.

Miller, Conrad, Jennifer Peck, and Mehmet Seflek. Forthcoming. "Missing Women, Integration Costs, and Big Push Policies in the Saudi Labor Market." *American Economic Journal: Applied Economics*.

Mincer, Jacob. 1962. "Labor Force Participation of Married Women: A Study of Labor Supply." In *National Bureau of Economic Research, Aspects of Labor Economics*, 63–106. Princeton NJ: NBER.

Mokyr, Joel. 2000. "Why 'More Work for Mother?' Knowledge and the Household Behavior, 1870–1945." *Journal of Economic History* 60 (1): 1–41.

Ngai, L.R., C. Olivetti, and B. Petrongolo. 2021. "Structural Transformation and the U-shaped Female Employment."

Ngai, L. Rachel, and Barbara Petrongolo. 2017. "Gender Gaps and the Rise of the Service Economy." *American Economic Journal: Macroeconomics* 9 (4): 1–44.

Ngai, L. Rachel, and Christopher A. Pissarides. 2008. "Trends in Hours and Economic Growth." *Review of Economic Dynamics* 11 (2): 239–56.

Ngai, L. Rachel, and Christopher A. Pissarides. 2011. "Taxes, Social Subsidies, and Allocation of Work Time." *American Economic Journal: Macroeconomics* 3 (4): 1–26.

Pidgeon, Mary E. 1937. "Women in the Economy of the United States of America: A Summary Report." United States Department of Labor Technical Report 155. Washington, DC: US Government Printing Office.

Reid, Margaret G. 1934. *Economics of Household Production*. New York: J. Wiley Sons, Incorporated.

Rendall, Michelle. 2018. "Female Market Work, Tax Regimes, and the Rise of the Service Sector." *Review of Economic Dynamics* 28: 269–89.

Rodrik, Dani. 2016. "Premature Deindustrialization." *Journal of Economic Growth* 21: 1–33.

Rogerson, Richard. 2008. "Structural Transformation and the Deterioration of European Labor Market Outcomes." *Journal of Political Economy* 116 (2): 235–59.

Ruggles, Steven. 2015. "Patriarchy, Power, and Pay: The Transformation of American Families, 1800–2015." *Demography* 52 (6): 1797–1823.

Sinha, J.N. 1965. *Dynamics of Female Participation in Economic Activity in a Developing Economy.* New York: UN Publications.

Vanek, J. 1973. *Keeping Busy: Time Spent in Housework, United States, 1920–1970.* PhD diss. University of Michigan.

Vidart, D. 2020. "Human Capital, Female Employment, and Electricity: Evidence from the Early 20th Century United States."

World Bank. 2021. *Enabling Women to Work and Their Children to Blossom: The Double Success Story of Mobile Childcare in Burkina Faso.* World Bank Technical Report.

Journal of Economic Perspectives—Volume 36, Number 1—Winter 2022—Pages 81–100

Young Adults and Labor Markets in Africa

Oriana Bandiera, Ahmed Elsayed, Andrea Smurra, and Céline Zipfel

Today, one of every five people who start looking for their first job is born in Africa. By 2050, it will be one in three, as illustrated in Figure 1. The large share of young Africans seeking employment is the outcome of different fertility trends across the world: fertility has been falling everywhere, but much more slowly in Africa (Lam et al. 2019). It is no exaggeration to say that the economic and political future of the continent, as well as the outcome of the fight against global poverty, depend on the job opportunities available to the children born in Africa today.

This paper presents and discusses facts about African labor markets, focusing on young labor market entrants in Africa and how they fare compared to older generations and to their counterparts in other low- and middle-income regions of the world.

To study macro patterns with micro data, we construct a new dataset of labor market indicators that can be aggregated at the geography, age, wealth and gender level. Our starting point is the Jobs of the World Project (Bandiera, Elsayed, and Smurra 2022), which harmonizes and collates the universe of Demographic and

■ *Oriana Bandiera is Sir Anthony Atkinson Professor of Economics, London School of Economics, London, United Kingdom. Ahmed Elsayed is a Senior Research Associate, IZA (Institute of Labor Economics), Bonn, Germany. Andrea Smurra is a PhD candidate in Economics, University College London, London, United Kingdom. Céline Zipfel is Assistant Professor of Economics, Stockholm School of Economics, Stockholm, Sweden. Their email addresses are o.bandiera@lse.ac.uk, elsayed@iza.org, a.smurra.11@ucl.ac.uk and celine.zipfel@hhs.se.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.36.1.81.

Figure 1
Share of World Population Age 15–24, by Region

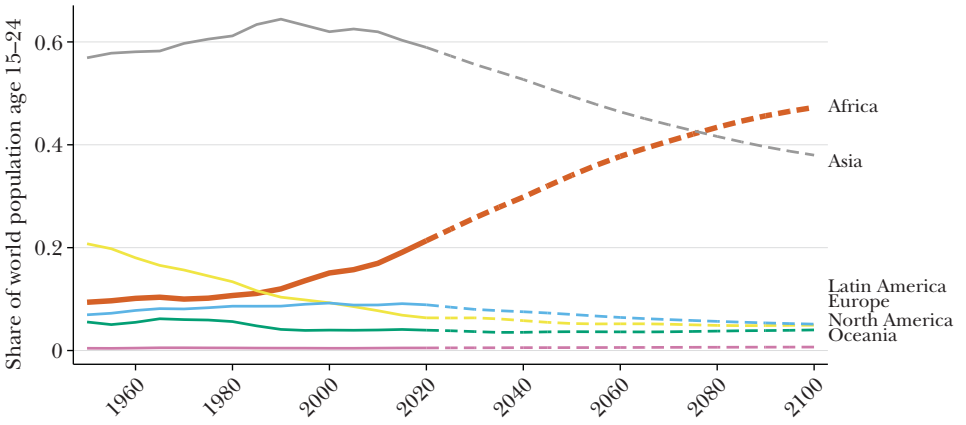

Source: UN World Population Prospects (2019).
Note: Regional shares of the world's total 15–24 year-old population, constructed from population estimates and projections from the UN World Population Prospects, 2019 (https://population.un.org/wpp/). The dashed parts of the curves plot projections. Regions correspond to the geographical units coded in the source data. The figure refers to the cohort aged 15–24 as population estimates and projections are provided for five year bands. In the remainder of the paper, where we examine aggregates produced from microdata, we will focus on the 18–24 age group. An analogous version of this graph, with absolute population sizes, is provided in Figure A1 in the online Appendix.

Health Surveys (ICF 1990–2017) and IPUMS censuses since 1990 (Minnesota Population Center 2020) in the public domain. We focus on countries for which we have sufficiently detailed information on respondents' occupations and keep only the latest wave available for each country. We study the employment patterns of the "young adults" group, defined as individuals aged 18-24. Our final sample covers 68 low- and middle-income countries: 28 in Africa and 40 in other regions. These comprise a total population of 593 million people in Africa and 1.18 billion people in the rest of the world. Our set of comparison countries is drawn from all other regions except North America, with South Asia the least represented (one country, Nepal, out of six) and Latin America and the Caribbean the most (22 out of 25 countries). Our sample covers a wide range of national incomes: real GDP per capita ranges from $1,051 (Liberia 2008) to $18,051 (Mauritius 2011) in the set of African countries and from $1,656 (Haiti 2003) to $19,285 (Romania 2011) in the comparison group of other countries.

The key pattern that emerges from this dataset is that the development process entails a transformation in the organization of labor: a myriad of micro-entrepreneurs turn into salaried workers hired by large firms. We show that this pattern starts with the youngest cohorts everywhere but in Africa, where young workers are as likely as their parents' generation to hold a salaried job. We then discuss possible explanations and policies to address this issue.

Jobs and Development

To begin, we document the evolution of jobs along the path of development and ask whether African countries are "on track" given their income levels—that is, if the organization of labor in African countries follows the same path as that of other countries conditional on income.

Figure 2 reports the share of young adults in different employment categories for 28 countries in Africa and our comparison group of 40 other low- and middle-income countries. For each group, the top panel of the figure shows the share of 18–24 year-old individuals who do any work whether paid or unpaid, with the remainder (students, unemployed workers, and homemakers) categorized as not working.[1] The lower panel then shows the composition of young workers' employment, according to their main sector of occupation (agriculture or manufacturing/services) and type of employment (self-employed or salaried). We classify individuals working for themselves or for family members as self-employed, and those who work for someone else for pay as salaried workers.

Young Adults in Africa Are Equally Likely to Work, but Less Likely to Get Paid

As the top panel of the figure makes clear, the 18–24 year-old population is almost equally split between those at work and those not working in both sets of countries. The main difference between Africa and the rest of the sample is the incidence of unpaid work. The share of those in unpaid work is 15.6 percent, over twice as large as in the group of comparison countries. The share of those in paid work is 32.4 percent in Africa compared to 41.5 percent in the comparison group. Disaggregating these results by gender largely mirrors our findings from the pooled sample.

Young Adults in Africa Are Less Likely to Have an Employer

In the lower bars in Figure 2, we divide workers (both paid and unpaid) into four groups: self-employed and salaried in both sectors of occupation. Here, we classify unpaid workers as self-employed, since these are typically individuals participating in household farm work or in the family business without monetary remuneration. Paid workers, on the other hand, are made up of employees earning regular wages and self-employed individuals who report taking home some monetary revenue from their activity.

We find a striking difference in the organization of labor across the two sets of countries: only 25.5 percent of labor market entrants in Africa have a salaried job, compared to 68.8 percent in our comparison group. This pattern also holds within sectors: the share of salaried jobs is over twice as large outside of Africa, both in agriculture (25.6 percent of young workers in agriculture have salaried

[1] Our characterization of workforce participation does not distinguish between unemployed individuals and those out of the labor force. In particular, the Demographic and Health Surveys This distinction is crucial to the definition of unemployment used by the International Labour Organization.

Figure 2
Occupational Structure of the 18–24 Year-Old Population

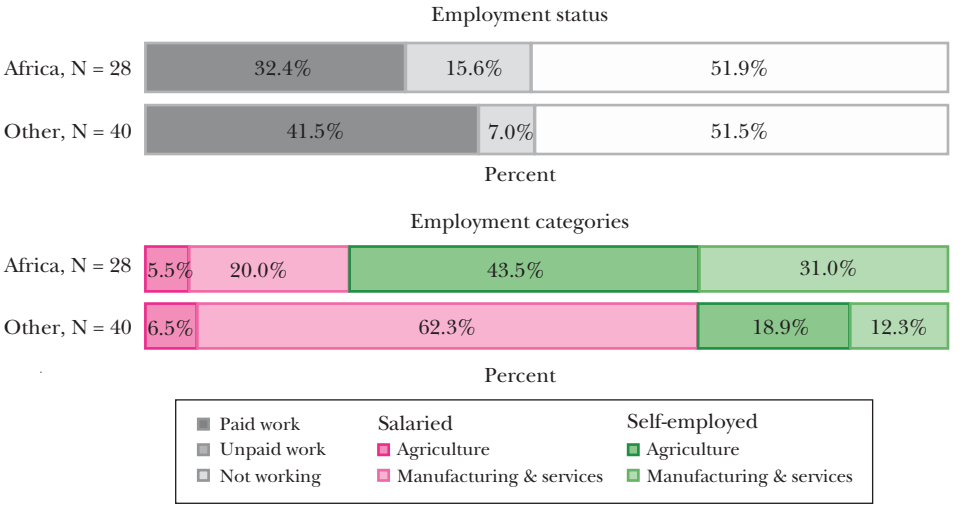

Source: Demographic and Health Surveys and IPUMS, harmonized via the Jobs of the World Project.
Note: Regional aggregates for the 18–24 year-old population in 68 low-income countries (28 countries from Africa and 40 countries from the rest of the world) constructed from the latest sample available for each country in the set of Demographic and Health Surveys and IPUMS censuses that contain the relevant labor outcomes for our exercise. The top panel plots the relative shares of three "extensive margin" categories: fraction of individuals aged 18–24 (i) working for pay, (ii) in unpaid work, and (iii) not working. The bottom panel plots the relative shares of four employment categories (defined according to sector and type of work), restricting the sample to working individuals (paid and unpaid). Regional averages are computed using countries' population size as weights; for the unweighted version, see Figure A2 in the online Appendix. For figures that disaggregate these results by gender, see Figure A3 in the online Appendix.

jobs compared to 11.2 percent in Africa) and in manufacturing and services (83.5 percent of young workers in these sectors have salaried jobs compared to 39.2 percent in Africa).

It will come to no surprise that agriculture employs a larger share of the population in Africa, where 49 percent of 18–24 year-old workers cite agriculture as their primary sector of occupation compared to only 25.4 percent in the comparison group of other low-income countries. Indeed, merging our sample of country-year observations with the World Development Indicators database reveals that in Africa, agriculture accounts for 20.2 percent of GDP on average—almost twice the average share in our comparison group (10.4 percent) (World Bank 2021).

If we split non-agricultural employment into manufacturing and services, we find that 16.9 percent of 18–24 year-old workers are employed in manufacturing in Africa compared to 25.6 percent in the low-income countries outside of Africa, and 31.8 percent of workers in this age group are employed in services in Africa compared to 48.7 percent in the comparison group.

Our database also allows us to break down employment patterns by occupational skill groups for 63 out of the 68 countries in our sample. Here, we can distinguish between high-skilled professionals, white-collar workers (for example, clerks) and blue-collar workers. We find that part of the higher share of salaried jobs in the employment structure of the comparison group of other low-income countries is also attributable to a higher prevalence of high-skilled jobs, of which the share is over twice as large in those countries as in Africa (12.9 percent versus 5.9 percent), and, to a lesser extent, white-collar jobs, which make up 25.8 percent of the workforce compared to 23.3 percent in Africa.[2]

Young Adults in Richer Countries Are More Likely to Have a Salaried Job

Are job differences due to differences in the level of development between African and other countries? To investigate this question, Figure 3 plots the shares of young workers in the four employment categories against GDP per capita. On the *y*-axis are country averages of these employment indicators for the 18–24 year-old population for each country in our dataset. The *x*-axis plots the natural logarithm of real GDP per capita (in 2017 US dollars).

Two patterns emerge from Figure 3. First, the process of development entails a shift in the organization of labor from self-employment to salaried jobs. This shift is more than just the structural change from agriculture to manufacturing and services: as the figure shows, higher levels of GDP per capita are associated with a fall in self-employment and rise in salaried employment within both agriculture and manufacturing/services. Perhaps the most striking change observed along the development path is a rapid rise in salaried jobs in manufacturing and services (shown in the second panel).

These findings resonate with those in Jensen (2021), who shows that the employee share increases as GDP grows. This is also in line with the view that economic growth stems from the expansion of the formal sector (Schoar 2010; La Porta and Shleifer 2008, 2014).

The fact that self-employment is the dominant activity in low-income countries implies a different interpretation of unemployment and labor force participation. Feng, Lagakos, and Rauch (2021) document that unemployment is largely a feature of advanced economies, while Donovan, Lu, and Schoellman (2021) show that a large share of workers in poorer countries frequently transition between self-employment and "non-employment."

Second, Figure 3 suggests that the differences presented in Figure 2 are largely in line with GDP differences between Africa and the rest of the world. The African countries with higher per capita income levels, such as Botswana, have a similar structure as non-African countries at similar levels of GDP per capita: that is, around 80 percent of 18–24 year-old workers in salaried employment in manufacturing and

[2]Figure A4 in the online Appendix illustrates the pattern described in the text from dividing non-agricultural employment into manufacturing and services. Figure A5 in the Appendix illustrates the breakdown by occupational skill group.

Figure 3
Share of Youth Employment Categories by GDP per Capita

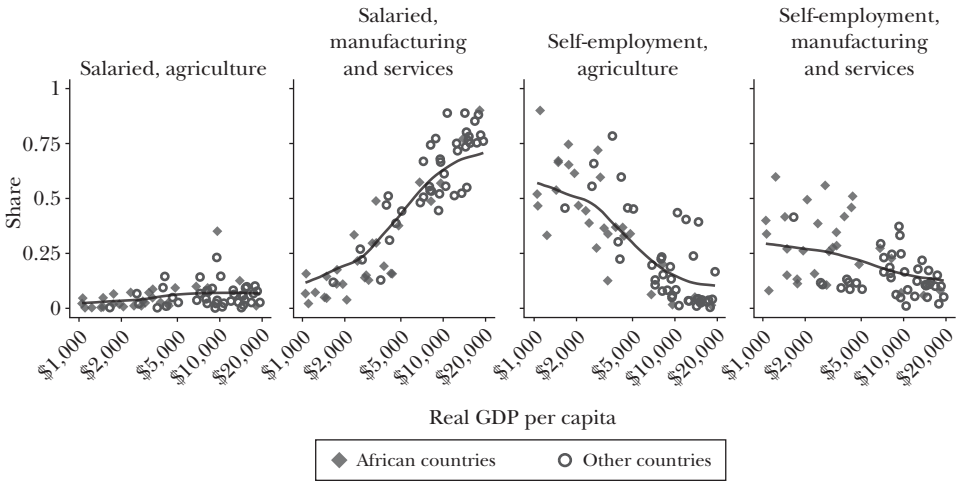

Source: Demographic and Health Surveys and IPUMS, harmonized via the Jobs of the World Project.
Note: Country-level aggregates for the 18–24 year-old population in 68 low-income countries (28 countries from Africa and 40 countries from the rest of the world) constructed from the latest sample available for each country in our dataset. Samples are selected if they contain the following information for both men and women: whether the individual engages in any economic activity at the time of the survey or census, type of employer (self/family/someone else), paid/unpaid work, and sector of occupation. The *x*-axis displays real GDP per capita at constant 2017 national prices (millions of 2017 US dollars) in logarithmic scale.

services and less than 5 percent of workers in this age group in self-employment in agriculture. Likewise, the employment structure of the lowest-income countries outside of Africa is similar to that of African countries at similar levels of development. For example, 45 percent and 41 percent of 18–24 year-old workers in Haiti are engaged in self-employment in agriculture and manufacturing and services respectively, with only 12 percent in salaried jobs in total.

The country-level data indicate that Africa is on average on the same development track as other regions. African countries are much poorer, but according to the patterns presented in Figure 3, their employment structures seem likely to undergo, eventually, a similar shift towards salaried jobs with future economic growth. Broken down by cohort, however, the data tells a different story.

Salaried Jobs Are More Common among the Young, but Not in Africa

Figure 4 plots the mean shares of each of the same four employment categories as in Figures 2 and 3 by age cohort, separately for African countries and our comparison group of other low-income countries. In the other low-income countries, the share of population in agriculture is essentially flat across age cohorts, while participation in manufacturing and services varies sharply with age. The second and fourth

Figure 4
Share of Employment Categories by Age Cohort

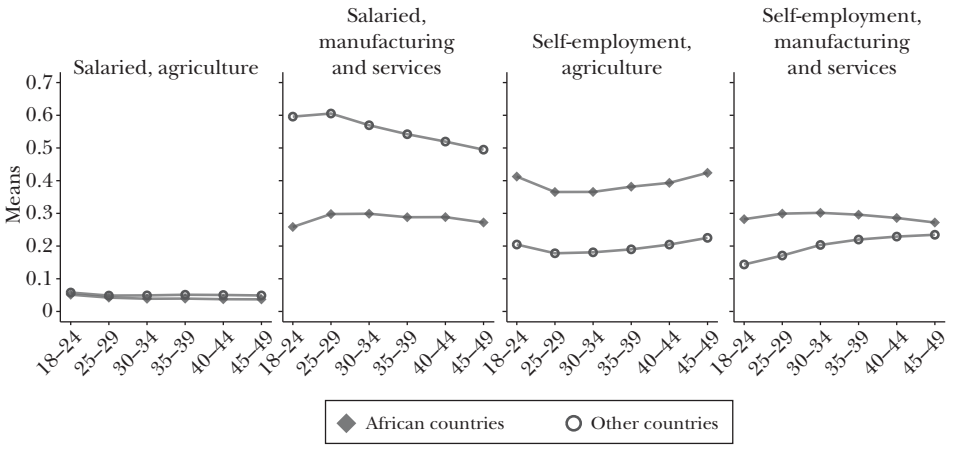

Source: Demographic and Health Surveys and IPUMS, harmonized via the Jobs of the World Project.
Note: Country-year-age cohort aggregates for 68 low-income countries (28 countries from Africa and 40 countries from the rest of the world) constructed from the latest sample available for each country in our dataset. Samples are selected if they contain the following information for both men and women: labor force participation, type of employer (self/family/someone else), paid/unpaid work, and sector of occupation. Each aggregate is constructed as a share of labor market participants within each age cohort.

panels show that the age gradient is positive for self-employment and negative for salaried jobs in these sectors: older people are more likely to be self-employed and less likely to have salaried jobs in these countries. This is likely due to positive trends in economic growth bringing about new salaried job opportunities for the younger cohorts, as suggested by the positive association between GDP per capita and salaried employment rates depicted in Figure 3.

In Africa, this pattern seems largely absent. While Africa's age-occupation profiles mirror those of the other countries for agriculture, they are essentially flat for manufacturing and services in the second and fourth panels. In particular, the trend presented in the second panel indicates that labor market entrants are no more likely to hold a salaried job than their older counterparts. This pattern seems particularly puzzling in light of existing evidence that many African economies seem to have undergone some structural change in recent decades (Diao, Harttgen, and McMillan 2017).[3]

The key takeaway of Figure 4 is thus that the jobs of many young people in Africa do not differ from that of their parents' generation. This casts doubt on the prediction that youth employment in the region will naturally follow the same path

[3] This study documents a shift out of agriculture by rural females over the age of 25 who have a primary education since the beginning of the 21st century in sub-Saharan Africa.

of transformation associated with economic growth that has been observed in other countries.

A key difference between Africa and other regions is that the number of young entrants to the labor market is much higher in Africa. Might this fact be responsible for the patterns we see or is it a consequence of the lack of salaried jobs? In principle, the relationship between cohort size and employment structure can go both ways. Indeed, the number of people in the same cohort determines the number of people competing for the same jobs. At the same time, the absence of salaried jobs lowers the cost of having children, likely keeping fertility high (Zipfel 2021). If both forces are at play, they will feed off each other keeping African youths trapped in a cycle of bad jobs and high fertility. We investigate this scenario next.

Cohort Size and the Competition for Salaried Jobs

A major difference between Africa and the comparison group of low- and middle-income countries in our database is the large size of the young cohort entering the labor market for the first time. This is a result of the region's slow demographic transition: since its onset, the fertility decline has been less steep in Africa relative to other developing regions, as recognized by demographers (Bongaarts and Casterline 2012; Bongaarts 2017).

We calculated the share of young people aged 15–24 out of the population aged 15–49 using data from the UN World Population Prospects. For the world as a whole, in 2020, the average share of the 15–24 age group is 31 percent. For the countries of Africa, the share is 40 percent. In our dataset, when we list the countries where the share of the 15–24 age group as part of the 15–49 age group is highest, seven of the top ten countries are in Africa: Eswatini (50 percent, 2006), Lesotho (46 percent, 2009), Honduras (45 percent, 2001), Zambia (44 percent, 2010), Cameroon (44 percent, 2005), Burkina Faso (44 percent, 1996), Mali (44 percent, 2009), Nicaragua (43 percent, 2005), Guinea (43 percent, 2014), and Guatemala (43 percent, 2014). Indeed, 21 of the top 30 countries in this ranking are in Africa.

Does the probability of finding a job depend on the size of one's age cohort? The answer depends on the several factors that determine the pool of workers who compete for the same job, that is, whether workers of different ages are substitutes. For low-skilled jobs with low returns to experience, the relevant pool is not tied to age; if experience matters, it is.

To shed light on the link between cohort size and the jobs of young adults we use variation in the share of youths across provinces—the smallest geographical level at which Demographic and Health Surveys surveys are representative—within African countries. This sample covers a total of 345 provinces across 28 countries. The number of provinces per country ranges from three (Malawi) to 34 (Guinea), with the median at ten. The share of the 18–24 age group across these regions varies between 21.6 percent and 43.3 percent.

Figure 5
Youth Employment Categories by Cohort Size

Source: Demographic and Health Surveys and IPUMS, harmonized via the Jobs of the World Project.
Note: Scatterplot of 345 province-level aggregates for 28 African countries constructed from the latest sample available for each country in our dataset. The *x*-axis plots the share of young people aged 18–24 relative to the overall sample aged 18–49 in a province, as a measure of cohort size. Samples are selected if they contain the following information for both men and women: labor force participation, type of employer (self, family, or someone else), paid/unpaid work, and sector of occupation. Each outcome plotted on the *y*-axis is constructed as a share of 18–24 year-old workers in each province. Lines plot the smoothed values of local polynomial regression estimates with an Epanechnikov kernel of optimal bandwidth. Details of the calculations are available in the online Appendix.

Figure 5 plots the province-level averages of each employment indicator against the share of the working-age province population who are aged 18–24 for all African provinces in our sample, as well as the fitted values from local polynomial regressions. It shows that the youth bulge is associated with a decrease in the share of salaried employment outside of agriculture among young workers (as shown in the second panel) and an increase in self-employment in both sectors (as shown in the third and fourth panels). That is, individuals born in large cohorts are less likely to find a salaried job when they join the labor force.

Figure 6 replicates the analysis across wealth quintiles.[4] In the left-hand graph, the share of workers in salaried agriculture jobs is small everywhere, but the poorest are particularly more likely to engage in this type of employment. The probability is relatively higher for regions with small cohorts and converges to a level very close to that of other quintiles as the young cohorts' share of the working-age population

[4]A major advantage of our dataset is that it uses detailed data about household assets and dwelling characteristics to construct household-level wealth proxies. We use these to create wealth quintiles (which are sufficiently large bins to ensure distinct socioeconomic differentiation) for each country-year sample for which the relevant variables exist. A large proportion of households in low-income countries are intergenerational, which explains why we observe young adults in each quintile.

Figure 6
Youth Employment Categories by Cohort Size and Wealth Quintile

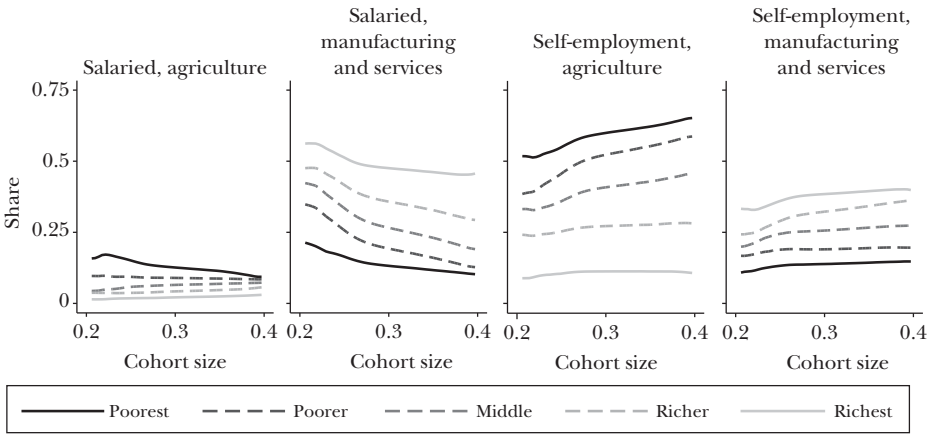

Source: Demographic and Health Surveys and IPUMS, harmonized via the Jobs of the World Project.
Note: Each graph plots a set of local polynomial regression estimates on a sample of 1,471 aggregates at the province-wealth quintile level. These cover 296 region cells from 26 African countries. Province is the term we use to refer to the largest subnational unit in each country—the smallest geographical level at which DHS surveys are representative. The *x*-axis plots the share of young people aged 18–24 in a province-wealth quintile cell relative to the overall sample aged 18-49 in this cell. Each outcome plotted on the *y*-axis is constructed as a share of 18–24 year-old workers in each province-wealth quintile cell.

increases. In the second figure, individuals born in larger cohorts are less likely to have a salaried job in manufacturing and services regardless of their wealth, but the gradient is steeper for the middle classes. The poorest are unlikely to have these jobs regardless of cohort size, while the top quintiles are also unaffected for the opposite reason: they are most likely to have these jobs regardless of cohort size. Wage employment falls especially for the second, third, and fourth quintiles, and the second converges to the bottom. Thus, it seems that wealth acts as a buffer, as the probability that richer individuals have higher-quality jobs is less sensitive to cohort size than those in the lower wealth classes.

What Can Policy Do?

Young labor market entrants in Africa are at a disadvantage relative to their counterparts in other low-income settings because their supply of labor outstrips demand. What can be done? Creating new jobs for the young is a policy priority everywhere: hence, we can draw lessons from a large number of evaluations. Policies can be divided in three groups: those that give workers skills that make them more attractive to prospective employers, those that increase the demand for labor, and those that help firms and workers to match.

Supply-Side Policies: Vocational Skill Training and Apprenticeships

Vocational skills training programs are a very commonly used tool to improve young adults' employment opportunities. Their rationale is that the demand for skilled workers is strong, and that the lack of skills impedes young workers to fill these vacancies. These programs come in different durations, on different topics, and are implemented in a variety of settings and countries.

Randomized evaluations of skills programs indicate that most of these policies have a modest impact on employment. McKenzie (2017a) provides an overview of these programs across developing countries. "Modest" means that three out of 100 trainees find a job they would have not found without the intervention and that out of 14 experiments, only one had a statistically significant impact on earnings. Increasing the supply of skilled workers does not seem sufficient to create its own demand, at least not when implemented at the scale that is commonly evaluated.

Do vocational skills training programs have equally modest effects in Africa? The answer is much more nuanced. There are studies that find modest effects, like Cho et al. (2013) who evaluate a vocational training program in Malawi. There are, however, also studies that find large effects, including Honorati (2015) in Kenya and Alfonsi et al. (2020) in Uganda. The programs in these two studies stood out for being much longer (three months technical training coupled with three months internship in Kenya and six months of vocational training in Uganda) and both were delivered by high-quality providers. In Alfonsi et al. (2020), outcomes were measured for a longer time period, which is key as the effects are shown to grow over time. This suggests that if these training interventions are properly administered, they could be a helpful tool towards getting young people into good jobs.

Why is Africa different? Along with accidental differences in program quality, the selection of trainees is likely to differ. Most young African adults are jobless (or at best engaged in low-quality jobs), so these training programs treat people from the entire support of the ability distribution, whereas in richer countries, the most able find employment on their own, and skill training programs treat individuals with lower returns.

Another common form of supply-side intervention is subsidies to firm training or apprenticeships. Externalities in human capital investment put a cap on firms' willingness to train workers because once the skills are embodied in workers, other firms can benefit without contributing to the cost. Subsidies are meant to compensate for that. The meta-analysis in McKenzie (2017a) indicates modest effects for these programs too. One possible reason is that firms are generally reluctant to train workers even with generous wage subsidies, especially those in high-skill sectors: Alfonsi et al. (2020) find evidence along these lines in their study of Uganda, as do Caicedo, Espinosa, and Seibold (2021).

Demand-Side Policies: Boosting Firm Growth to Create Jobs

To expand the pool of salaried jobs available to young labor market entrants, another set of policies seek to relax constraints to firm entry and growth. Where subsistence self-employment dominates the employment structure, this involves subsidizing access to business training, capital, and labor.

Business training programs have proven to be largely ineffective in creating jobs (McKenzie and Woodruff 2014; Mackenzie et al. 2020). Africa is no exception in this regard. Premand et al. (2012), for example, provide evidence from Tunisia that although entrepreneurship trainings increase the likelihood of starting a new business and reduce the chance of business failure, the increases in self-employment are likely to be offset by reductions in wage employment, pointing to a partial substitution between the two types of work.

Similarly, microcredit has had only very limited effects on the growth of microenterprises on average. A meta-study of seven randomized evaluations of micro-finance shows that benefits concentrate among the larger and more profitable firms (Meager 2019). Of these, two studies are based in Africa. In rural Ethiopia, Tarozzi, Desai, and Johnson (2015) show no evidence that microcredit improves labor market outcomes despite a substantial increase in borrowing. On the contrary, in Morocco, the take-up rate of microcredit in rural areas tends to be extremely low, despite the availability of no formal credit alternatives (Crépon et al. 2015). However, among those who take up the microcredit, there are proportionally higher—albeit highly heterogeneous—impacts on self-employment investments and profits. These gains are offset by correspondingly large declines in employment income stemming from a fall in labor supplied outside the household.

A plausible alternative to create jobs and spur salaried employment would be to subsidize existing firms to hire workers. However, not many studies point to the existence of hiring frictions constraining firm size in developing countries. In Sri Lanka, offering wage subsidies to micro-enterprises to hire new workers had low take-up and only increased employment for as long as the subsidy remained in place (de Mel, McKenzie, and Woodruff 2019). One possibility is that firms are unable to find workers with the right skillset. Anderson and McKenzie (forthcoming) show that in Nigeria, helping firms recruit an accounting or marketing specialist improved business practices, increased product innovations activities, and raised profits and sales.

The average effects of programs seeking to boost microenterprise growth hide substantial heterogeneity. For instance, firms in manufacturing tend to be more responsive to wage subsidies, and interventions seem to be more effective for firms that are well-managed in the first place. This is in line with estimated returns to capital for microenterprises that, for instance, have been found to vary from negative to 13.8 percent per month in Sri Lanka (de Mel, McKenzie, and Woodruff 2008).

Indeed, recent research on Tanzania suggests that programs directed at stimulating entrepreneurship in low-income countries must take into account the very heterogeneous nature of these small firms. Using one of the only representative surveys of small firms in Africa, Diao, Magalhaes, and Mcmillan (2018) document that the share of the rural labor force working in nonfarm employment tripled (from 6.8 percent to 20.5 percent) between 2002 and 2012. But they also find extreme dispersion in labor productivity among these businesses, with many being less productive than agricultural activities. In another study, Diao, Kweka, and McMillan (2018) find that 94 percent of the country's labor productivity growth

over that period came from a very small subset of informal firms that share characteristics with firms in the formal sector.

If most self-employed workers in low-income countries run businesses out of necessity while for a small fraction with entrepreneurial talent it is a calling, a policy priority is to focus on identifying these "gazelles" (Grimm, Knorringa, and Lay 2012; Jayachandran 2021). Targeting entrepreneurship interventions towards this group could have important multiplier effects if the "out-of-necessity" microentrepreneurs eventually get hired by the firms whose growth was fostered by these policies. The results of a large-scale national business plan competition in Nigeria lend support to this recommendation. Random assignment of grants provided each winner with approximately US$50,000. Winning these grants led to firm entry and growth, including increases of over 20 percentage points in the likelihood of a firm having ten or more workers (McKenzie 2017b).

Increasing the size of the market to which firms cater can also be a way to boost salaried employment growth. McCaig and Pavcnik (2018) show that in Vietnam, a low-income country with significant self-employment, an increase in access to export markets promoted employment in larger firms.[5]

Another way of generating salaried jobs is by attracting foreign firms. This can increase labor demand directly or indirectly through its effects on domestic firms. In Costa Rica, Méndez-Chacon and Van Patten (2021) find that domestic firms experience a 26 percent increase in their number of employees after selling to multinational buyers. In Africa, a recent paper by Mendola, Prarolo, and Sonno (2021) finds a clear positive association between proximity to multinational enterprises and employment. A review of the effects of foreign direct investment on local jobs also suggests that foreign direct investment is associated with an increase in the number of "good jobs" in developing countries, because the jobs created by foreign firms tend to pay more and tend to offer more training than local firms, and inflows of foreign direct investment boost the aggregate productivity of the host country (Javorcik 2015).

Matching Demand and Supply

Other market failures, resulting from information frictions, impede efficient matching of skilled workers to existing jobs. Information asymmetries in credit markets affect both workers who cannot afford to invest in skills and firms that cannot afford to invest to expand their business and hence have low labor demand. Such frictions also directly increase how firms perceive the cost of hiring labor, especially for new entrants whose type is unknown to the firm. Evidence of this includes the use of entrance fees as screening mechanisms in West African apprenticeships. In Ghana, most firms claim that these serve to force young workers to signal investment in the apprenticeship (Hardy and McCasland 2021).

[5]The positive impacts of offering firms access to foreign markets extend to increased profits and improvements in product quality, as indicated by an experiment on small rug manufacturers in Egypt which randomly assigned the opportunity to export to high-income markets (Atkin, Khandelwal, and Osman 2017).

Interventions that allow workers to signal their skills in a credible way relax this information constraint—namely, they reduce the uncertainty of employers on workers' type and make it worthwhile for them to hire. For example, Abebe et al. (2021b) in their study in Addis Ababa look at the effects of a job application workshop that includes advice on how to apply and interview, along with standardized personal selection tests that offer a verifiable measure of skills. Young workers randomly assigned to the workshop were significantly more likely to find stable salaried jobs. Bassi and Nansamba (2021), using a randomized trial in Uganda, find that making the results of a verifiable test of noncognitive skills available to both workers and employers encourages hiring. Assessing the skills of jobseekers and giving them their assessment results in a certificate they can credibly share with firms also increased employment in South Africa (Carranza et al. 2021).

These policies might also change workers' expectations and subsequently, through these, outcomes. Bandiera et al. (2021) show that vocational training makes young workers in Uganda more optimistic about finding employment in high-wage firms, thus they are more likely to search and more likely to be employed in such firms. In the same setting, an offer to introduce workers to firms backfired by lowering expectations, limiting search, and consequently, increasing the likelihood of working for a lower wage in a low-wage firm.

Expanding Job Search

Enabling young liquidity-constrained workers to expand their search is perhaps the most promising route. Paying application costs (about $4.50 in US dollars) turned out to be as effective as doubling the wage rate in increasing the quality and the number of applicants for a clerical job in Ethiopia's capital city of Addis Ababa (Abebe et al. 2021a). Similarly, interventions as simple as subsidizing the search for salaried jobs have also had some positive results. Examples include the reimbursement of travel costs to attend an interview or to find a job in a nearby city. Offering transport subsidies for young jobseekers in Addis Ababa, where salaried jobs are typically advertised in the city center, had large positive effects on the probability of finding a formal job (Franklin 2018; Abebe et al. 2021b). A small transfer covering the bus fare was sufficient to induce a large number of households to send one of their members to work in the city during the lean season in Bangladesh, resulting in higher consumption and regular migration even after the transfer was removed (Bryan, Chowdhury, and Mobarak 2014).

Progress in digital technology has opened up employment possibilities globally. The arrival of fast internet has been shown to increase employment rates on the African continent due to the technology's impact on firm entry, productivity, and exports (Hjort and Poulsen 2019a). The expansion of the internet network encouraged firm entry, primarily in sectors with a reliance on information and communications technology, such as finance, in South Africa. More broadly, the spread of information and communications technology in Africa in recent years holds promise for creating jobs and reallocating workers into higher-return

Figure 7
Internet Access and Salaried Jobs for Young Adults in Africa

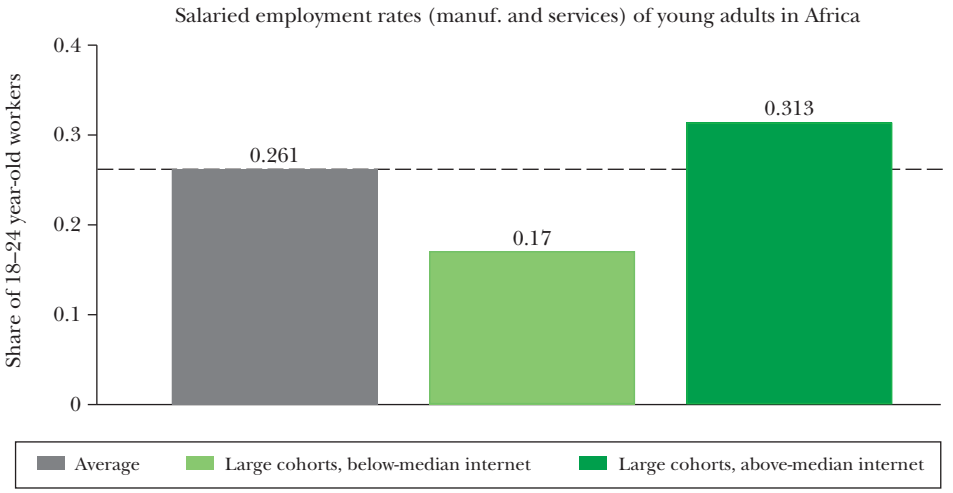

Source: Demographic and Health Surveys and IPUMS (harmonized via the Jobs of the World Project) and ITU World Telecommunication/ICT Indicators Database.

Note: This figure is based on sample of 22 African countries constructed by merging the country-level aggregates from our dataset with Internet Telecommunications Union data at the country-year level. We keep only the latest survey year available for each country. The first bar plots the average of the country-level shares of 18–24 year-old workers with salaried jobs in manufacturing/services. The second bar plots the average share in the group of countries with above-median cohort size and below-median levels of internet penetration, defined as the percentage of households using internet. The third bar plots the average share in the countries with above-median cohort size and above-median levels of internet penetration.

occupations in Africa. For example, in Kenya, the expansion of the mobile money technology M-PESA has led women to shift away from agriculture and into business (Suri and Jack 2016).

Merging the country-level means from our dataset with internet use data provided by the International Telecommunications Union across African countries yields correlations that corroborate these findings.[6] In countries where internet use is more widespread, young adults are more likely to have a salaried job. Most importantly, internet access can undo, in part, the effect of large cohorts.

This is illustrated in Figure 7, which we construct from our final merged sample covering 22 African countries. The gray bar plots the average share of 18–24 year-old workers who have salaried jobs in manufacturing and services in the latest year available in the dataset: 26.1 percent. The light-green bar in the middle plots this average for the countries characterized by relatively large young cohorts compared

[6]The percentage of households using internet is a country-year-level indicator provided by ITU World Telecommunication/ICT Indicators Database. This data was retrieved from Hjort and Poulsen (2019b).

to the region median and relatively low internet penetration (namely, Guinea 2014, Mali 2009, Rwanda 2012, Cameroon 2005, Lesotho 2009, Kenya 2008). This number is lower than the overall average by 9 percentage points, at 17 percent. Finally, the last bar in dark green plots the average share of young adults in salaried jobs outside of agriculture for countries with relatively large cohorts but above-median rates of internet use (which includes Benin 2013, Egypt 2006, Namibia 2006, Senegal 2013 and Zambia 2010). For this group, the average lies at 31.3 percent, about 5 percentage points above the regional average in this sample. The positive association between young adults' chance of finding a salaried job outside of agriculture and the spread of internet therefore seems present in the data even for relatively large cohorts.

Lessons from Theory: Multiple Equilibria

Taken together, active labor market policies tend to yield stronger outcomes in Africa compared to other developing countries. However, looking at the broader picture across low-income countries, the evidence from active labor market policies seems puzzling: the average worker does not fare better after being trained and firms do not hire more workers even when the cost of labor is low. At the same time, grants as cheap as a bus ticket or a standardized test may transform the lives of those who receive them by connecting them to jobs at a tiny fraction of the cost of vocational training. This disparate set of findings can be connected in a coherent picture if the underlying economic model has multiple equilibria and poverty traps.

The existence of multiple equilibria implies the existence of a threshold of resources below which individuals are trapped in the low equilibrium and above which they escape. Micro interventions that tap into a pool of misallocated resources, such as credit-constrained job-seekers in Ethiopia, can have huge effects because they allow large groups of people to move from one equilibrium to another.

At the macro level, the low-productivity equilibrium has many small firms and only a few low-wage salaried jobs. Market failures create a vicious cycle: if firm owners cannot borrow to invest, there will be a handful of firms, and among these, only low-productivity, poorly managed firms can survive because entrepreneurs with higher growth potential cannot enter the market (Bandiera et al. 2017b; Bloom and Van Reenen 2010). Besides its implications for the misallocation of entrepreneurial talent, the lack of labor specialization keeps productivity and wages low.

In the high-productivity equilibrium, larger firms employ the majority of the population. Individuals with entrepreneurial talent can access funds to start their own firm, and at any point in time only the most productive firms are present in the market. Because skills are rewarded, individuals have incentives to invest, and because wages are high, they can afford to.

The existence of these multiple equilibria can explain why apparently identical active labor market policy programs can have very different effects and why the same program can have very different effects on apparently identical people:

changes that shift people across equilibria have very large and permanent effects. The facts presented in this paper suggest that many African countries may be in the low equilibrium, where improving the supply of skilled workers is ineffective because the demand for salaried labor is low.

Big Push Policies

Policies that seek to assist the low-productivity firms that exist in a low-productivity equilibrium often do not help move the economy to the high equilibrium. Policies that enable people with entrepreneurial talent to start and grow firms or facilitate the exit of low productivity firms all have the potential to shift the economy to the high productivity equilibrium. Examples of such policies include large credit expansion (Burgess and Pande 2005) and large-scale infrastructure investment in transport, energy, and communications. In Ethiopia, bundled investments in roads and electrification shifted labor out of agriculture and into manufacturing and services (Moneke 2020). In Indonesia, electrification has been shown to foster positive selection by driving the least productive firms out of the market to be replaced by new, more productive entrants (Kassem 2020). In addition, better access to foreign capital reduces misallocation of resources in the economy and stimulates growth and employment, with the highest gains recorded in areas where local capital markets were least developed (Bau and Matray 2021). Because family ownership lowers the quality of management practices (Bandiera et al. 2017b; Lemos and Scur 2019), tax policies aimed at curbing intra-family transfers of businesses could also lead to improvements in long-term productivity. Such policies could come in different shapes, including wealth taxes, inheritance taxes, or a reduction in the exemptions that family firms often enjoy (Tsoutsoura 2015). Other "big-push" institutional policies include improvements in contract enforcement that allow owners to delegate management to professionals (Grobovšek 2020).

Taken together, both empirical evidence and lessons from theory suggest that a policy focus on the creation of institutions and infrastructure that support firms in growing to larger sizes can help broaden the pool of salaried employment opportunities available to labor market entrants in Africa.

Conclusion

Many African economies are in a vicious cycle, where most people run subsistence enterprises because there are no salaried jobs and there are no salaried jobs because most enterprises operate at subsistence levels and are not profitable enough to expand. Up-skilling young labor market entrants will not make subsistence entrepreneurs any more likely to hire them, and subsidizing subsistence entrepreneurs will not make them more productive. These policies fail because they take the organization of labor as given. But no high-income economy is made of a myriad of entrepreneurs running tiny firms. Rather, most economic activity takes place in large organizations that put together workers with complementary skills.

How might the countries of Africa get there? "Getting there" requires incentives for microenterprises to merge, which will only happen if there is sufficient demand. The additional demand can be subsidized by the state, perhaps even via direct procurement. The alternative is to give workers access to international markets where salaried jobs are easily available. Until recently this possibility required physical migration, which came with high transaction costs. Today, services can be provided online. If young African workers can sell their services to high-income countries, their earnings could create the boost in demand that their countries need to jump-start growth and a modernized organization of labor.

■ *We thank Gabriel Leite Mariante and Caterina Soto Vieira for their exceptional contribution in assembling the data, and Anders Jensen and Tommaso Sonno for sharing their data. We thank the editors, Heidi Williams, Nina Pavcnik, Erik Hurst and Timothy Taylor, for comments and suggestions that substantially improved this paper. The data building and harmonization have been funded by the joint IZA/FCDO Gender, Growth, and Labour Markets in Low Income Countries Programme ($G^2LM|LIC$), we thank Hilmar Schneider, Marco Caliendo, and Benedetta Musillo for supporting this endeavour.*

References

Abebe, Girum, A. Stefano Caria, and Esteban Ortiz-Ospina. 2021a. "The Selection of Talent: Experimental and Structural Evidence from Ethiopia." *American Economic Review* 111 (6): 1757–1806.

Abebe, Girum, A. Stefano Caria, Marcel Fafchamps, Paolo Falco, Simon Franklin, and Simon Quinn. 2021b. "Anonymity or Distance? Job Search and Labour Market Exclusion in a Growing African City." *Review of Economic Studies* 88 (3): 1279–1310.

Alfonsi, Livia, Oriana Bandiera, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali. 2020. "Tackling Youth Unemployment: Evidence from a Labor Market Experiment in Uganda." *Econometrica* 88 (6): 2369–2414.

Anderson, Stephen J., and David McKenzie. Forthcoming. "Improving Business Practices and the Boundary of the Entrepreneur: A Randomized Experiment Comparing Training, Consulting, Insourcing, and Outsourcing." *Journal of Political Economy*.

Atkin, David, Amit K. Khandelwal, and Adam Osman. 2017. "Exporting and Firm Performance: Evidence from a Randomized Experiment." *Quarterly Journal of Economics* 132 (2): 551–615.

Bandiera, O., Elsayed, A., and Smurra, A. 2022. "Jobs of the World Project: Towards Better Understanding of Labor Markets in Low- and Middle-Income Countries." Unpublished.

Bandiera, Oriana, Renata Lemos, Andrea Prat, and Raffaella Sadun. 2017. "Managing the Family Firm: Evidence from CEOs at Work." *Review of Financial Studies* 31 (5): 1605–53.

Bandiera, Oriana, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali. 2021. "The Search for Good Jobs: Evidence from a Six-Year Experiment in Uganda." https://be716900-bf6c-4d3f-9bdb-e679de36b900.filesusr.com/ugd/997323_ae8ebbf27d06496aa09629710580856e.pdf.

Bassi, Vittorio, and Aisha Nansamba. 2021. "Screening and Signaling Non-Cognitive Skills: Experimental Evidence from Uganda." *Economic Journal*. https://doi.org/10.1093/ej/ueab071.

Bau, Natalie, and Adrien Matray. 2021. "Misallocation and Capital Market Integration: Evidence from India." NBER Working Paper 27955.

Bloom, Nicholas, and John Van Reenen. 2010. "Why Do Management Practices Differ Across Firms and Countries?" *Journal of Economic Perspectives* 24 (1): 203–24.

Bongaarts, John, and John Casterline. 2012. "Fertility Transition: Is Sub-Saharan Africa Different?" *Population and Development Review* 38 (1): 153–68.

Bongaarts, John. 2017. "Africa's Unique Fertility Transition." *Population and Development Review* 43 (S1): 39–58.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2014. "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh." *Econometrica* 82 (5):1671–1748.

Burgess, Robin, and Rohini Pande. 2005. "Do Rural Banks Matter? Evidence from the Indian Social Banking Experiment." *American Economic Review* 95 (3): 780–95.

Caicedo, Santiago, Miguel Espinosa, and Arthur Seibold. 2021. "Unwilling to Train? Firm Responses to the Colombian Apprenticeship Regulation." Unpublished.

Carranza Eliana, Robert Garlick, Kate Orkin, and Neil Rankin. 2021. "Job Search and Hiring with Two-Sided Limited Information About Workseekers' Skills." Unpublished.

Cho, Yoonyoung, Davie Kalomba, Ahmed Mushfiq Mobarak, and Victor Orozco-Olvera. 2013. "Gender Differences in the Effects of Vocational Training: Constraints on Women and Drop-Out Behaviour." World Bank Policy Research Working Paper 6545.

Crépon, Bruno, Florencia Devoto, Esther Duflo, and William Parienté. 2015. "Estimating the Impact of Microcredit on Those Who Take It Up: Evidence from a Randomized Experiment in Morocco." *American Economic Journal: Applied Economics* 7 (1): 123–50.

de Mel, Suresh, David McKenzie, and Christopher Woodruff. 2008. "Returns to Capital in Microenterprises: Evidence from a Field Experiment." *Quarterly Journal of Economics* 123 (4): 1329–72.

de Mel, Suresh, David McKenzie, and Christopher Woodruff. 2019. "Labor Drops: Experimental Evidence on the Return to Additional Labor in Microenterprises." *American Economic Journal: Applied Economics* 11 (1): 202–35.

Diao, Xinshen, Kenneth Harttgen, and Margaret McMillan. 2017. "The Changing Structure of Africa's Economies." *World Bank Economic Review* 31 (2): 412–33.

Diao, Xinshen, Josephat Kweka, and Margeret McMillan. 2018. "Small Firms, Structural Change and Labor Productivity Growth in Africa: Evidence from Tanzania." *World Development* 105 (1): 400–15.

Diao, Xinshen, Eduardo Magalhaes, and Margaret Mcmillan. 2018. "Understanding the Role of Rural Non-Farm Enterprises in Africa's Economic Transformation: Evidence from Tanzania." *Journal of Development Studies* 54 (5): 833–55.

Donovan, Kevin, Will Jianyu Lu, and Todd Schoellman. 2021. "Labor Market Dynamics and Development." Unpublished.

Feng, Ying, David Lagakos, and James E. Rauch. 2021. "Unemployment and Development." Unpublished.

Franklin, Simon. 2018. "Location, Search Costs and Youth Unemployment: Experimental Evidence from Transport Subsidies." *Economic Journal 128 (614): 2353–79.*

Grimm Michael, Peter Knorringa, and Jann Lay. 2012. "Constrained Gazelles: High Potentials in West Africa's Informal Economy." *World Development* 40 (7): 1352–68.

Grobovšek, Jan. 2020. "Managerial Delegation, Law Enforcement, and Aggregate Productivity." *Review of Economic Studies* 87 (5): 2256–89.

Hardy, Morgan, and Jamie McCasland. 2021. "Are Small Firms Labor Constrained? Experimental Evidence from Ghana." Unpublished.

Hjort, Jonas, and Jonas Poulsen. 2019a. "The Arrival of Fast Internet and Employment in Africa." *American Economic Review* 109 (3): 1032–79.

Hjort, J. and Poulsen, J. 2019b. Replication data for: The Arrival of Fast Internet and Employment in Africa. Nashville, TN: American Economic Association [publisher]. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2019-10-12. https://doi.org/10.3886/E113156V1.

Honorati, Maddalena. 2015. "The Impact of Private Sector Internship and Training on Urban Youth in Kenya." World Bank Policy Research Working Paper 7404.

ICF. 1991-2017. Demographic and Health Surveys (various) [Datasets]. Funded by USAID. Rockville, Maryland: ICF [Distributor].

Javorcik, Beata S. 2015. "Does FDI Bring Good Jobs to Host Countries?" *World Bank Research Observer* 30 (1): 74–94.

Jayachandran, Seema. 2021. "Microentrepreneurship in Developing Countries." In *Handbook of Labor, Human Resources and Population Economics*, edited by Klaus F. Zimmermann, 1–31. Berling: Springer.

Jensen, Anders. 2021. "Employment Structure and the Rise of the Modern Tax System." NBER Working Paper 25502.

Kassem, Dana. 2020. "Does Electrification Cause Industrial Development? Grid Expansion and Firm Turnover in Indonesia." Unpublished.

Lam, David, Murray Leibbrandt, and James Allen. 2019. "The Demography of the Labor Force in Sub-Saharan Africa: Challenges and Opportunities." Growth and Labour Markets in Low Income Countries Synthesis Paper 10.

La Porta, Rafael, and Andrei Shleifer. 2008. "The Unofficial Economy and Economic Development." *Brookings Papers on Economic Activity* 2008 (Fall): 275–363.

La Porta, Rafael, and Andrei Shleifer. 2014. "Informality and Development." *Journal of Economic Perspectives* 28 (3): 109–26.

Lemos, Renata, and Daniela Scur. 2019. "**The Ties That Bind: Implicit Contracts and the Adoption of Management Technology in the Firm**." Unpublished.

McCaig, Brian, and Nina Pavcnik. 2018. "Export Markets and Labor Allocation in a Low-Income Country." *American Economic Review* 108 (7): 1899–1941.

McKenzie, David. 2017a. "How Effective Are Active Labor Market Policies in Low-income countries? A Critical Review of Recent Evidence." *World Bank Research Observer* 32 (2): 127–54.

McKenzie, David. 2017b. "Identifying and Spurring High-Growth Entrepreneurship: Experimental Evidence from a Business Plan Competition." *American Economic Review* 107 (8): 2278–307.

McKenzie, David, and Christopher Woodruff. 2014. "What Are We Learning from Business Training and Entrepreneurship Evaluations around the Developing World?" *World Bank Research Observer* 29 (1): 48–82.

McKenzie, David, Christopher Woodruff, Kjetil Bjorvatn, Miriam Bruhn, Jing Cai, Juanita Gonzalez Uribe, Simon Quinn, Tetsushi Sonobe, and Martin Valdivia. 2020. "Training Entrepreneurs." *VoxDevLit* 1 (1): 3.

Meager, Rachael. 2019. "Understanding the Average Impact of Microcredit Expansions: A Bayesian Hierarchical Analysis of Seven Randomized Experiments." *American Economic Journal: Applied Economics* 11 (1): 57–91.

Méndez-Chacon, Esteban, and Diana Van Patten. 2021. "Multinationals, Monopsony, and Local Development: Evidence from the United Fruit Company." Unpublished.

Mendola, Mariapia, Giovanni Prarolo, and Tommaso Sonno. 2021. "Curse or Blessing? Multinational Corporations and Labour Supply in Africa." Unpublished.

Minnesota Population Center. Integrated Public Use Microdata Series, International: Version 7.3, various countries and years. Minneapolis, MN: IPUMS, 2020. https://doi.org/10.18128/D020.V7.3, (accessed December 29, 2021).

Moneke, Niclas. 2020. "Can Big Push Infrastructure Unlock Economic Development? Evidence from Ethiopia." Unpublished.

Premand, Patrick, Stefanie Brodmann, Rita Almeida, Rebekka Grun, and Mahdi Barouni. 2012. "Entrepreneurship Training and Self-Employment Among University Graduates: Evidence from Randomized Trial in Tunisia." World Bank Policy Research Working Paper 6285.

Schoar, Antoinette. 2010. "The Divide Between Subsistence and Transformational Entrepreneurship." *Innovation Policy and the Economy* 10 (1): 57–81

Suri, Tavneet, and William Jack. 2016. "The Long-Run Poverty and Gender Impacts of Mobile Money." *Science* 354 (6317): 1288–92.

Tarozzi, Alessandro, Jaikishan Desai, and Kristin Johnson. 2015. "The Impacts of Microcredit: Evidence from Ethiopia." *American Economic Journal: Applied Economics* 7 (1): 54–89.

Tsoutsoura, Margarita. 2015. "The Effect of Succession Taxes on Family Firm Investment: Evidence from a Natural Experiment." *Journal of Finance* 70 (2): 649–88.

UN-World Population Prospects. 2019. "United Nations, Department of Economic and Social Affairs, Population Division." (accessed December 29, 2021).

World Bank. 2021. "Agriculture, Forestry, and Fishing, value added (% of GDP)". World Bank Group. https://data.worldbank.org/indicator/NV.AGR.TOTL.ZS (accessed June 10, 2021).

Zipfel, Céline. 2021. "Defusing a Population Explosion? Jobs and Fertility in sub-Saharan Africa". Unpublished.

Journal of Economic Perspectives—Volume 36, Number 1—Winter 2022—Pages 101–124

Political Distortions, State Capture, and Economic Development in Africa

Nathan Canen and Leonard Wantchekon

A classic question in economics is how to explain differences in economic growth and development across countries. For researchers focused on sub-Saharan Africa, the corresponding question has long been how to explain what Easterly and Levine (1997) memorably called the "African growth tragedy."

A standard starting point to study these differences has been to look at disparities in the factors of production like human capital and physical capital, as well as technology. However, growth accounting exercises suggest that such explanations account for, at most, 50 percent of this variation in salient settings (for example, Hsieh and Klenow 2009). Thus, the question became why the economies of countries were performing so much better or worse than one might expect based on their factors of production. For much of the developing world, newly developed explanations over the past 20 years include frictions in transportation, market access and information, inefficient regulation, as well as in the disparity in the enforcement of the rule of law, to name a few. That is, politics, technology, and history might all be relevant in explaining such differences in growth through inducing suboptimal choices by firms and citizens. With such mechanisms, there is a clear scope for policy to improve welfare by alleviating these frictions.

■ *Nathan Canen is Assistant Professor of Economics, University of Houston, Houston, Texas. Leonard Wantchekon is James Madison Professor of Political Economy, Professor of Politics and International Affairs, Princeton University, Princeton, New Jersey. He is also a Research Associate at the National Bureau of Economic Research, Cambridge, Massachusetts, and the Founder and President of the African School of Economics, Abomey-Calavi, Benin. Their email addresses are ncanen@uh.edu and lwantche@princeton.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.36.1.101.

However, when it comes to Africa, the explanations for its economic development are mostly focused on long-run structural factors. In particular, mechanisms like colonial and pre-colonial institutions, deep-rooted historical factors, and ethnic diversity have been proposed as the key explanations for slow economic growth in Africa (for a review, see Michalopoulos and Papaioannou 2020). It has been extensively argued that certain political and judicial institutions set up by colonial powers—such as judicial systems, colonial concessions, and the organization of government—were designed for resource extraction, and thus, such specific institutions did not secure property rights and have had persistently negative effects on development outcomes since then (for example, Acemoglu, Johnson, and Robinson 2001). Another strand of the literature argues that the fractionalization and heterogeneity of cultures in Africa has induced conflict and the distortion of public good provision (for example, Alesina et al. 2016; Baldwin and Huber 2010 and references therein).

We find that such African-centered and historical arguments are limited in their ability to explain key facts and narratives about current African economic performance. Most notably, they miss the fast GDP growth rates in many countries (for example, Rwanda, Ghana, and Ethiopia have each doubled their GDP over the past 20 years), the changing political institutions in these countries (including the wave of democratization across Africa in the 1990s), and the regional variation in outcomes, even in contexts with the same colonial past or ethnic compositions. For instance, a very productive information technology sector in Kenya, or the increasingly productive agricultural sectors in Ethiopia, can coexist with less productive industries in the same country. As a result, such long-run theories often explain little of the underlying variation in the data: for example, slavery only explains 15–25 percent of the variation in current trust levels (Nunn and Wantchekon 2011), while pre-colonial institutions might explain less than 10 percent of the variation in current economic outcomes (Michalopoulos and Papaioannou 2013). This is not to say the structural explanations are unimportant; historical experiences certainly play a role in current outcomes. However, in our view, they are better thought of instead as constraints on current government, firm, and citizen decision-making. Within such a framework, researchers can think of counterfactual policies that may ease such constraints, without overemphasizing historical and deterministic explanations at the expense of policies taken by those agents.

As a result, in this paper we argue for researchers and policymakers to focus on political distortions to address the nature of economic development and growth in sub-Saharan Africa and elsewhere, and to move away from placing disproportionate emphasis on historical arguments. A political distortion is any situation in which a special interest group can direct economic development towards its own exclusive ends rather than towards increasing general welfare. It also captures situations in which the provision of public goods, public investments, and redistribution is instead primarily motivated by the narrow political interests of office-holders, or of its political connections, rather than by considerations of public interest and general welfare. This concept thus encompasses phenomena like "state capture,"

whereby private interests significantly influence state decision-making (whether through control of the bureaucracy, campaign contributions, connections, or some other channel), as well as clientelism and targeted redistribution, whereby politicians' actions reflect social ties and exchanges of government transfers for political support. Examples of such behavior are widespread in Africa and elsewhere, and frequently exact large costs on economic outcomes and citizen welfare. Salient cases include the Gupta brothers' capture of the South African state under the former President Zuma and the concession of the port of Lomé (Togo) to the French conglomerate Bolloré in exchange for undercharging consulting fees. Other examples, including at the local level, are discussed below.

Discussing economic development in Africa through the lens of political distortions has important advantages: it provides an umbrella concept which integrates the study of economic development across different regions, whether sub-Saharan Africa, Latin America, Western Europe, or elsewhere. As institutions and economic characteristics around the world look increasingly alike (as we discuss in the next section), the same can be said about political and economic incentives. This approach also leaves a wide scope for policy analysis, which is often largely absent from historical and more deterministic accounts. Political distortions are induced by strategic choices and influenced by incentives, which can be eased with alternative policies such as restrictions on campaign contributions, bureaucratic reform, audits, multinational initiatives for free trade, the availability of information through debates, and the introduction of new technologies within government. Such policies have been shown to be effective in a variety of settings and may present possible welfare gains in sub-Saharan Africa and the continent more generally.

In fact, the emphasis on political distortions can be complementary to other types of distortions studied in the economics literature. For example, the challenges of addressing market failures, or credit, labor, or information frictions, or providing banking credit to certain sectors, or implementing improved educational choices, all depend both on existing economic scenarios (including information, public infrastructure, and so on), but also on political institutions that are required to implement them, including those institutions in charge of contract enforcement. Hence, political distortions are critical in helping sustain and expand policy effectiveness. Indeed, political distortions are particularly prominent, because they are typically present whenever other frictions arise due to government policy.

In the next section, we review the arguments for long-run/structural explanations for African development and what they may miss in the data. In the following section, we provide examples of the way political distortions shape market structure and other economic outcomes. This includes how political connections can hinder public good and infrastructure provision, contract enforcement, and "Schumpeterian growth" (where innovative firms should be replacing old firms). Such examples are interpreted relative to theoretical models and discussed together with empirical work. Although our focus here is on sub-Saharan Africa in particular, we also mention some evidence from other countries and contexts and discuss how they can be informative for the sub-Saharan context under this common theoretical

framework. We then use the theoretical and empirical results to sketch an array of policy reforms that could limit political distortions. Some possibilities include banning corporate campaign contributions, reforms insulating regulators from firm influence, how the new African Continental Free Trade Agreement will help to reform governance, more widespread use of tools of participatory democracy (like town hall meetings), randomized audits of politician behavior, and the use of digital technologies for identification and tax collection.

Historical Explanations, Novel Facts, and Searching for a New Framework

A substantial literature in recent years has focused on long-run "structural" mechanisms to explain patterns of African economic development, which focus on deep-rooted historical, technological, or social factors, including historical patterns of colonial and pre-colonial institutions. Such approaches have great value: for example, they can help make sense of patterns of persistence and path-dependency.

However, over the last few decades, Africa's political institutions and growth patterns have been converging toward global norms, with the emergence of constitutional democracies and the consolidation of the rule of law, alongside incredibly high rates of economic growth in some regions. In addition, there is extensive within-country variation in economic outcomes—such as GDP, trade, and consumption—even across regions subject to similar historical institutions (like slavery, colonial institutions, and others). Explanations of African economic growth need to incorporate these facts more than they do at present, as they deviate from Africa-specific historical or structural factors.

For example, sub-Saharan Africa saw a dramatic rise in democratic institutions of governance during the third wave of democratization in the 1990s, with Zambia, Cape Verde, and Benin as salient examples. This was spurred by the spread of democratic ideas, the end of the Cold War and the fall of the Soviet Union, the creation of robust local democratic communities, and the implementation of economic reforms (Huntington 1991). While only Botswana and Mauritius held regular multiparty elections by 1989, 33 of the region's countries had held at least two sets of elections by late 2003 (Crawford and Lynch 2012). Figure 1 illustrates this change with data from the widely used Polity V database, produced by the Center for Systemic Peace, which collects components of governing institutions in 167 countries. These components are merged into an overall scale ranging from –10 (think "hereditary monarchy") to +10 (consolidated democracy). Autocracies are scored from –5 to –10, and as Figure 1 shows, sub-Saharan Africa as a whole was in that category for much of the 1970s and 1980s. Since then, the Polity score for sub-Saharan Africa has risen substantially, approaching average world levels.

Economic outcomes in sub-Saharan Africa have also been converging to world norms. During the past 20 years, average GDP per capita in sub-Saharan Africa has more than doubled: from about $600 to close to $1600 (comparison using current

Figure 1
Institutions in Sub-Saharan Africa Converging to World Levels since the 1990s
(average Polity score by region)

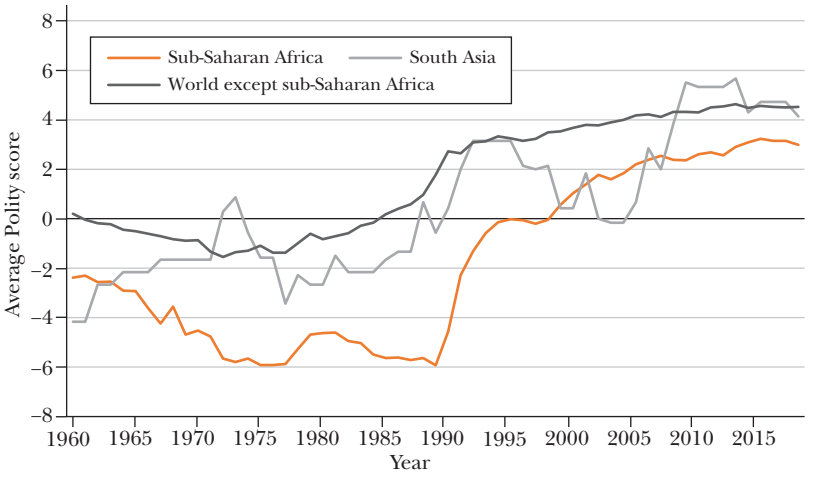

Source: Center for Systemic Peace (https://www.systemicpeace.org/inscrdata.html).
Note: The graph presents the evolution of the average Polity scores (Center for Systematic Peace) per region, across selected regions. The score measures the level of democratic institutions and ranges from –10 to 10, with –10 to –6 defining "autocracies", and +6 to +10 "democracies." Levels of executive recruitment, constraints on executive authority, and political competition are all part of the score.

US dollars, World Development Indicators data, as of June 2021). This wave of economic growth across sub-Saharan Africa is admittedly uneven. But while some countries still lag, economic growth rates in Rwanda, Ghana, and Ethiopia over the past 20 years resemble those in China and India, and the regional growth rates for Africa are comparable to those in regions like East Asia and Latin America. Kremer et al. (2021) document this convergence in economic performance, institutions, and policies around the world since the 1960s. Along the same lines, Patel et al. (2021) provide basic evidence on unconditional convergence and sustained growth performance in developing countries. Finally, Transparency International reports a Corruption Perception Index which measures how country experts, entrepreneurs, and firm-owners perceive corruptions in government and in the public sector. The measure aggregates 13 different datasets and ranges from 0 (highest perceived corruption) to 100 (lowest perceived corruption). As shown in Figure 2, while corruption is perceived on average to be higher in sub-Saharan Africa than other regions, those levels are similar in magnitude to those in Eastern Europe/Central Asia and the Middle East.

Figure 2
Corruption Perception Index by Region

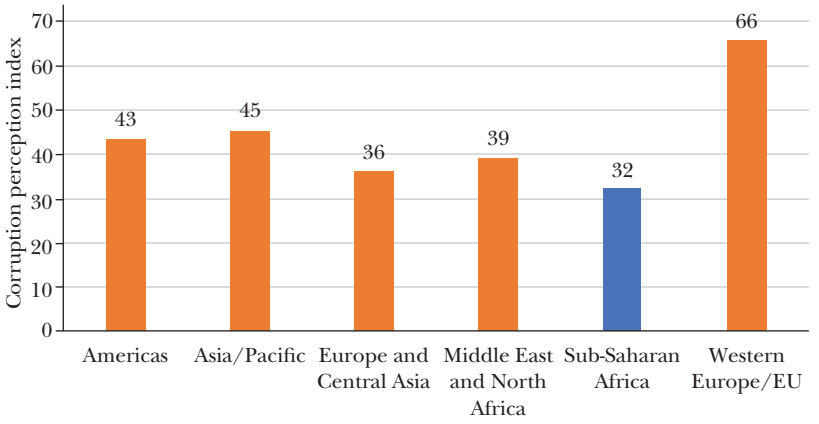

Note: The graph presents the average Corruption Perception Index (from Transparency International) by region. The score measures how citizens in each country perceive the corruption levels to be. The score ranges from 0 (highest perceived corruption) to 100 (lowest perceived corruption).

We do not find that such patterns can be easily explained by existing theories of development in Africa centered around pre-colonial or colonial institutions, or those based on ethnic compositions. To clarify our perspective, we focus on two of the main arguments about long-term historical and structural factors affecting economic growth in Africa: the types of colonial (or pre-colonial) institutions and the degree of ethnic fractionalization.

Consider the literature on how colonial and pre-colonial institutions are also perceived as one of the main drivers of development challenges in Africa (Michalopoulos and Papaioannou 2020). In one prominent example, Acemoglu and Robinson (2010, p. 22–23) argue that Africa's development has lagged because it has historically lacked strong centralized states, and state institutions were highly absolutist and patrimonial when they did exist. They argue that colonialism made the situation worse by imposing extractive land and labor institutions (as in South Africa and Kenya) and through excessive taxation and policies of indirect rule giving local elites even more power (as in Ghana). In a similar vein, Michalopoulos and Papaioannou (2013) find that pre-colonial institutions matter for development. More specifically, ethnic groups that had more centralized or more complex pre-colonial forms of political organizations have higher levels of development today.

The framing of these arguments suggests that Africa was predisposed to underdevelopment due to powerful historical constraints that cannot be easily overcome, although the evidence outlined above suggests that other mechanisms beyond long-run factors are at play.

Clearly, past political institutions (for example, regarding protection of property rights) have been a factor in disparities in public good allocations, with consequential outcomes today. However, they may not be the main mechanism by which such distortions occur. First, such institutions cannot explain the extensive variation in outcomes within countries with the same colonial or pre-colonial institutions. For example, even within the same country or regions, economic outcomes can differ significantly across sectors: the success of the information and flower sectors in Kenya, the airline sector in Ethiopia, or the movie industry in Nigeria coexists with less productive sectors in the same regions within the country. This is reflected in how such approaches typically explain only a minor part of the variation in per capita income across nations. A variance decomposition analysis in Nunn and Wantchekon (2011) suggests that slavery can explain about 15–25 percent of the variation in trust levels, leaving over 75 percent of them to be explained by short-term factors. Meanwhile, the relevant variable in the study on pre-colonial institutions by Michalopoulos and Papaioannou (2013, Table 2, column 1) has an R^2 of around 6 percent.

From our perspective, history clearly matters—but history is unlikely to matter more in Africa than elsewhere. States like Bihar, Jharkhand, and Uttar Pradesh within India have similar economic outcomes to the poorest African countries (Drèze and Sen 2013), but a very different colonial history. Other countries, like many in Latin America, also went through colonization with extractive institutions, but have reached quite different levels of development since then. In 1975, one could have easily made similar arguments about China—that it was underdeveloped due to a history of absolutism and communism (Mühlhahn 2019). However, studies on China soon started to emphasize how trade and public investment can spur economic growth (for example, Hu and Khan 1997).

Furthermore, the colonial persistence explanation might not apply equally well throughout the African continent. Huillery (2009) studies the persistent impact of colonial investments in education in Francophone Africa and finds that during the period of colonialism, investments tended to go to less politically centralized and prosperous areas. In Benin, except for the current capital city of Porto Novo, the first colonial schools and medical facilities were created in the periphery of the capital of the precolonial states, such as Abomey, Nikki, and Kouande (Garcia 1971), in contrast to the centralization argument outlined above.

The salient historical and structural considerations in the context of Africa have also often involved ethnic conflict (for an overview, see Michalopoulos and Papaioannou 2020), with consequential impacts on public good allocations and development (for example, Alesina et al. 2016; Baldwin and Huber 2010). This literature often assumes that stronger ethnic divisions will cause voters to have a stronger taste for redistribution targeted only to certain groups, together with support for clientelism and patronage. This point is developed in Easterly (2001), and then summarized by Besley and Ghatak (2006, p. 292): "[I]f externalities are limited to within ethnic groups, then there will be less total demand for public goods that benefits all groups such as roads and education." In the case of geographical

segregation, whereby each ethnic group occupies a different region, there would then be less demand for interregional travel or trade in fractionalized societies, suggesting less provision of countrywide public goods. A similar argument would imply less support for nationwide public education in the case of fractionalization in culture or language (Alesina et al. 2003).

While these mechanisms do have consequences on development in sub-Saharan Africa today, as discussed above, they struggle to explain the variation in outcomes within regions with the same historical ethnicities. That is, other factors beyond historical political institutions and ethnicity drive political behavior. For example, it seems that clientelistic appeals to co-ethnic citizens may depend on a variety of political, demographic, and urban factors.[1] Historically, and even today, many of the wealthiest areas and metropoles in Africa and beyond were also the most diverse ones due to the interactions from international trade and the resulting gains, suggesting the non-determinism of this factor. Indeed, the effects of ethnic conflict can be shaped over time, by public policy like fines or by social sanctions (Miguel and Gugerty 2005). Finally, it seems likely that conflicts across ethnic groups in power may be addressed when agents have a rich enough set of strategies. For example, looking at data from 15 African countries that together account for almost half of the continent's population, Francois et al. (2015) find that even autocratic leaders negotiate broader coalitions with other ethnicities in exchange for stability and rents.

Hence, we consider ethnic fractionalization as a common constraint on achieving optimal infrastructure or other public investments, but current policies can be effective even in the presence of ethnic divisions. For a concrete example, past public investments do have persistent effects over time: it is likely that roads built today will be concentrated in similar areas as roads built during colonial times. However, current policies may also help to correct distortions in government choices given existing infrastructure. In the case of Kenya, Burgess et al. (2015) find that during times of dictatorship, politicians preferred to spend on road-building in certain locations based on ethnicity. However, in periods of democracy, such distortions were not present. This occurred in a setting with the same colonial background and institutions, suggesting a positive role for policies on welfare consequences.

Another example is in international trade, where the lack of trade opportunities due to a variety of frictions may play a more significant role in welfare losses and lack of economic growth in sub-Saharan Africa than historical factors. Startz (2018) shows that significant welfare losses are incurred due to frictions in matching between intermediaries in trade in Nigeria. For example, a Nigerian sales representative may rely on production by intermediaries in China, but finding the appropriate matching producer in the market and enforcing production when effort

[1] Wantchekon (2003) offers evidence from Benin on the role of demographic and political factors, including local versus national appeals of such promises, while Ichino and Nathan (2013) provide evidence on rural-urban divides in ethnic appeal in Ghana, including the role of geographic factors and local context.

is unobservable may be very costly. Hence, decreasing information frictions through technology or by increasing the possibility of in-person visits through removing visa restrictions for business travellers may increase welfare gains by almost 15 percent. Relatedly, Bhandari (2021) conducts an experiment in Senegal and shows that the asymmetry in contract enforcement induces a breakdown in deals between buyers and sellers, particularly when one party knows the other is politically connected. In Togo and Benin, Blimpo (2015) shows how a lack of contract between owners and drivers of moto-taxis (and a simultaneous overreliance on kinship ties) can lead to excessive depreciation. Such results are in line with the long theoretical literature on moral hazard and trust, which indeed could be due to different ethnic or cultural backgrounds. Nevertheless, these newer results leave scope for policy interventions that could remove such spatial and contract enforcement frictions.

The Importance of Political Distortions: Political Connections, Patronage, and State Capture

We suggest an attention shift in the literature toward the role of distortions by the political process in understanding economic performance and development paths in sub-Saharan Africa. Political distortions are defined by a situation in which a special interest group can direct economic development towards its own exclusive ends rather than towards increasing general welfare. This may include provisions of public goods and redistribution that are chosen due to narrow political interests rather than for the welfare of the larger public. Political distortions affect growth by affecting resource allocation and economic decision-making: they often imply misallocation of public funds, inefficient public investments, distorting taxation, and altered market structure that arise through policies benefiting some firms instead of others. For example, some companies obtain concessions, such as beneficial tax breaks, custom duties, or regulation, through their connections to the bureaucracy and politicians. This can explain why regions with the same colonial or pre-colonial history may have very different growth trajectories. If subsidizing education or building infrastructure can significantly alter productivity, then the way such decisions are made, and whether they are efficient or represent citizens' preferences may be a powerful determinant of outcomes.

The term "distortion" alludes to the negative influence on welfare of these political decisions. However, quantifying their role is only possible relative to some established (efficient) benchmark. A theoretical framework that is sufficiently rich to unify all mechanisms by which political processes affect economic and growth outcomes does not yet exist. To illustrate, this section considers several examples that model how a particular political distortion can lead to a suboptimal economic outcome and how they relate to existing empirical evidence. These models are discussed side-by-side with empirical evidence. The former allows us to interpret the latter as distortions: they are deviations from a benchmark and negatively affect citizen welfare. While an exhaustive literature review on political distortions is

beyond the scope of this text (for discussion, see Martinez-Bravo and Wantchekon 2021) and empirical work in sub-Saharan Africa is still lacking, the following results establish empirical regularities that are targets of policy analysis and interpretable through the lens of established economic models with direct relevance to the region we study.

Political connections refer to a situation in which a firm has privileged access to a politician and may use that connection for its benefit. From the theoretical point of view, political connections may reduce costs and/or allow firms to enter new markets and become more profitable. In Akcigit et al. (2018), an incumbent firm can entrench itself by investing in political connections, which lowers its cost of entry (for example, by reducing red tape or through tax concessions). Those connections reduce the entry of innovating firms, because there is a connected incumbent already in the market. Hence, industries with more political connections are less innovative and see less productivity growth. Such connections have been shown to benefit firms through increased profits and stock prices: for example, Fisman (2001) shows that the stock market performance of firms with connections to the Suharto regime in Indonesia during the early 1990s suffered with negative reports about his health status. More recently, firms have been shown to gain from political connections in a variety of different contexts: for example, easier access to import licenses in Indonesia (Mobarak and Purbasari 2006); preferential lending in a dataset of 90,000 firms in Pakistan (Khwaja and Mian 2005); increased loans from state banks for firms in Brazil that make influence-building campaign contributions (Claessens et al. 2008); increased probability of bailouts for the Faccio et al. (2006) data on political connectedness across 47 countries; regulatory capture and favorable regulatory policies (as reviewed by Dal Bó, 2006); and more procurement contracts using data on political contributions by firms in Lithuania (Baltrunaite 2020). A substantial body of evidence, of which this is just a sampling, suggests that distortions from political connections can have significant effects on the allocation of resources.

One salient instance in Africa is the "godfatherism" phenomena documented by Joseph (1987) and Omotola (2007) in Nigeria. It is defined by a widespread and elaborate system of contractual arrangements, involving politicians, electoral brokers, and firms, in which the latter provide campaign funds for politicians in exchange for cabinet positions, influence over the bureaucracy, or market gains, such as increased market power in key sectors or procurement reservations.

Variation in contract enforcement, a main role for government emphasized in the political economy literature, is also a political distortion that may drive differential development. For instance, Goldstein and Udry (2008) find that agricultural producers in Ghana that are subject to stricter enforcement of property rights produce and invest more. In the macroeconomic framework of Aguiar and Amador (2011), the political friction is a party's inability to commit due to not knowing whether the party will be in power in the future. As a result, the incumbent government overvalues present consumption. The authors show that the speed of convergence and the levels of steady-state income in this economy

depend on the extent of political disagreement: with more disagreement, there is slower convergence. This environment embodies the notion that the availability of contract enforcement mechanisms can induce higher steady state income and higher economic growth, even in a democracy with political turnover. This finding is confirmed in other settings, such as in Senegal (Bhandari 2021, described above) and by Boehm and Oberfeld (2020) in a study of manufacturing plants in India. In the latter, weak contract enforcement increases a hold-up problem for plants requiring customized inputs, which in turn implies that both the intensive and extensive margins of input use are distorted. The aggregate welfare implications of such problems appear quantitatively meaningful.

Patronage refers to the discretionary appointment of individuals to public office or government positions.[2] The close connections of firms and private interests to state hiring practices may affect who gets hired, the quality of the bureaucracy, and which policies get implemented. These connections have been extensively verified in the literature. For example, Xu (2018) shows that connections measured by family ties and shared attendance of elite academic institutions induced higher salaries and worse public finance performance for connected bureaucrats in the British imperial civil service, but this effect disappeared after a ban on discretionary appointments. Often, such bureaucrats may receive future career benefits through "revolving doors" when they exit government for the private sector, as Blanes i Vidal et al. (2012) show in the case of lobbyists who formerly worked as staff for US Senators. To the best of our knowledge, the literature has not statistically shown the prevalence of this behavior in sub-Saharan Africa. Nevertheless, even in countries like Ghana and Nigeria, there are few laws, if any, guaranteeing merit-based appointments. Most of them are discretionary, as we discuss in the next section.

Another political distortion is the *targeted distribution of public goods*. For instance, politicians may prefer to spend more on infrastructure in preferred locations—say, for political loyalty or ethnic relationships—than would be optimal in a first-best world, as in the case (alluded to earlier) of road-building in Kenya (Burgess et al. 2015). This was also the case with the extensive public expenditure on roads, utilities, and white elephants in Côte d'Ivoire's Yamoussoukro, chosen to be a capital because it was the birthplace of President Houphoët-Boigny.

We can interpret such expenditures as distortionary within the early set-up of Barro (1990), which provides a clear and explicit role for politicians in the provision of public services. The author constructs a growth model that includes public services as a productive input for private producers. The model includes many producers (firms), with each firm choosing its own capital given government spending, although production depends on both private capital and public services. In the case of public services that can become congested and are excludable to some extent, such as highways, water and sewer systems, police services, and so on,

[2] While this is typically viewed as a negative phenomenon, it could be theoretically beneficial if the nominating party has private information that is used appropriately in selecting employees, as seen in Dal Bó et al. (2021) in Paraguay.

a firm's decision to expand its own capital increases the risk of congestion in public services available for other producers. This externality implies that governments have an explicit role in driving efficient and higher growth through providing public infrastructure, but distortions will arise if such taxation and public investment choices are not optimal, such as rents to certain groups in the population. Jedwab and Moradi (2016) show large effects of railroads in Ghana, as does Okoye et al. (2019) in Nigeria. In Ethiopia, Gebresillasse (2018) finds that the extension of rural roads generates increased productivity of 6 percent, with important mechanisms being increased information transmission (like take up of advice and modern inputs), as well changes in labor and crop allocation.

Finally, *regulation and technological adoption* are also areas subject to political distortions. The growth model of Acemoglu and Robinson (2006) centers around the role of an incumbent politician who may choose whether to pursue economic innovation or to block it. Technological innovation may change political competition by affecting the rents incumbents receive from being in office. The population chooses whether to keep or replace the incumbent after observing the latter's choice. The authors show that, in one equilibrium, political elites may block technological innovation for fear the innovation will reduce their term in power. By blocking technological innovation, these politicians become further entrenched and less productive firms are kept in the market, implying welfare losses.[3]

Many other examples of political distortions exist. Perhaps the most salient recent example of state capture—the control of state agencies and bureaucrats by special interest groups—in Africa involves the Zuma presidency and its relationship with the Gupta family's business empire in South Africa. After Zuma's election in 2009, the Gupta family used its friendly ties to the presidency to run the government as a "private piggy-bank" (as reported in Gevisser 2019). This involved the granting of numerous government contracts to Gupta-affiliated firms (Alence and Pitcher 2019), using influence with the president and his appointed (manipulable) bureaucrats. It is even alleged that the Zuma government outsourced picking the Minister of Finance and members of the Treasury to the Guptas—an especially obvious example of what Canen, Ch, and Wantchekon (2021) term "direct" capture.

Other times, the mechanisms by which political distortions occur may be hard to observe or done through intermediaries. To capture the latter, Canen, Ch, and Wantchekon (2021) define the notion of *indirect capture*, which occurs when firms use intermediaries (that is, non-bureaucrats) to influence policy implementation rather than when there are clear and direct interactions between firms and the

[3]This insight can also be viewed an example of the importance of property rights. In the set-up of Aghion, Alesina, and Trebbi (2007), improved protection of property rights implies a lower ability of incumbents to block new entrants and innovation. They further assume that innovation matters more for countries closer to the technological frontier. Hence, technologically advanced sectors benefit the most from the protection of property rights as it stimulates further competition and entry, inducing a more positive correlation between democracy and growth with proximity to the technological frontier. This correlation is validated in cross-country comparisons. Although this model implies the potential importance of political distortions, such distortions are not explicitly discussed in the paper.

bureaucrats in charge of policy. An example of indirect capture is lobbying. The United States is one of the few settings where data on lobbying expenditures is available: in that context, there appears to be sizable premia for lobbyist connections regardless of field expertise (Bertrand et al. 2014). Furthermore, Huneeus and Kim (2021) find that lobbying affects firm size and entry inducing large misallocation effects (in that study, 10 percent of the US economy). We refer the reader to Martinez-Bravo and Wantchekon (2021) for further examples.

In short, political economy incentives can meaningfully distort economic outcomes and welfare, which is the heart of the idea of political distortions. Theoretical models show how political distortions lead to suboptimal equilibria with effects on economic growth and consumption, debt profiles and fiscal sustainability, and endogenous technological innovation and the resulting firm productivity. These kinds of models structure empirical analysis by framing the empirical results relative to such benchmarks.

To identify effective policy solutions, we must be clear on the incentives and political factors behind the distortions at play. In our view, the framework of political distortions suggests we should think of the long-run structural conditions in Africa as constraints that affect regional inequalities and present-day choices by citizens, firms, and political elites. For example, a certain distribution of goods will arise through the investments colonial governments made or did not make, according to their state-making goals or historical accidents (for example, Huillery 2009; Okoye et al. 2019; Wantchekon et al. 2015). In terms of public goods, colonial administrations created schools and built roads in some places but not in others in order to maximize revenues to be extracted from the colonies (Ricart-Huguet 2021). The political distortions framework above allows us to think about which current institutional and policy reforms can help to correct these misallocations and the inequalities inherited from the past, along with achieving other development goals.[4]

Policy and Institutional Reforms to Reduce Political Distortions

Certain policies hold a promise of reducing political distortions and state capture given the framework and evidence above. The hope is that if such policies are enacted, they will reduce the effects of market distortions and offset some effects of past historical and structural events. We begin with three reforms that involve broad institutional change: campaign finance rules, regulatory and bureaucratic reform, and constraints imposed by multilateral reform (in this case, the new African Continental Free Trade Agreement). However, these kinds of systemic

[4]One could suggest that such political distortions may result from the historical factors outlined earlier. However, we are unaware of systematic evidence to that regard. In fact, the types of political distortions outlined above often vary across sectors or even across firms in a sector, all subject to the same historical and ethnic considerations.

reforms may be politically unpalatable or difficult to implement. Thus, we then turn to alternative policies that may be implemented in the short run and even at the local level and that may still curb state capture, mostly through revealing information to voters about previously unobserved/hard-to-observe behavior between firms and politicians. First suggestions include: audits of political jurisdictions, increasing the availability of information to citizens (say, via town hall meetings or debates), and using technology to improve transparency within the government.

Campaign Finance Rules

Political connections are often reflected in campaign donations: firms that have some preference or relationship with a politician (or wish to have one) offer funds for campaigns in exchange for beneficial policies or access to politicians if elected. It is then natural to consider that policies that limit such contributions would limit undue influence of political connections on policy. By curbing campaign contributions, firms would lose a mechanism by which it can influence policy, thereby restricting political connections and subsequent market distortions. Indeed, the positive welfare effects of such policies have been reinforced in the models of Prat (2002), Coate (2004), and Ashworth (2006).

There are many ways to implement restrictions on campaign contributions, from public finance of campaigns, limiting the amount of donations, or banning influential economic actors like corporations. Baltrunaite (2020) showed that a 2012 ban on corporate contributions in Lithuania decreased the winning rate of government contracts by "connected" firms (that is, those previously contributing to campaigns) relative to non-contributing firms. While pre-ban there was a 5-percentage point gap in the probability of winning a contract between such connected and non-connected firms, even after controlling for observables, this gap disappeared within a year of the policy. The ban also decreased the price the government paid in those procurement auctions. In sub-Saharan Africa, there is room to implement reforms of this type: Table 1 shows that less than 15 percent of African countries have banned corporate donations to candidates (compared to 26 percent elsewhere), and less than one-third restrict the amount candidates can spend, compared to 56 percent beyond the continent. In fact, even less restrictive policies are not as widely implemented: almost two-thirds of countries in Africa do not ban donations to candidates from corporations with government contracts. By comparison, vote-buying is already banned in most of the continent (almost 90 percent of countries), even though its financial cost (and arguably, its influence on policy) is much smaller than those from campaign contributions.[5]

[5] In the context of Africa, vote-buying has been widely studied and is often considered pervasive in Africa (for a review, see Vicente and Wantchekon 2009). However, we believe this political distortion is overemphasized. Vote-buying is already banned in most of the continent, as shown in Table 1. Moreover, the amount of money spent by firms in campaign donations is often magnitudes larger than those in clientelism and vote-buying attempts. Indeed, in Kenya, politicians are estimated to spend $2 for each vote, implying a total of $100 million if every voter received such a transfer (IDEA and Falguera et al.

Table 1
Campaign Finance Rules in Africa and Beyond

	Africa			*Countries beyond Africa*		
	Yes	*No*	*No data*	*Yes*	*No*	*No data*
Ban on corporate donations to candidates	14%	78%	8%	26%	68%	5%
Limits on the amount a candidate can spend?	31%	61%	8%	56%	39%	5%
Ban on donations from corporations with government contracts to candidates	12%	65%	23%	37%	54%	9%
Ban on donors to political parties/candidates participating in public tender/procurement processes	0%	37%	63%	5%	64%	30%
Ban on vote buying	86%	4%	10%	95%	3%	2%

Source: IDEA Political Finance Database
Note: The table shows the proportion of countries in Africa and those outside of Africa that have a certain policy on elections and political finances in place ("Yes") or that do not have them ("No"). The policies are the rows. Totals do not always sum to 100 percent because of rounding.

Of course, rules about campaign finance can often be circumvented, especially in low state-capacity settings. In sub-Saharan Africa, campaign contributions by firms are difficult to enforce and observe. Would such rules still be effective if firms could donate in alternative forms of campaign contributions that are unobserved or left unrestricted? Or if firms simultaneously donate to many candidates? Empirical findings, such as those in Avis et al. (2020), suggest that restricting campaign spending is still effective in those cases. They show that such restrictions in mayoral races in Brazil induced higher levels of political competition and more entry by less wealthy and established candidates. As a result, changes to political finance rules may be likely to constrain state capture in the African context when faced with similar incentives, even though implementation and enforcement is not perfect.

Insulating Bureaucrats and Regulators

Insulating bureaucrats and regulators from political influence is another possible reform to reduce political distortions. For example, one possibility is the use of elections to elect regulators directly, as argued in Besley and Coate (2003) using evidence from US states, thus giving regulators some insulation from separate political pressures to implement policies not supported by voters. It is also likely that results on bureaucratic reform from past studies could further inform such decisions. In a study mentioned earlier, Xu (2018) showed that the banning of discretionary appointments in the United Kingdom in 1930 yielded an improvement in public finance outcomes (with increases in tax revenues and public investments),

2014). However, each party in Kenya may already spend such an amount in a campaign. Put together, this suggests an increasingly important role played by policymaking and firms' interactions with policymakers.

and an equalization of salaries and promotions across previously connected and non-connected bureaucrats. Similar results have been obtained by Ujhelyi (2014) in the context of state-level reforms in the United States guaranteeing tenure to bureaucrats.

To the best of our knowledge, there is less systematic evidence of this in African countries, as policies protecting bureaucrats still appear to be widely below desired levels (as reported in the context of electrical power by Rodriguez 2017), although there has been an increasing focus on technical qualifications. While Nigeria, Ghana, and Kenya have civil service exams for entry and promotions, most bureaucratic appointments in Ghana are discretionary, for instance (Brierley 2021). And while changes in Kenya's 2010 constitution gave more control of cabinet and senior positions to the legislature, it is unclear whether that is necessarily welfare-improving (Opalo 2019). This seems to be a promising avenue for future research.

Multi-Country Initiatives: The African Continental Free Trade Area

Multi-country and continent-wide agreements are another source of policy change that can stimulate economic growth through limiting political distortions. In Africa, a salient example is the just-established African Continental Free Trade Area, whose goal is to stimulate business competition by reducing the costs of conducting business in the continent. It was founded in 2018, with trade under its terms starting in 2021. Participation in the new free trade area entails reforms such as the simplification of customs procedures and the reduction of contractual and costly licensing. These steps may reduce the uncertainty around contract enforcement and in search, as the common set of regulations implemented beyond the country's political system leaves less scope for political connections to play a role, reducing uncertainty in exchanges. For example, simplification of import duties restricts the possibility that they may be differentially given to politically connected firms.

Our argument suggests that gains from the African Continental Free Trade Area may come in two ways: the conventional gains from trade liberalization, and benefits from a reduction in political distortions and rent seeking. The standard arguments suggest that trade liberalization increases productivity within liberalized regions by reshuffling resources to more productive plants/firms (for an example from Chile, see Pavcnik 2002), through the removal of inefficient quotas (for an example from Chinese exporters, see Khandelwal, Schott, and Wei 2013), or through the decrease in markups due to increased competition (for an example from Taiwanese producers, see Edmond, Midrigan, and Xu 2015). Such reforms still have potential trade-offs, like an increase in income inequality (overviewed in Pavcnik 2017), or countervailing effects on misallocation (as modelled for Chinese manufacturing firms in Bai et al. 2019). However, the previous mechanisms suggest that the African Continental Free Trade Agreement would provide less opportunity for political discretion and hence, less returns from state capture for connected firms in participating countries. As a result, it holds the potential to spur valuable

institutional reforms even before conventional international trade effects are even considered.

Audits

Audits of political jurisdictions can increase accountability. This mechanism has been most prominently studied in Brazil due to its randomized implementation at the municipality level (for example, Ferraz and Finan 2008, 2011; Avis et al. 2018). In the Brazilian context, an independent oversight body at the federal level randomly audits municipalities, investigating their accounts and spending patterns for evidence of fraud or corruption. These audits appear to be effective at decreasing future corruption by making information on the politicians' behavior widely available, which in turn allows voters, lawsuits, and official institutions to punish the guilty parties. For example, Avis et al. (2018) find that audits may reduce future acts of corruption by approximately 10 percent. However, our own follow-up work extending the results of Avis et al. (2018), and available by request from the authors, finds that the effects are strongest for municipalities that are already below the median in the number of corruption acts. Hence, audits are effective at decreasing acts of corruption and malfeasance particularly in settings where such politicians are likely responsive, but they are unlikely to trigger the same extent of welfare changes without combining them with fundamental institutional reforms.

We are not aware of extensive evidence of audits being applied or evaluated within sub-Saharan Africa. However, the economic development and structure of some municipalities in Brazil seem comparable, for example, to regions of South Africa and elsewhere on the continent (Huchzermeyer 2002), and recent evidence from South Africa (Berliner and Wehner forthcoming) together with common political and economic incentives (discussed earlier) suggest they would be effective there as well.

Increasing Available Information

Increasing the quantity and quality of information available to citizens can happen in other ways, as well, and several African countries have taken steps along these lines. In Sierra Leone, Bidwell et al. (2020) found that screening public debates led to increasing information to voters—and the effect is significant enough to influence voting behavior, campaigns' responses, and ultimately policy outcomes. In Benin, Fujiwara and Wantchekon (2013) found that town hall meetings focused on programmatic debates also increase political competition (reducing the incumbency advantages in political strongholds) and reduce clientelism without affecting turnout.

Both town hall meetings and public debate screenings are low-cost policies that can increase information availability to voters with policy consequences. Such increases in information might not only affect the outcome of the election, but may also encourage the entry of different candidates. As discussed in Canen et al. (2021), the entry of new candidates with different preferences for consumer

welfare relative to firm behavior appears to be one of the few ways to decrease significantly the effects of political incentives on economic outcomes and voter welfare.

Technology and Transparency

Finally, technology can be used to improve transparency in government and, in this way, to reduce the potential for state capture through increasing available information to citizens and decreasing information frictions. This approach has received renewed attention, for example, in recent empirical studies on the use of biometric identification cards in the provision of public services in India (Muralidharan et al. 2016; Banerjee et al. 2018). In fact, Banerjee et al. (2018) find that the introduction of biometric identification cards in India led to beneficiaries receiving 26 percent more subsidies, with no effect on ineligible citizens. Another example is in its use for tax collection: see Okunobe and Santoro (2021) for an overview of such applications for Africa. For instance, technology can help identify the tax base, help enforcement of tax collection, and decrease the cost of compliance. Technology can also be used for creating satellite imagery maps that can improve the coverage of tax collection, as Ali et al. (2018) demonstrate for the city of Kigali in Rwanda. Within the context of our theoretical framework, such technology would reduce the scope for political influence and corruption (Kochanova et al. 2020). There has been increased adoption of technology in government policies, including the use of the well-known M-PESA mobile technology that has been implemented in revenue collection in Kenya (Safaricom Media Release 2019), and similar mobile technology is used in the Growth Enhancement Scheme providing subsidies to farmers in Nigeria. However, we are unaware of systematic evidence evaluating such policies in economics. We encourage further study in this area as well.

Conclusion

In this paper, we studied how political incentives may distort economic growth and development, with a focus in Africa and, in particular, sub-Saharan Africa. In doing so, we used the general concept of political distortions to provide a unified framework through which extant empirical evidence, in Africa and elsewhere, could be interpreted as inhibiting economic development. While this connection has been well-known in other regions, such as the United States or Latin America, this has been less emphasized in the African context relative to long-run structural explanations. As sub-Saharan African countries converge in development to others, the relationship between governments and firms also become more like that in other regions, as well. Hence, a unified theoretical framework based on political distortions allows us to transport the experiences, results, and policy considerations from other countries to the African context.

In our view, historical factors may act as constraints on political distortions, making it more likely that political distortions may be observed in some places.

However, historical constraints can be overcome. Recent papers and evidence that leverage variation *within* African countries clearly show that African economic development responds to policy interventions. Despite strife and conflict, countries like Ethiopia, Rwanda, and Côte d'Ivoire, to name only a few, have all maintained very successful economic trajectories over the last decade. Kenya has generated sustained growth because of newly founded institutions that have, in turn, encouraged a booming information technology sector, while Ghana has been solidly democratic for decades. This has been the case despite colonial pasts that would have suggested far less optimistic outcomes. Historical analysis is essential, but to understand how historical factors affect the present, it must explicitly tie in the strategies and policies chosen by current governments, firms, and civil society organizations. Thus, we propose that political economy research in Africa should focus on laws, markets, and social interactions between current key economic and political actors.

■ *We thank Gregory Amusu, Julio Arce Solis, Marco Castradori, Weonhyeok Chung, Kurtis Lockhart, Matthew Merrigan, and Shourya Sen for excellent research assistance. We thank Erik Hurst, Dozie Okoye, Ken Opalo, Nina Pavcnik, Heidi Williams and Timothy Taylor for very helpful comments.*

References

Acemoglu, Daron, and James A. Robinson. 2006. "Economic Backwardness in Political Perspective." *American Political Science Review* 100 (1): 115–31.

Acemoglu, Daron, and James A. Robinson. 2010. "Why Is Africa Poor?" *Economic History of Developing Regions* 25 (1): 21–50.

Acemoglu, Daron, Simon Johnson, and James A. Robinson. 2001. "The Colonial Origins of Comparative Development: An Empirical Investigation." *American Economic Review* 91 (5): 1369–1401.

Aghion, Philippe, Alberto Alesina, and Francesco Trebbi. 2007. "Democracy, Technology, and Growth." NBER Working Paper 13180.

Aguiar, Mark, and Manuel Amador. 2011. "Growth in the Shadow of Expropriation." *Quarterly Journal of Economics* 126 (2): 651–97.

Akcigit, Ufuk, Salomé Baslandze, and Francesca Lotti. 2018. "Connecting to Power: Political Connections, Innovation, and Firm Dynamics. NBER Working Paper 25136.

Alence, Rod, and Anne Pitcher. 2019. "Resisting State Capture in South Africa." *Journal of Democracy* 30 (4): 5–19.

Alesina, Alberto, Arnaud Devleeschauwer, William Easterly, Sergio Kurlat, and Romain Wacziarg. 2003. "Fractionalization." *Journal of Economic Growth* 8 (2): 155–94.

Alesina, Alberto, Stelios Michalopoulos, and Elias Papaioannou. 2016. "Ethnic Inequality." *Journal of Political Economy* 124 (2): 428–88.

Ali, Daniel Ayalew, Klaus Deininger, and Michael Wild. 2018. "Using Satellite Imagery to Revolutionize Creation of Tax Maps and Local Revenue Collection." World Bank Policy Research Paper 8437.

Ashworth, Scott. 2006. "Campaign Finance and Voter Welfare with Entrenched Incumbents." *American*

Political Science Review 100 (1): 55–68.

Avis, Eric, Claudio Ferraz, and Frederico Finan. 2018. "Do Government Audits Reduce Corruption? Estimating the Impacts of Exposing Corrupt Politicians." *Journal of Political Economy* 126 (5): 1912–64.

Avis, Eric, Claudio Ferraz, Frederico Finan, and Carlos Varjão. 2020. "Money and Politics: The Effects of Campaign Spending Limits on Political Entry and Competition and Incumbency Advantage." https://drive.google.com/file/d/1FbhUnhSFYJ7nsXIry-FzXxIpebj3JqOg/view.

Bai, Yan, Keyu Jin, and Dan Lu. 2019. "Misallocation under Trade Liberalization." NBER Working Paper 26188.

Baldwin, Kate, and John D. Huber. 2010. "Economic versus Cultural Differences: Forms of Ethnic Diversity and Public Goods Provision." *American Political Science Review* 104 (4): 644–62.

Baltrunaite, Audinga. 2020. "Political Contributions and Public Procurement: Evidence from Lithuania." *Journal of the European Economic Association* 18 (2): 541–82.

Banerjee, Abhijit, Rema Hanna, Jordan Kyle, Benjamin A. Olken, and Sudarno Sumarto. 2018. "Tangible Information and Citizen Empowerment: Identification Cards and Food Subsidy Programs in Indonesia." *Journal of Political Economy* 126 (2): 451–91.

Barro, Robert J. 1990. "Government Spending in a Simple Model of Endogeneous Growth." *Journal of political economy* 98 (5, Part 2): S103–125.

Berliner, Daniel, and Joachim Wehner. Forthcoming. "Audits for Accountability: Evidence from Municipal By-Elections in South Africa." *Journal of Politics*.

Bertrand, Marianne, Matilde Bombardini, and Francesco Trebbi. 2014. "Is It Whom You Know or What You Know? An Empirical Assessment of the Lobbying Process." *American Economic Review* 104 (12): 3885–3920.

Besley, Timothy, and Stephen Coate. 2003. "Elected versus Appointed Regulators: Theory and Evidence." *Journal of the European Economic Association* 1 (5): 1176–1206.

Besley, Timothy, and Maitreesh Ghatak. 2006. "Public Goods and Economic Development." In *Understanding Poverty*, edited by Abhijit Vinayak Banerjee, Roland Bénabou, and Dilip Mookherjee, 285–303. Oxford: Oxford University Press.

Bhandari, Abhit. 2021. "Political Determinants of Economic Exchange: Evidence from a Business Experiment in Senegal." *American Journal of Political Science* https://doi.org/10.1111/ajps.12593.

Bidwell, Kelly, Katherine Casey, and Rachel Glennerster. 2020. "Debates: Voting and Expenditure Responses to Political Communication." *Journal of Political Economy* 128 (8): 2880–2924.

Blanes i Vidal, Jordi, Mirko Draca, and Christian Fons-Rosen. 2012. "Revolving Door Lobbyists." *American Economic Review* 102 (7): 3731–48.

Blimpo, Moussa P. 2015. "Kinship, Trust and Moral Hazard in the Motorcycle-Taxi Market in Togo and Benin." *Journal of African Economies* 24 (2): 173–92.

Boehm, Johannes, and Ezra Oberfield. 2020. "Misallocation in the Market for Inputs: Enforcement and the Organization of Production." *Quarterly Journal of Economics* 135 (4): 2007–2058.

Brierley, Sarah. 2021. "Combining Patronage and Merit in Public Sector Recruitment." *Journal of Politics* 83 (1): 182–97.

Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró i Miquel. 2015. "The Value of Democracy: Evidence from Road Building in Kenya." *American Economic Review* 105 (6): 1817–51.

Canen, Nathan, Rafael Ch, and Leonard Wantchekon. 2021. "Political Uncertainty and the Forms of State Capture." https://dx.doi.org/10.2139/ssrn.3567523.

Claessens, Stijn, Erik Feijen, and Luc Laeven. 2008. "Political Connections and Preferential Access to Finance: The Role of Campaign Contributions." *Journal of Financial Economics* 88 (3): 554–80.

Coate, Stephen. 2004. "Pareto-Improving Campaign Finance Policy." *American Economic Review* 94 (3): 628–55.

Crawford, Gordon, and Gabrielle Lynch, eds., 2012. *Democratization in Africa: Challenges and Prospects*.

Dal Bó, Ernesto. 2006. "Regulatory Capture: A Review." *Oxford Review of Economic Policy* 22 (2): 203–25. London: Routledge.

Dal Bó, Ernesto, Frederico Finan, Nicholas Y. Li, and Laura Schechter. 2021. "Information Technology and Government Decentralization: Experimental Evidence from Paraguay." *Econometrica* 89 (2): 677–701.

Drèze, Jean, and Amartya Sen. 2013. *An Uncertain Glory: India and Its Contradictions*. Princeton: Princeton University Press.

Easterly, William. 2001. "Can Institutions Resolve Ethnic Conflict." *Economic Development and Cultural*

Change 49 (4): 687–706.

Easterly, William, and Ross Levine. 1997. "Africa's Growth Tragedy: Policies and Ethnic Divisions." *Quarterly Journal of Economics* 112 (4): 1203–50.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu. 2015. "Competition, Markups, and the Gains from International Trade." *American Economic Review* 105 (10): 3183–3221.

Faccio, Mara, Ronald W. Masulis, and John J. McConnell. 2006. "Political Connections and Corporate Bailouts." *Journal of Finance* 61 (6): 2597–2635.

Falguera, Elin, Samuel Jones, and Magnus Ohman. 2014. *Funding of Political Parties and Election Campaigns: A Handbook on Political Finance.* https://www.idea.int/sites/default/files/publications/funding-of-political-parties-and-election-campaigns.pdf.

Fawehinmi, Feyi. 2017. "Africa's Richest Man Has a Built-in Advantage with Nigeria's Government." *Quartz*, October 11. https://qz.com/africa/1098137/africas-richest-man-has-a-built-in-advantage-with-nigerias-government/.

Ferraz, Claudio, and Frederico Finan. 2008. "Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes." *Quarterly Journal of Economics* 123 (2): 703–45.

Ferraz, Claudio, and Frederico Finan. 2011. "Electoral Accountability and Corruption: Evidence from the Audits of Local Governments." *American Economic Review* 101 (4): 1274–1311.

Fisman, Raymond. 2001. "Estimating the Value of Political Connections." *American Economic Review* 91 (4): 1095–1102.

Francois, Patrick, Ilia Rainer, and Francesco Trebbi. 2015. "How Is Power Shared in Africa?" *Econometrica* 83 (2): 465–503.

Fujiwara, Thomas, and Leonard Wantchekon. 2013. "Can Informed Public Deliberation Overcome Clientelism? Experimental Evidence from Benin.» *American Economic Journal: Applied Economics* 5 (4): 241–55.

Garcia, Luc. 1971. «L'organisation de l'instruction publique au Dahomey, 1894–1920.» *Cahiers d'Études Africaines* 11 (41): 59–100.

Gebresilasse, Mesay. 2018. "Rural Roads, Agricultural Extension, and Productivity." https://intranet.weatherhead.case.edu/document-upload/docs/1664.pdf.

Gevisser, Mark. 2019. "'State Capture': The Corruption Investigation that Has Shaken South Africa." *The Guardian*. July 11, https://www.theguardian.com/news/2019/jul/11/state-capture-corruption-investigation-that-has-shaken-south-africa.

Goldstein, Markus, and Christopher Udry. 2008. "The Profits of Power: Land Rights and Agricultural Investment in Ghana." *Journal of Political Economy* 116 (6): 981–1022.

Hsieh, Chang-Tai, and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics* 124 (4): 1403–48.

Hu, Zuliu F., and Mohsin S. Khan. 1997. "Why Is China Growing So Fast?" *Staff Papers* 44 (1): 103–31.

Huchzermeyer, Marie. 2002. "Informal settlements: Production and Intervention in Twentieth-Century Brazil and South Africa." *Latin American Perspectives* 29 (1): 83–105.

Huillery, Elise. 2009. "History Matters: The Long-Term Impact of Colonial Public Investments in French West Africa." *American Economic Journal: Applied Economics* 1 (2): 176–215.

Huneeus, Federico, and In Song Kim. 2021. "The Effects of Firms' Lobbying on Resource Misallocation." https://www.fedehuneeus.com/s/misallocation.pdf.

Huntington, S.P. 1991. "Democracy's Third Wave." *Journal of Democracy* 2 (2): 12–34.

Ichino, Nahomi, and Noah L. Nathan. 2013. "Crossing the Line: Local Ethnic Geography and Voting in Ghana." *American Political Science Review* 107 (2): 344–61.

IDEA, Political Finance Database. 2021. https://www.idea.int/data-tools/data/political-finance-database (accessed December 14, 2021).

Jedwab, Remi, and Alexander Moradi. 2016. "The Permanent Effects of Transportation Revolutions in Poor Countries: Evidence from Africa." *Review of Economics and Statistics* 98 (2): 268–84.

Joseph, Richard A. 1987. *Democracy and Prebendal Politics in Nigeria: The Rise and Fall of the Second Republic.* Cambridge: Cambridge University Press.

Khandelwal, Amit K., Peter K. Schott, and Shang-Jin Wei. 2013. "Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese Exporters." *American Economic Review* 103 (6): 2169–95.

Khwaja, Asim Ijaz, and Atif Mian. 2005. "Do Lenders Favor Politically Connected Firms? Rent Provision in an Emerging Financial Market." *Quarterly Journal of Economics* 120 (4): 1371–1411.

Kochanova, Anna, Zahid Hasnain, and Bradley Larson. 2020. "Does E-Government Improve Government

Capacity? Evidence from Tax Compliance Costs, Tax Revenue, and Public Procurement Competitiveness." *World Bank Economic Review* 34 (1): 101–20.

Kremer Michael, Jack Willis, and Yang You. 2021. "Converging to Convergence." *NBER Macroeconomic Annual* 36.

Martinez-Bravo, Monica, and Leonard Wantchekon. 2021. "Political Economy and Structural Transformation: Democracy, Regulation and Public Investment". https://www.cemfi.es/~martinez-bravo/mmb/Research_files/Inception%20Paper%20%20STEG%20PE%20MMB%20LW%202%202%2021%20.pdf.

Michalopoulos, Stelios, and Elias Papaioannou. 2013. "Pre-Colonial Ethnic Institutions and Contemporary African Development." *Econometrica* 81 (1): 113–52.

Michalopoulos, Stelios, and Elias Papaioannou. 2020. "Historical Legacies and African Development." *Journal of Economic Literature* 58 (1): 53–128.

Miguel, Edward, and Mary Kay Gugerty. 2005. "Ethnic Diversity, Social Sanctions, and Public Goods in Kenya." *Journal of Public Economics* 89 (11–12): 2325–68.

Mobarak, Ahmed Mushfiq, and Denni Puspa Purbasari. 2006. "Corrupt Protection for Sale to Firms: Evidence from Indonesia." Unpublished.

Mühlhahn, Klaus. 2019. *Making China Modern: From the Great Qing to Xi Jinping*. Cambridge, MA: Harvard University Press.

Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar. 2016. "Building State Capacity: Evidence from Biometric Smartcards in India." *American Economic Review* 106 (10): 2895–2929.

News Wires. 2021. "Bolloré Group Fined €12 Million in Togo Corruption Case." *France 24*, February 26. https://www.france24.com/en/europe/20210226-bollor%C3%A9-group-fined-%E2%82%AC12-million-in-african-corruption-case.

Nunn, Nathan, and Leonard Wantchekon. 2011. "The Slave Trade and the Origins of Mistrust in Africa." *American Economic Review* 101 (7): 3221–52.

Okoye, Dozie, Roland Pongou, and Tite Yokossi. 2019. "New Technology, Better Economy? The Heterogeneous Impact of Colonial Railroads in Nigeria." *Journal of Development Economics* 140: 320–54.

Okunogbe, Oyebola, and Fabrizio Santoro. 2021. "The Promise and Limitations of Information Technology for Tax Mobilization." World Bank Policy Research Working Paper 9848.

Omotola, J. Shola. 2007. "Godfathers and the 2007 Nigerian Elections." *Journal of African Elections* 6 (2): 147–8.

Opalo, Ken Ochieng'. 2019. *Legislative Development in Africa: Politics and Postcolonial Legacies*. Cambridge, UK: Cambridge University Press.

Patel, Dev, Justin Sandefur, and Arvind Subramanian. 2021. "The New Era of Unconditional Convergence." *Journal of Development Economics* 152: 102687.

Pavcnik, Nina. 2002. "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants." *Review of Economic Studies* 69 (1): 245–76.

Pavcnik, Nina. 2017. "The Impact of Trade on Inequality in Developing Countries." NBER Working Paper 23878.

Polity Annual Time Series, Center for Systemic Peace. 2021. https://www.systemicpeace.org/inscrdata.html (accessed June 9, 2021).

Prat, Andrea. 2002. "Campaign Advertising and Voter Welfare." *Review of Economic Studies* 69 (4): 999–1017.

Ricart-Huguet, Joan. 2021. "Colonial Education, Political Elites, and Regional Political Inequality in Africa." *Comparative Political Studies* 54 (14).

Rodriguez, Ignacio. 2017. "The Role of Independent Regulators in Bringing Power to Africa." *Tetra Tech*, June 12. https://www.tetratech.com/en/markets/international-development/news-and-impact/insights/the-role-of-independent-regulators-in-bringing-power-to-africa.

Safaricom Media Release 527. 2019. Kwale County Automates Revenue Collection with M-PESA. https://www.safaricom.co.ke/about/media-center/publications/press-releases/release/527. (accessed December 2021).

Startz, Meredith. 2018. "The Value of Face-to-Face: Search and Contracting Problems in Nigerian Trade." https://dx.doi.org/10.2139/ssrn.3096685.

Transparency International, Corruption Perception Index. 2020. Data File Retrieved from https://www.transparency.org/en/cpi/2020/index/ (accessed December 14, 2021).

Ujhelyi, Gergely. 2014. "Civil Service Rules and Policy Choices: Evidence from US State Governments." *American Economic Journal: Economic Policy* 6 (2): 338–80.

Vicente, Pedro C., and Leonard Wantchekon. 2009. "Clientelism and Vote Buying: Lessons from Field Experiments in African Elections." *Oxford Review of Economic Policy* 25 (2): 292–305.

Wantchekon, Leonard. 2003. "Clientelism and Voting Behavior: Evidence from a Field Experiment in Benin." *World Politics*: 399–422.

Wantchekon, Leonard, Marko Klašnja, and Natalija Novta. 2015. "Education and Human Capital Externalities: Evidence from Colonial Benin." *Quarterly Journal of Economics* 130 (2): 703–57.

World Bank Open Data. 2021. Data File Retrieved from https://data.worldbank.org/ (June 9, 2021).

Xu, Guo. 2018. "The Costs of Patronage: Evidence from the British Empire." *American Economic Review* 108 (11): 3170–98.

Journal of Economic Perspectives—Volume 36, Number 1—Winter 2022—Pages 125–150

The Price of Nails Since 1695: A Window into Economic Change

Daniel E. Sichel

Nails are a basic manufactured product whose form and quality have changed relatively little since the late 1600s, and nails produced then would be quite recognizable today (as would nails produced in ancient Rome). Yet the manufacturing process for nails has changed dramatically, with a shift from artisanal to factory production, a change in power source from hand to water to steam to electricity, and a shift in materials from iron to steel.

Coincident with those changes, the price of nails fell significantly. In this paper, I construct a price index for nails, relative to an index of overall consumer prices and based on data going back to 1695. The real price of nails fell by a factor of about 10 from the late 1700s to the middle of the 20th century, averaging a decline of about 1½ percent per year. These declines are paltry compared with more dramatic examples: for example, Nordhaus (2007) calculated that the real cost of computing dropped by a factor of at least 2 trillion times from 1850 to the early 2000s, while Nordhaus (1997) showed that the real cost of lighting fell by a factor of about 3400 times from 1800 to 1992. However, while these vivid examples provide important insights, they do not particularly shed light on the developments and changes that affect more pedestrian products. We need to look elsewhere for that, and one product that hits the nail on the head is, wait for it, nails. Indeed, these declines in the price of nails provide a useful entry point for deepening our understanding of changes in manufacturing processes, key sources of those changes, and the evolving

■ *Daniel E. Sichel is Professor of Economics, Wellesley College, Wellesley, Massachusetts. He is also a Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts. His email address is dsichel@wellesley.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.36.1.125.

role of nails in the US economy. Moreover, nails are ideal for this analysis because there has been relatively little change in the product itself (unlike the production of computing power or lighting), thereby greatly simplifying the task of adjusting for changes in quality over time.

Over the same span of 300 plus years during which these transformations occurred, the place of nails in the economy (and in popular accounts) also underwent a huge shift. In 1810 (the earliest year for which I could assemble necessary data), the use of nails in the US economy (measured as production plus imports minus exports) was 0.4 percent of nominal GDP, as shown in Figure 1. To put this share into perspective, in 2019 household purchases of personal computers and peripheral equipment amounted to roughly 0.3 percent of GDP and household purchases of air travel amounted to about 0.5 percent. That is, back in the 1700s and early 1800s, nails were about as important in the economy as computers or air travel purchased by consumers are today.

The high value of nails during earlier periods is reflected in the practice of recovering used nails. According to Temin (1964, p. 42), during the 1700s abandoned buildings were sometimes burned down to facilitate the recovery of nails. Even a century later, nails were still highly valued. The well-known novel *Little House on the Prairie* (Wilder 1935, p. 124) contains a description of attaching a roof to a log home on the frontier during the 1870s (after the price of nails had already fallen significantly from the late 1700s), which highlights the value placed on nails:

> Now Pa carefully took the nails one by one from his mouth, and with ringing blows of the hammer he drove them into the slab. It was much quicker than drilling holes and whittling pegs and driving them into the holes. But every now and then a nail sprang away from the tough oak when the hammer hit it, and if Pa was not holding it firmly, it went sailing through the air.
>
> Then Mary and Laura watched it fall and they searched in the grass till they found it. Sometimes it was bent. Then Pa carefully pounded it straight again. It would never do to lose or waste a nail.

Today, 150 years after the period depicted in *Little House,* the price of nails has fallen far enough (and alternative means of fastening objects together have been developed) so that an individual nail seems cheap and disposable.[1] The importance of nails in the economy, as shown in Figure 1, had dropped to a *de minimis* share of less than 0.01 percent by 2017.

[1] Despite the very substantial declines in the relative price of nails documented in this paper, some builders still recover them from construction sites—although this now largely is seen as a sign of frugality. For example, the *New York Times* reported in the obituary of Fred Trump (Donald Trump's father) that the senior Trump at one of his construction sites "would walk through the studs and across the plywood floors, picking up unused nails to hand back to his carpenters the next day" (Rozhon 1999).

Figure 1
Nails as a Share of Nominal GDP

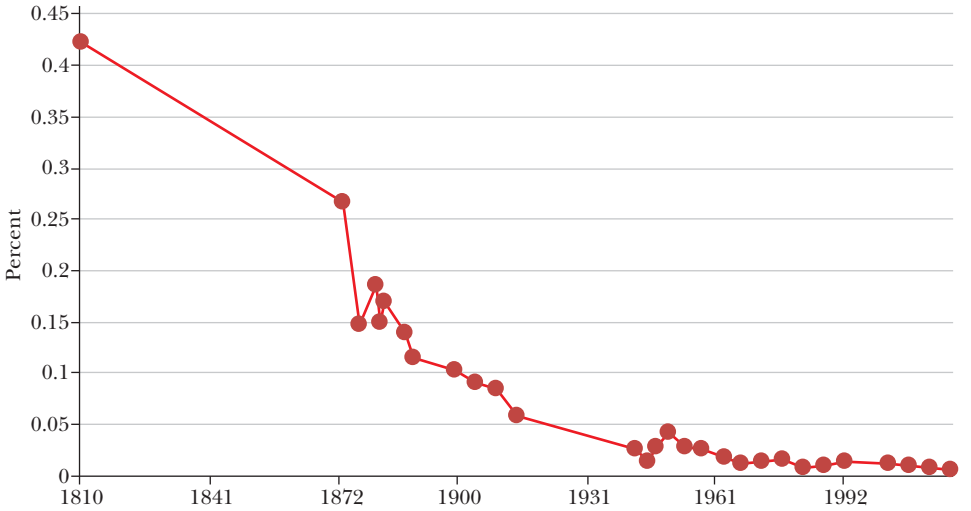

Source: Author's calculations based on French (1858) and data from the Census of Manufactures and from the American Iron and Steel Institute, as detailed in the online Appendix available with this article at the *JEP* website.
Note: The figure plots domestic absorption (production plus imports minus exports) as a share of nominal GDP. Data include nails, spikes, tacks, and staples. Data points are shown as dots connected by line segments. For underlying details of data and methods, see online Appendix, available with this article at the *JEP* website.

This essay investigates the long-term patterns in nail prices and their underlying causes. Some of the facts that emerge include: 1) real prices for domestically produced US nails fell dramatically up until about 1930; 2) improvements in multifactor productivity are the main cause of price decline in nails up though the 1930s, although declining prices of materials also play a role; 3) reported real nail prices have risen since the 1930s owing partly to rising materials prices but also likely to difficulties in tracking prices of basic nails as US nail producers shifted toward specialty nails in the wake of import competition; and 4) including the complementary technology of nail guns, the price of *installed* nails has risen much less than the price of nails themselves in recent decades. A recurrent theme that runs throughout the discussion is that changes in the production of even a humble item like the nail are often intertwined with larger technological changes in the US economy that touch on multiple areas of both of micro- and macroeconomics.

A Brief History of Nails

Nails fall into three broad types—hand-forged, machine cut, and wire—with each of these types dominant in each of three overlapping periods as described in

Adams (2002), Lewis (1998), and Wells (1998). (Additional details on the timeline and additional source citations are in Appendix Table A1.) Hand-forged nails have been made at least since Roman times and continued to be made in relevant quantities through about 1820. Forged nails are made by a blacksmith (or nailsmith), hammering the nail from a rod of iron and hammering a head on the top.

Cut nails are made by a bladed machine that cuts nails from thin strips of iron or steel. The first patents for cut nails in the United States were granted in the 1770s and 1780s, and a flood of patents followed in subsequent years. The manufacturing technology for cut nails improved dramatically during the 1800s as production shifted to factories, mirroring many of the developments of the broader industrial revolution. The power source shifted from water to steam and later electricity, and more and more of the individual tasks of nail-making became mechanized. In the 1880s, production shifted from iron to steel nails.

By the 1880s, wire nails became more prevalent, with the first US patent for wire nails granted in 1877. Initially, wire nails were made from iron wire. By the late 1880s and early 1890s, wire nails were being produced from stronger steel wire in sizable quantities as the technology for producing steel improved. Wire nails are made by cutting each nail from a coil of drawn wire, sharpening a tip, and adding a head. Wire nails remain the dominant type used for most purposes today, though cut nails are still used for some specialty applications such as period architecture and furniture.[2,3] For wire nails, the manufacturing technology also improved considerably in the decades after the 1880s. Today, a significant part of nailmaking has become continuous-process manufacturing, with large wire rolls fed into a machine that automatically transforms that wire into finished nails.

Hand-forged and cut nails look rather similar and, indeed, they have similar "holding power" (or resistance to being withdrawn after being hammered in). One advantage of hand-forged nails over early machine-cut nails is that forged nails could be "clinched": that is, the tip of the nail that extended through the pieces of material being joined could be bent over, or clinched, thereby increasing the holding power of the nail. Early machine-cut nails had the grain of the metal running perpendicular to the length of the nail, and cut nails would break if an effort was made to clinch them. Later, cut nails were made with the grain of the metal running parallel to the length of the nail, and these cut nails could be clinched. Wire nails have considerably less holding power than forged or cut nails, but because each nail is lighter, shipping costs per nail were less. The basic wire nail has changed relatively little since the 1890s, with the graphic in late 19th century Sears catalogues depicting a nail that looks much like one that could be purchased at Home Depot today.

Even for a run-of-the-mill item like nails, the changes in manufacturing processes were a big deal. Prior to the industrial revolution, nails were produced

[2] One US manufacturer—Tremont Nail Company in Wareham, Massachusetts— still makes cut nails, in some cases using vintage machinery with some components dating back to the early 19th century.
[3] Based on a factory tour and conversation with Gary Anderson in early 2020.

one at a time by a blacksmith, and according to Rybczynski (2000, p. 70–71), it took about a minute for a skilled blacksmith or nailsmith to produce a single hand-forged nail. Adam Smith (1776) highlighted early process improvements for pin manufacturing, and many of the same developments would have been applied to nail manufacturing as well. Changes since then have been more dramatic. Currently, a typical nail-making machine with a footprint of about three square feet can produce 300 to 450 wire nails in a minute, while the newest machines can produce 2000 nails per minute. Assume 500 nails per minute, on average, and that a worker can operate seven machines at once.[4] Thus, the number of nails produced per minute of worker time has increased by a factor of 3500 times since the era of hand-forged nails.

These shifts from hand-forged to cut to wire nails track with broader changes across the US manufacturing sector described by Goldin and Katz (1998). They highlighted the evolution from artisanal to factory to continuous-process production technologies, and nails provide a canonical example of that transformation.

The rapid price declines for nails documented in this paper also affected economic activity more broadly in the 19th and early 20th centuries. Those effects were especially pronounced in the construction industry (including railroad beds and fences) during a time when construction represented a larger share of economic activity than it does today: for example, residential and nonresidential construction amounted to 16 percent of nominal GDP in 1839 (Gallman 1966), compared with about 6¾ percent of GDP in 2019. One particularly important transformation within construction is that the drop in prices of nails enabled "balloon-frame construction" to develop in the 19th century as a considerably lower-cost style of construction than the long-used post-and-beam approach (Jackson 1987). Balloon-frame construction—in which a building's walls are made up of dimensional-lumber studs (typically 2" x 4" today) that carry the weight of the structure—use many nails and became feasible only after the relative price of nails had fallen far enough. Outside of construction, other downstream sectors benefitting from price declines for nails include furniture, wooden containers and boxes, and many other products made from wood.

Beyond reductions in cost of production, the variety of nails produced has expanded significantly over the years, including nails made from specialized materials for particular applications, nails made with coatings to prevent rust, and nails made with rings around the shank to increase holding power. Indeed, the website of Maze nails, one of a handful of current American nail manufacturers, lists 110 distinct varieties of nails, not counting different sizes and colors of the same type of nail.[5]

[4] Roelif Loveland, President of Maze Nails, provided information in a January 2020 email on the number of nails per minute made by a typical nail-making machine as well as the number made by the latest machines. In addition, he reported that a worker can operate seven machines at once.

[5] Catalog available at https://www.mazenails.com/assets/pdf/maze-nails-price-list-catalog.pdf. The list of nail varieties is on page 4 of the 2020 catalog.

Figure 2
Nominal Price of Nails, 1695–2018

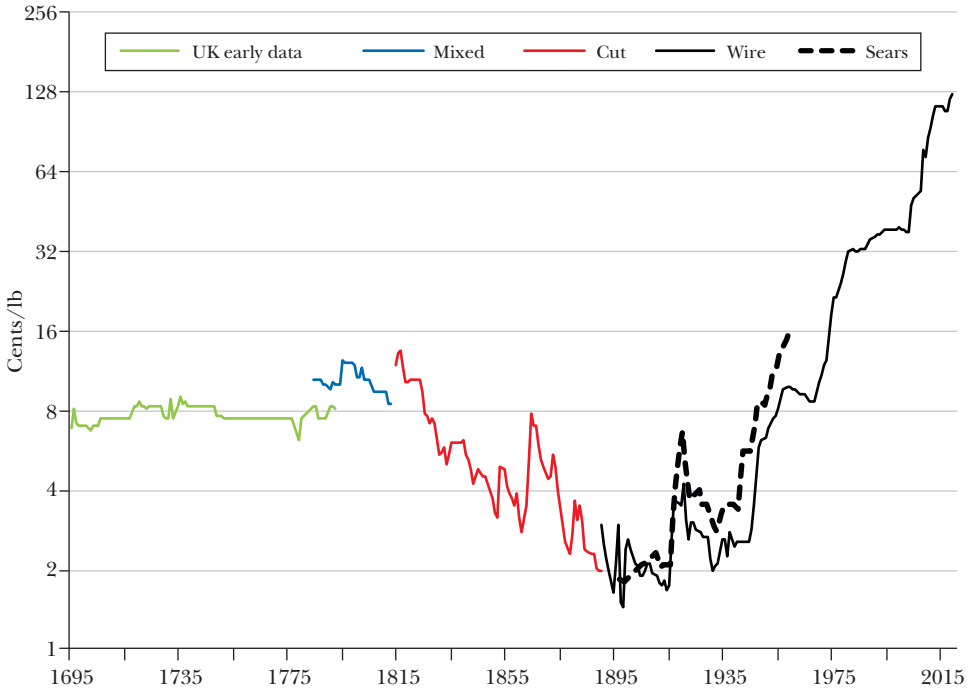

Source: Author's calculations based on Beveridge (1939), Cole (1938), Bureau of Labor Statistics Producer Price Index reports, and other sources as described in the online Appendix.
Note: For underlying details of data and methods, see online Appendix.

Nail Prices

Raw Nominal Price Data

Figure 2 plots the raw data on the nominal price of nails in cents per pound; nearly all the price quotes I found for nails were on a price per pound basis. In the figure, different colors capture the different "regimes" of data. The green segment refers to prices from Beveridge (1939) for nails in the United Kingdom from 1695 to 1792. Beveridge collected these prices from log books from a hospital in Greenwich, United Kingdom, as well as from schools and government departments. Given the time period, these quotes must have been for hand-forged nails, and the quotes cover a wide range of sizes, with prices provided in UK shillings per 12 pounds of nails. I converted these prices to US dollars using an exchange rate from UK pounds to US dollars for 1792.[6] Thus, before 1792, these prices are capturing movements in prices in the United Kingdom, indexed to the 1792 value in cents per pound.

[6]The exchange rate used for 1792 is $4.47 per UK pound sterling, downloaded from the Measuring Worth website at www.measuringworth.com (Measuring Worth 2011a,b).

The blue segment in Figure 2 captures prices from Cole (1938) for the period from 1784 to 1813. These data represent wholesale prices in the Philadelphia market.[7] The type of nail is not specified, but, given the time period, these quotes probably cover a mix of hand-forged and cut nails. I refer to this segment of data as "mixed." The quotes in Cole are for various size lots, and they were all converted to cents per pound of nails.

The red segment covers machine-cut nails from 1814 to 1890. These data are from Cole from 1814 to 1828 and from various other sources in the later part of the period. The price quotes cover wholesale prices for New York City through 1849 and prices posted by a Pennsylvania nail manufacturer through 1890. The quotes are for dollars per 100 pounds and were converted to cents per pound.

The black segment in Figure 2 covers the period from 1890 to the present and refers to wire nails. This segment incorporates data from the Bureau of Labor Statistics, reflecting a number of different reports for the earlier periods and Producer Price Indexes for more recent years. In the data, the quotes are for different varieties of nails and are quoted for lots of various sizes. All of these quotes were converted to cents per pound.

The choice of breakpoints in price quotes across the different types of nails is largely driven by data availability. That said, the switchovers to cut nails in 1814 and to wire nails in 1890 are consistent with what historical archaeologists characterize as the eras in which each type of nail was prominent (for example, as described in Wells 1998, Figure 8).

At the data breakpoints, as can be seen in Figure 2, there are some discontinuities in prices. In the late 1700s, the series for forged nails in the United Kingdom (shown in green) is below the series for "mixed" nails (shown in blue) in the United States. This gap could reflect any of a number of factors. Perhaps nails were produced more cheaply in the United Kingdom in this period. However, shipping nails was expensive, so if shipping charges were added to the UK prices, they might look more like the higher US prices during this period. In addition, the descriptions of the nails for which prices were collected often are limited, so there likely are differences in what is being priced. The break after 1813 reflects the data source that began explicitly pricing cut nails starting 1814. The implied jump in price in that year is not so surprising given that it occurred during the War of 1812, when many commodity prices rose. The other interesting overlap is that between the series for cut nails (red) and wire nails (black). It appears that wire nails were more expensive, raising the question of why buyers shifted to wire nails, particularly given

[7]The online Appendix provides details on the sources from which all of the prices in Figure 2 were drawn. As highlighted in the text, the price quotes from nails are taken from disparate sources, often reflecting prices in major cities that may or may not be representative of nail prices in other locations or purchased through other distribution channels. Overcoming this challenge would be difficult given available data. That said, Rothenberg (1979) found that movements in prices received by farmers for agricultural products in rural Massachusetts were similar to price movements for those products in Boston and New York City.

their less impressive holding power. The resolution of this puzzle is discussed in the section below on quality adjustment.

An Alternative Price Series since 1897

As a check on the nail prices from the US Bureau of Labor Statistics covering the period since the late 19th century, I also collected prices for 2" nails from Sears catalogs. These quotes extend from 1897 to 1960, the last year the Sears catalog included steel wire nails in a variety of sizes. One appeal of the Sears prices is that it is straightforward to compare like to like by pricing the same nail over time. For example, it is possible to track the price of 2" nails—that is, size 6d—over this period, where a number followed by the letter d is a standard for nail sizing. Moreover, economists have a long tradition of using Sears catalogs (and those from other retailers) to track prices over time—including work by Rees and Jacobs (1961) and Gordon (2008)—so these catalogs are a natural source to use as a check on the government price data.

These Sears prices are plotted as the dashed black line in Figure 2.[8] Starting in 1897, these prices are nearly indistinguishable from the US Bureau of Labor Statistics prices shown by the solid line. However, subsequently, the Sears series starts to rise more rapidly than does the Bureau of Labor Statistics series and by 1960 the Sears price is 70 percent higher than the other price (16.4 versus 9.6 cents). While it is difficult to ascertain definitively the source of this growing gap, a likely factor is a shifting wholesale/retail margin at Sears. The price quotes forming the basis for my preferred price series are producer prices. The Sears price is, technically, a retail price, though in the late 1800s and the first part of the 1900s, Sears sold nails in large volumes (100-pound kegs) and sold them directly to homebuilders (as inferred from the advertising copy in the catalogs). In the more recent period, it appears that the Sears catalog was catering almost exclusively to small-scale retail purchasers of nails. In particular, starting in the mid-1930s, Sears began quoting prices for one-pound packages and in the early 1940s began quoting prices for five-pound packages; by this time, Sears no longer quoted prices for 100-pound kegs of nails in its catalog.

These changes support the view that Sears gradually shifted from being more of a wholesaler to more of a retailer in the market for nails. Such a shift would be consistent with the more rapid increase in nail prices at Sears relative to the Bureau of Labor Statistics prices and would suggest that Sears prices after about 1930 may not be so useful for tracking producer or wholesale prices of nails. This gap also emphasizes the importance of care in choosing what price of a product to track over long spans of time.[9]

[8] The Sears prices of nails are plotted on a cents per pound basis to match the units of the other raw series shown in Figure 2. As noted below, these prices can be readily converted to a cents/nail basis.

[9] I suspect that this issue of shifting wholesale/retail margins could affect other products in the Sears catalogs as well. In a different context, Gordon (2009) found evidence for a later period that prices for apparel in the Sears catalog increased relative to the Consumer Price Index as Sears was increasingly undercut on prices by new competitors such as Target and Walmart. My conjecture about margins, along

Real Prices

The real price of nails over time is constructed by deflating the nominal nail price index with an index of consumer prices. I construct a consumer price index back to 1695 by using ratio splices to link together series for the United Kingdom retail price index from 1695 to 1784, the US Consumer Price Index for 1784–1928, and the chain price index for Personal Consumption Expenditures from the US Bureau of Economic Analysis for 1929–2018. Of course, as Gordon (2009), Nordhaus (1997), and many others have noted, comparisons of price indexes over very long spans of time raise a host of difficult issues. Nonetheless, it seems more relevant to focus on real prices—despite the inherent limitations—rather than nominal prices, particularly given the central interest in the prices of nails relative to those of other goods and services.

Real prices of nails, relative to the consumer prices, are shown in Figure 3 on a cents-per-pound basis from 1695 to 2018 by the dashed line segments. On this basis, without any adjustment for quality changes, the real price of nails was relatively stable from 1695 through the early 1800s (at least compared with the large decline that followed) with a peak in the mid-1700s. Then during the 1800s, the real price fell substantially through the 1930s before rising, on balance, over subsequent decades.

Quality Adjustment

Ideally, price indexes are constructed on a quality-adjusted basis, so that price comparisons over time are like-to-like. For example, directly comparing prices of a car today to a Model T in 1908 would be problematic because the quality of the vehicles is so different that it would not represent, without some adjustments, a like-to-like comparison. Quality-adjusting prices of nails is much easier because, as noted above, nails have changed very little over the centuries. In my calculations, just two adjustments are made to construct a quality-adjusted price index for nails.

First, nail prices are adjusted so as to track as accurately as possible the same size nail over time. The series on nail prices on a cents-per-pound basis does not price a homogenous product, but rather is conflating prices of nails of many different sizes. To standardize the size of nails in the index, I convert everything to be as equivalent as possible to a 2" (size 6d) nail. For the earlier periods before nails sizes were standardized, I also standardize on 2" nails. To do this, I convert pounds of nails to the number of 2" nails using an estimate of the number of 2" nails per pound. I use a count of 85 nails per pound for both forged and cut nails for the period spanning 1695–1889 based on an average of counts from multiple sources (as detailed in the online Appendix). Wire nails are considerably thinner, and as the price quotes switch to wire nails in 1890, the count jumps to 150 nails per pound for 2" nails in that year. The count briefly increases further to 181 during World War II when nails available to the public became thinner, presumably to conserve essential

with Gordon's results, suggests that researchers using catalog prices need to be mindful of what catalog prices represent.

Figure 3
Real Price of Nails, Cents/pound (Dashed) and Cents/Nail (Solid), Relative to Overall Consumer Prices
(in 2012 US Dollars)

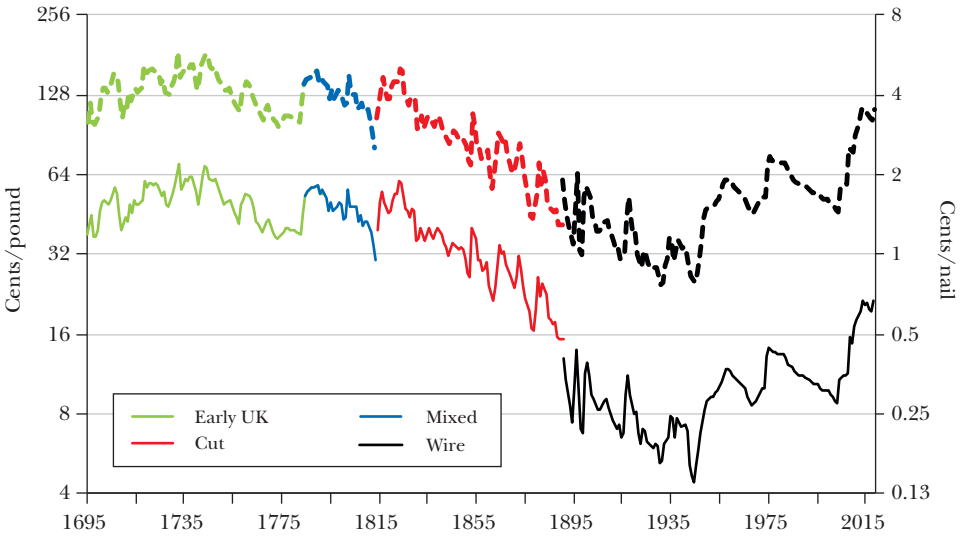

Source: Based on author's calculations using prices from Figure 2, counts of the number of nails per pound as described in the text, and an index of consumer prices described in the online Appendix.
Note: For underlying details of data and methods, see online Appendix.

war materials. After the war, the count dropped back to 168, where it has remained since. With these counts, I convert prices in real cents per pound to real cents per nail, and the solid line segments in Figure 3 show these prices.

One key difference is evident between the two lines in Figure 3 as a result of this quality adjustment. Prices fell by a larger multiple on a *per nail* basis than on a *per pound* basis. In particular, note what happened around the period of transition in price quotes in 1890 from cut nails to wire nails. On a *per pound* basis, wire nails look more expensive than cut nails, but on a *per nail* basis, wire nails look less expensive than cut nails. The reason for this difference is the higher count of wire nails per pound than of forged or cut nails. Temin (1964) pointed out that the lighter weight of wire nails meant lower shipping costs per nail, enhancing the attractiveness of wire nails. This discussion highlights the importance, for the purposes of price measurement, of being as precise as possible about the product being priced.

The second adjustment accounts for differences in other characteristics of hand-forged, cut, and wire nails. Key differences include "holding power" and shipping costs.[10]

[10]The online Appendix provides some additional information on the holding power of cut and wire nails, as well as some additional analysis of shipping costs.

Regarding holding power, wire nails have less holding power than cut or forged nails. This outcome occurs primarily because the cross section of a cut or forged nail is rectangular and tapered compared with the round and untapered cross section of a wire nail. The greater holding power comes from the wedging action arising from the shape of the cut nails when pounded into wood. In the world of wood engineering, these differences in holding power have been measured, and the literature suggests that cut nails have about twice the holding power of wire nail. The prices plotted in Figure 3 could be adjusted explicitly for differences in holding power. However, for many applications, the smaller holding power of wire nails is perfectly sufficient—so that the greater holding power of cut nails would be unnecessary and may not be valued as highly as would be suggested by an explicit adjustment for holding power.

Shipping costs are another characteristic that would be important to some buyers. For example, the 1897 Sears catalog indicates that shipping costs for a 100-pound keg of nails from Chicago to Boston amounted to about 20 percent of the price of the nails. Given the greater number of wire nails that would be in a keg compared with the number of cut nails, shipping costs per nail would have been considerably lower for wire nails than for cut nails. Again, with some additional data and calculation, the prices in Figure 3 could be explicitly adjusted for differences in shipping costs.

How best to account for identifiable differences in holding power, shipping costs, and any other relevant characteristics of different types of nails? One modern approach would be to construct a so-called hedonic price index that would use statistical techniques to account explicitly for all relevant characteristics of nails and for how the marketplace valued those characteristics. Data limitations make that impossible in this case. Thus, I use another common technique, a matched-model procedure. This methodology links the cents-per-nail prices in Figure 3 across the switchover points from one type of nail to another; this linking will accurately adjust for quality change to the extent that the price/performance ratio for nails was equalized by the market in periods when multiple types of nails were available. This approach is a sensible and widely used methodology to construct a consistent price index that adjusts for changes in quality over time so that like can be compared to like, given available data. Specifically, I start with the prices for the most recent period (the black segment for wire nails), and then link backwards. In each of the crossover years (1890, 1814, and 1784), I use the price from the more recent type of nail and link and extend backwards from that year using percent changes in prices for the earlier type of nails in the earlier years.

The matched-model price series is plotted in Figure 4. It shows that real prices of nails fell by a factor of about 10 from the late 1700s to its low point in the middle of the 20th century, amounting to a 1.6 percent annual average rate of decline from 1792 to 1930. Declines in real prices during this period were uneven, with especially rapid declines during two periods: from the early 1820s to 1860 and from the early 1880s through about 1930. Since the middle of the 20th century, real prices have risen considerably on balance.

Figure 4
Real Price of Nails: Matched-Model Index, 1695–2018

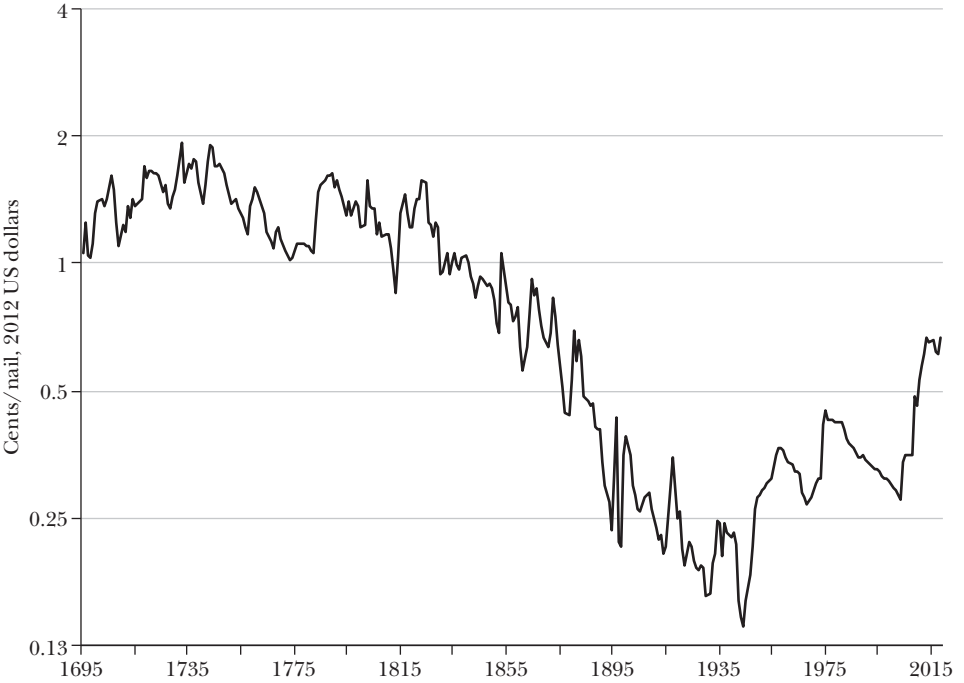

Source: Based on author's calculations from prices shown as solid lines in Figure 3.
Note: For underlying details of data and methods, see online Appendix.

What Accounted for Changes in the Real Price of Nails?

Over long spans of time, changes in nail prices, as for any product, should largely depend on changes in the cost of inputs and on advances in the production technology (very roughly captured through multifactor productivity). Economists capture these relationships through decompositions that break down changes in prices of a product into its proximate sources. While such decompositions are not causal in any sense, they do provide valuable insights; for example, quantifying linkages between changes in prices of nails and changes in the prices of key inputs such as iron or steel.

The basic setup for this decomposition begins with a conventional growth accounting relationship that decomposes quantities. In a textbook version, growth in output (value-added) is broken down into contributions from capital, labor, and multifactor productivity. Typical assumptions in this setup include a Cobb-Douglas production function with constant returns to scale (factor shares summing to one)

and perfect competition in input and output markets. Multifactor productivity captures all elements other than inputs that affect output and often is taken by macroeconomists to be a rough measure of the state of technology; that is, multifactor productivity rises if it is possible to produce more output from the same quantity of inputs. With these assumptions, the growth in output equals a weighted average of growth in capital and labor (with the weights equal to factor shares) plus the growth in multifactor productivity.

For the purpose here, a more complex version of this relationship, the KLEMS variant of growth accounting as described in Jorgenson, Fukao, and Timmer (2016), is more useful. This approach decomposes growth in gross output (Y) into—in addition to capital (K) and labor (L) inputs—energy inputs (E), materials (M), and purchases of services (S). Purchased services include any services purchased from outside the business, such as repair services. The full relationship, including multifactor productivity is given by:

$$\dot{y} = \alpha_K \dot{k} + \alpha_L \dot{l} + \alpha_E \dot{e} + \alpha_M \dot{m} + (1 - \alpha_K - \alpha_L - \alpha_E - \alpha_M)\dot{s} + (\dot{mfp})$$

where lower case variables with dots over them represent growth rates (log differences) of inputs and α are the factor shares. In this setup, both the growth rates of inputs and the factor shares vary year to year.

Because I am focusing on changes in prices rather than quantities, I turn to the dual representation of the production function. While the production function focuses on the relationship between the quantities of output and inputs, the dual is a corresponding relationship between the price of output and the cost of inputs. In percent change terms the dual relationship is:

$$\dot{p} = \alpha_K \dot{rc} + \alpha_L \dot{w} + \alpha_E \dot{p}_e + \alpha_M \dot{p}_M + (1 - \alpha_K - \alpha_L - \alpha_E \alpha_M)\, \dot{p}_S - \dot{mfp}$$

where P is the price of nails plotted in Figure 4, RC is the rental costs of capital, W is wages, P_E is the price of energy, P_M is the price of materials, and P_S is the price of purchased services, respectively, with all prices measured in real terms. Again, lower case variables with dots over them represent growth rates (log differences) of inputs and the terms are the factor shares. Notice that growth of multifactor productivity enters the dual equation with a negative sign. When focusing on quantities, gains in multifactor productivity raise output; when focusing on prices, gains in multifactor productivity hold down prices relative to the contributions of factor costs. In short, this formulation says that changes in the price of nails reflects a weighted average of changes in input (factor) prices using factor shares as weights, with an adjustment for multifactor productivity growth.

To implement this dual equation, I proceeded as follows. The series on quality-adjusted nail prices from Figure 4 is used for the left-hand side. For the right-hand side terms for each input, I compiled an annual dataset on factor prices and factor shares for capital, labor, energy, materials, and services as described below. Finally, in this type of analysis, growth in multifactor productivity is calculated as a residual

by subtracting the terms for changes in input costs from changes in prices. As noted above, multifactor productivity captures all elements other than changes in factor prices or shares that affect the price of nails, such as rearranging the physical layout of production once electricity became widely used as a power source.

Data on Factor Prices

For factor prices, I assembled annual data back to 1790 on the rental cost of capital, wages, the price of energy as likely used in manufacturing, the price of steel and iron for earlier years, and an estimate of the price of purchased services. I used 1790 as the initial year because that is when the wage series started. Briefly, I used data from various sources, including the Bureau of Labor Statistics, Gallman and Rhode (2019), measuringworth.com, the NBER MacroHistory database, Temin (1964), and Warren and Pearson (1933). Each factor price was converted to real 2012 dollars using the index for consumer prices that was used to calculate real nail prices. Additional details on sources and data construction are described in the online Appendix.

Figure 5 plots the resulting real factor prices for each input with each price indexed to 100 in 1790. Real wages rose substantially from 1790 to 2018, as did the real price of purchased services. (No surprise that purchased services rose in a similar way to services, given that wages were used as an extrapolator for costs of services prior to 1949.) The real rental price of capital was relatively flat from 1790 until about 1900, and then moved lower through 2018. The real price of energy followed a similar trajectory to that of capital, though it declined more rapidly during the 20th century until the early 1970s energy crisis, after which it moved higher. The real price of materials fell through about 1950 and then began moving higher.

These long spans of factor prices rely on solid research by statistical agencies and economic historians. That said, they undoubtedly are measured with some error. Again, linking together price statistics over long spans of time can miss the benefits of more revolutionary changes in inputs. For example, the benefits of the switch from water-power to steam-power likely are missed in my methodology, given that fuel prices were used to extend back into the period when water was a principal power source. The linking methodology also may miss benefits of the switch from steam to electric power. Moreover, I pulled data from many different sources, likely creating inconsistencies in the definitions and data collection procedures used.

One area of particular concern is the rental price of capital, because the capital stock deflators that I rely on were not constructed with the attention to quality adjustment that is typical for many types of capital in a modern price measurement framework. Nonetheless, these factor prices capture important trends over time. In addition, the price decomposition I implement relies only on growth rates (log differences) in these factor prices, so some inconsistencies in levels will be washed out. Finally, as shown by the sensitivity analysis described below, even substantial allowances for measurement error in the cost of capital would not alter the basic story that emerges from my analysis.

Figure 5
Real Factor Prices, 1790–2018

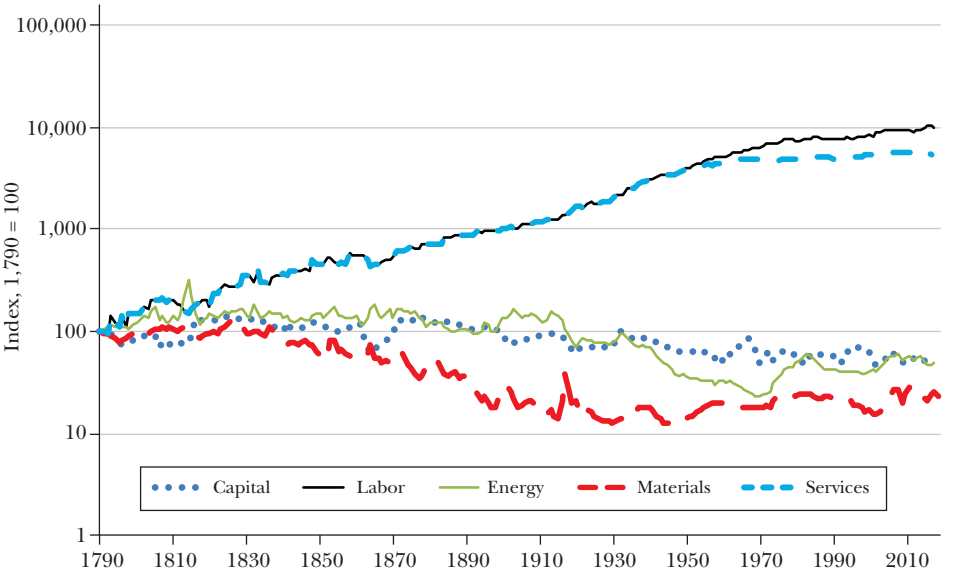

Source: Based on author's calculations in the online Appendix.

Data on Factor Shares

For factor shares for capital, labor, energy, materials, and purchased services, I also had to rely on a range of sources, including a decomposition for fabricated metal products since 1949 into these inputs published by the Bureau of Labor Statistics; the *Hand and Machine Labor* study published in 1899 that compared machine production for many products in 1897 to hand production for the same products in 1813[11]; and the Bateman, Weiss, and Atack extracts from the 1850 and 1870 Census of Manufactures as described in Atack and Bateman (1999). From these sources, I obtained annual factor shares from the period 1949–2018 and shares for 1813, 1850, 1870, and 1897. For the intervening years, I linearly interpolated shares. For the years prior to 1813, I used the 1813 shares. Because I am using a decomposition into five factors (plus multifactor productivity), the labor and capital shares will not be the familiar two-thirds and one-third, although the five factor shares still sum to unity under the maintained assumption of constant returns to scale.

[11] For background on the *Hand and Machine Labor* study, see Atack, Margo, and Rhode (2019, 2020). They recently digitized the voluminous information in that study, making it much more widely accessible to researchers. Although not directly related to the measurement of factor shares, see Margo (2015) for a discussion of likely measurement error for the labor input of smaller firms relative to larger firms.

Estimating factor shares for 1850 and 1870 from the Bateman, Weiss, and Atack Census extracts has not been done often and is not entirely straightforward so a brief explanation here is warranted, with details of this and other data sources and procedures described in the online Appendix.[12] For 1850 and 1870, the extracts provide firm-level accounting information for a sample of firms in those years. I identified the firms in the extracts for which nails were the dominant output: seven firms in 1850 (three hand-powered, three water-powered, and one steam-powered) and four firms in 1870 (three steam-powered and one water-powered). For each firm, the extract provided figures for the value of output, the wage bill, the cost of energy inputs, and the cost of materials inputs. From these, I could calculate factor shares for labor, energy, and materials. The extract also provided an estimate of firm "capital," although it is not entirely clear what is included in that figure. Nonetheless, I used that number and the user cost of capital formula (along with some assumptions about components of the formula) to estimate the factor share for capital. For purchased services, I calculated the factor share as a residual, subtracting the value of all other inputs from the value of nails produced. The numbers used for my decomposition are the unweighted average factor shares across the seven firms in 1850 and across the four firms in 1870.

Figure 6 plots the estimates of factor shares from 1790 to 2018. Two observations fall out immediately from this figure. First, the capital intensity of nail production rises over the period, with capital's share rising steadily over time from a low between 3 and 4 percent at the beginning of the sample to about 15 percent by 2018. Second, the energy share remains low throughout the sample, though it is lowest in the earliest years. Other observations are more subtle, given the ups and downs of factor shares. Labor intensity was relatively high early in the sample during the period of hand production (about 34 percent in 1790). As nail manufacturing began to become mechanized, the share fell to about 25 percent in 1850. The share then rose over the next 100 years to reach just above 40 percent in the early 1950s. Since then, my estimate of the labor share drops back to about 30 percent by 2018, likely owing to increased mechanization as reflected by the rise in the capital share during this period. The materials share generally increases and then falls back, rising from around 36 percent in 1790 to about 50 percent by 1870, drifting lower to about 40 percent by the late 1940s, and then fluctuating around that value through the end of the sample. Finally, the purchased services share generally declines and then increases part way back; specifically, it falls from around 25 percent early in the sample to around 10 percent in the 1980s before rising back up to about 15 percent by the end of the sample.

In my view, the factor share estimates are the weakest element of the price decomposition, given the sparseness of years for which I obtained estimates, the range of very different sources and methodologies, and the incompleteness of data even for the years for which I obtained estimates. Indeed, one could argue that

[12] One other paper that used these data to identify key aspects of production processes is Margo and Atack's (2019) paper on US blacksmithing from 1850 to 1870.

Figure 6
Factor Shares, 1790–2018

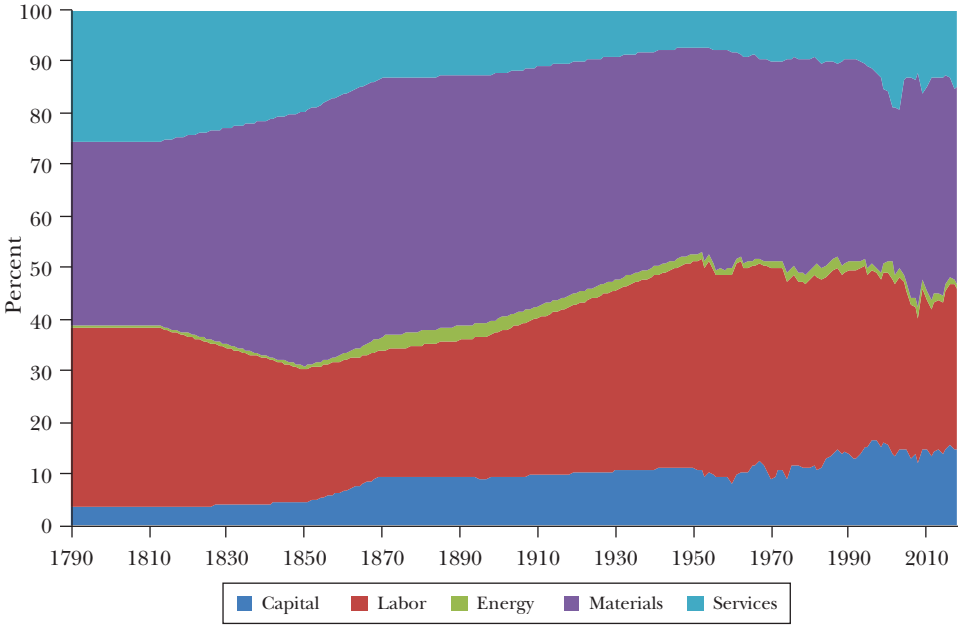

Source: Based on author's calculations from sources described in the online Appendix.
Note: For underlying details of data and methods, see online Appendix.

we do not know enough about early factor shares to draw confident conclusions. However, the sensitivity check described below suggests that any plausible measurement error in factor shares would not significantly affect the conclusions drawn from the decomposition of proximate sources of changes in real prices of nails. In addition, my estimates for factor shares are not that different from those of other researchers using different sources (again as described in the online Appendix).

What Accounted for Price Changes?

With data on factor prices and shares in hand, I can calculate the price term on the left-hand side and the term for each input on the right-hand side of the price decomposition for each year, estimating multifactor productivity growth as a residual. For these annual estimates, I follow usual practice in the growth accounting literature in which the factor shares used in the decomposition are the average of shares for the two years covered by each annual percent change in a factor price; for example, the factor shares entering the price equation for, say, 1872 are the averages of the annual factor shares for 1871 and 1872. Although I calculate the decomposition for each year, I only report averages over longer spans of time, given the short-run volatility in markups, tariffs, taxes, and other factors.

Table 1
Sources of Real Price Change for Nails, 1790–2018
(annual averages over periods shown)

	Percent change	*Contribution of (in percentage points)*					
Period	*Nails (1)*	*Capital (2)*	*Labor (3)*	*Energy (4)*	*Materials (5)*	*Purchased services (6)*	*Multifactor productivity (7)*
1791–1820	–0.37	0.03	0.66	0.01	–0.04	0.49	1.51
1821–1860	–1.89	0.00	0.84	0.00	–0.75	0.62	2.60
1861–1880	0.45	0.08	0.30	–0.02	–0.11	0.15	–0.06
1881–1930	–2.89	–0.12	0.62	–0.02	–1.30	0.23	2.31
1931–2018	1.60	–0.04	0.75	–0.01	0.31	0.10	–0.48
Memo:							
1791–2018	–0.36	–0.03	0.69	–0.01	–0.31	0.28	–0.97

Source: Based on author's calculations from data plotted in Figures 4, 5, and 6.
Note: For each period shown, the percent change in the price of nails in column 1 is, by definition, the sum of the contributions of the factors of production (columns 2 to 6) less the contribution of multifactor productivity (column 7). All prices are real prices, relative to the consumer price index described in the text. For underlying details of data and methods, see online Appendix.

Table 1 reports the price decomposition for nails. The selected periods shown correspond to the eras I identified in trends for real prices of nails: the modest declines from the late 18th century through about 1820, the first wave of very rapid declines that came with increasing mechanization through 1860, the flattening out during the next two decades, the second wave of very rapid price declines between 1881 and 1930, and then the substantial increases in real prices from the 1930s through the end of the sample in 2018.

Looking across periods, the most interesting and important result in the decomposition is the large contribution of multifactor productivity. In the early period (1791–1820) when real nail prices declined about 0.4 percent a year, rapid advances in multifactor productivity pulled down nail prices by an average of 1.5 percentage points per year. (Recall that increases in multifactor productivity hold down nail prices.) The improvement in multifactor productivity was even more substantial during the first (1821–1860) and second (1881–1930) waves of rapid decline in real nail prices, with multifactor productivity growth holding down the annual average rate of change of real prices by 2.6 and 2.3 percentage points in these periods, respectively.[13] Then, in more recent decades, declines in multifactor productivity growth accounted for nearly 0.5 percentage point per year of the *increase* in real nail prices.

[13]The decomposition used here abstracts from shifts in the complementarity of capital to skilled and unskilled labor as the technology of nail-making evolved, a factor highlighted by Goldin and Katz (1998). To the extent that such shifts are not accounted for by input prices or shares, they will be reflected in the multifactor productivity term in the decomposition.

The very large role of multifactor productivity growth during the two waves of rapid price declines in real nail prices suggests that developments other than changes in input costs or factor shares contributed importantly to declining nail prices. Interpreting multifactor productivity can be difficult because it is calculated as a residual, so it is impossible to know exactly what accounted for a given change. Moreover, any measurement error in factor prices or shares would be reflected in multifactor productivity. That said, other research highlights one important explanation for the big role of multifactor productivity. For the 19th century, Atack, Margo, and Rhode (2019, 2020), using their digitized version of the *Hand and Machine Labor* study, provide evidence that mechanization per se accounted for only about one-third of the productivity advance in manufacturing between 1813 and 1897. In their analysis, the balance of 19th century productivity advance reflected other factors, such as the increasing specialization of labor, which would be captured in multifactor productivity.

Similarly, the outsized role of multifactor productivity likely reflects, at least in part, the story told by David (1990) for the transition from steam to electric power early in the 20th century. In particular, he highlights the important role of reorganization of manufacturing processes in boosting labor productivity, a factor that also would show up as multifactor productivity growth.

Materials prices also contribute importantly to declines in real nail prices, especially during the first and second wave of rapid price declines. Specifically, falling real materials prices pull down real nail prices by an average of about 0.75 and 1.3 percentage points per year in the periods 1821–1860 and 1881–1930, respectively. These rapid price declines in materials prices—iron in the earlier period and steel in the latter period—likely reflect the dramatic advances in the manufacture of those inputs. This pattern of declining materials prices reversed more recently, consistent with the general upturn in commodity prices after 1950 (as documented by Jacks 2013). The upturn in materials prices in the decades after 1930 contributed about 0.3 percentage points per year to the average annual increase in real nail prices.

Among other factors, changes in the rental price of capital account for relatively little of the change in the real prices in any period, reflecting the modest changes in the price of capital and its relatively small factor share. This result may seem surprising, but it does fit with the story that there was more than just mechanization to the changes in nail manufacturing. Real wages make a consistent and noticeable positive contribution to nail prices ranging from about 0.3 to 0.8 percentage point per year on average, partially offsetting other sources that pushed down nail prices in some periods. Interestingly, declines in the price of energy make only very modest contributions to declines in the real price of nails through 1930 and also very little contribution to the upturn in nail prices in more recent decades. The factor share of energy is sufficiently small so that even substantial swings in energy prices do not (and cannot) account for much of the movement in real nail prices.

Sensitivity Analysis and Measurement Error

As noted, the estimates of factor shares are based on incomplete source data and so could well be subject to measurement error, especially for periods prior to 1949 (the first year for which more detailed data from the Bureau of Labor Statistics on shares are available). In addition, concerns could be raised that the price deflator for capital in the earlier periods did not fully account for quality improvements in successive vintages of nail-making machinery. To assess the effect of some of the more plausible measurement errors, I did two sensitivity analyses.

First, I held factor shares prior to 1949 fixed at their 1949 levels and recalculated the decomposition of price changes. In this "fixed-share" scenario, the overall pattern of contributions is very similar to that in the decomposition reported above with my estimates of factor shares. This result suggests that measurement error in factor shares would likely have little effect on the conclusions drawn from the decomposition of price changes.

Second, I considered an alternative scenario for capital prices in which the deflator for capital falls more rapidly than in the decomposition above to capture the possibility that significant improvements in capital were not reflected in its deflator. If that were the case, then the contribution of the rental cost of capital to price declines for nails would be larger than the negligible effect in the numbers reported above and the contribution of multifactor productivity to price declines—which is calculated as a residual—would be smaller. However, even significant measurement error in the price deflator for capital would not affect the overall story. Between 1790 and 1948, the real rental price of capital used in the decomposition declined at an average annual rate of 0.29 percent. (By comparison, the average rate of decline of this variable from 1949–2018 is 0.32 percent.) Suppose that the true real rental price of capital declined 4 percentage points faster per year than is reflected in the data I relied on, a change that would represent a massive reassessment of deflators for capital in the 19th century. With capital's factor share averaging in the neighborhood 10 percent, adopting this alternative cost of capital would make capital's contribution to declining real nail prices larger by 0.4 percentage point (=4 x 0.10) a year with an offsetting reduction in the amount by which multifactor productivity growth held down nail prices. While certainly noticeable, this shift is modest enough to suggest that possible measurement error in the price deflator for capital is unlikely to change the overall story dramatically.

What Accounted for the Increases in Real Nail Prices since the Mid-20th Century?

The decomposition of real price changes in Table 1 suggests that the increases in nail prices since the 1930s primarily reflects two factors: the upturn in materials costs after decades of declines and a *decline* in multifactor productivity after sizable increases in earlier periods. (Again, recall that in this price-based decomposition, increases in multifactor productivity hold down prices, while declines in multifactor productivity boost prices.) Real wages also rose in this period and made a positive contribution to the change in nail prices since the mid-20th century, although that positive contribution is not so different from that in earlier periods. The increase

in steel prices amidst a more general rise in commodity prices is straightforward to understand. However, the decline in multifactor productivity is puzzling. Did the ability of nail manufacturers to organize production really go backwards? This question raises the possibility that other factors, not reflected in the earlier decomposition, may also have played a role. Possibilities along these lines include the difficulty of tracking prices of US-produced basic nails amidst a shift in domestic production toward specialty nails in the wake of import competition for basic nails.

Some straightforward calculations show that the import share for nails (imports/domestic absorption) began a dramatic uptrend in the 1950s, rising to about 70 percent by the 1980s before dropping back somewhat. This increase in imports came earlier than for many other manufactured goods, likely reflecting that by 1950, the technology for producing basic wire nails was widely understood.

Several pieces of evidence suggest that this increase in imports led domestic producers to shift to more specialized, higher-value products. First, the US Bureau of Labor Statistics discontinued specific series for prices of domestically produced nails after 1998 and shifted to pricing a broader category of steel hardware, presumably because the agency no longer could find enough price quotes for domestically produced basic nails. Second, a current US nail manufacturer confirmed the shift in product mix; specifically, the Maze company, one of the few remaining US producers of nails, now produces mostly specialized nails. In addition, by the 2000s the US military was having difficulty finding domestically produced basic hardware (as noted in Mandel 2011).

The final evidence supporting the view that domestically produced nails represented a different product mix than imported nails comes from a comparison of prices of imported and domestically produced nails. Relevant import price data are available starting in 1974, though consistent with the earlier theme prices are available only for broader and broader categories over time. (From 1974 to 2005, the available series covers "nails, screws, nuts, bolts, and rivets of iron, steel, copper or aluminum"; from 2005 to 2010, the series reports prices for "hardware manufacturing"; and since 2010, it covers the even broader category of "fabricated metal products.") On this measure, real import prices fell nearly 60 percent from 1974 to 2018. During this time period, the matched-model index for prices of domestically produced nails *increased*, on balance. This sharp divergence in the trajectory of prices suggests that the import price series and the Producer Price Index for domestically produced items were tracking different products. (Additional information on issues related to nail imports and prices is provided in the online Appendix.)

Taken together, this evidence suggests that in recent decades, my series on real nail prices was no longer tracking the 2" nail that I standardized on in earlier periods. While I do not have access to the underlying detail for the relevant Producer Price Index series, it is easy to imagine scenarios in which prices of specialty nails would rise faster than those for basic nails that were now imported in significant quantities. If so, then that shifting product mix also contributed to the run-up in recent decades in the real nail prices reported in this paper.

Nail Guns and the Price of Installed Nails

Nail guns, which first became widely available in the early 1980s, are an important innovation in the construction industry, and they had sizable effects on the all-in "installed" price of a nail. Indeed, nail guns raise the question of what nail-related product should be priced. If it is an individual nail, then nail guns can be regarded as a distinct piece of complementary capital equipment; however, if the product to be priced is an *installed* nail, then nail guns should be considered an integral part of the process of installing nails. An all-in price for installed nails would include materials (the nails), capital costs (hammer versus nail gun), labor costs, shipping, and everything else. Although I have not calculated an all-in price across all years, the following illustrative example highlights that nail guns are a big deal.

In mid-2020 on Amazon.com, a high-quality pneumatic nail gun that shoots 2" nails could be purchased for $247 and a small compressor and air hose for $170. Packs of nails for the gun are about $75 per 5,000 nails—about 1.5 cents per nail and about twice the price of the same size nail not packed in strips for a nail gun. In contrast, a standard claw hammer costs about $16.[14] For illustrative purposes, assume that a nail gun and compressor last six years in commercial applications and experiences straight-line depreciation over that period (assume the same for the hammer). Further, assume that the hourly wage for a construction worker is $25 per hour, somewhat below what the Bureau of Labor Statistics reported as the average hourly earnings of a construction worker in July 2020, on the assumption that a nail installer is somewhat less skilled than an average construction worker and is paid a bit less. Finally, I assume (based on personal experimentation with a hammer and nail gun) that a worker with a hammer can install six nails per minute and that worker with a nail gun can install 20 nails per minute. With these assumptions, a plausible estimate of the cost per *installed* nail for a worker using a hammer is 7.7 cents per nail (including the cost of the nail, the cost of capital per nail, and the wage per nail). The cost per installed nail for a worker using a nail gun is 3.6 cents per nail, less than half the cost using a hammer.[15] This calculation indicates that the advent and diffusion of nail guns offset some of the rise in the price of an installed nail thereby mitigating the upward trajectory of nail prices in recent decades.

[14] Prices were pulled from Amazon in August 2020 for a DeWalt DWF83PL framing nailer ($247), a Bostitch BTFP02012-WPK compressor and 50' air hose ($170), and a standard fiberglass 16 oz. claw hammer.

[15] Additional assumptions are that a full-time worker is employed for 2000 hours a year and that the worker spends 500 hours installing nails and 1500 hours arranging materials and undertaking other tasks; accordingly, the labor cost allocated to installing nails is $12,500 (= $25x500). On these assumptions, the worker would install 600,000 nails in a year with a nail gun (= 20x60x500) and 180,000 nails with a hammer (= 6x60x500).

Conclusion

This paper focuses on nails—a basic, non-revolutionary manufactured product whose form has changed relatively little over the last three centuries—and constructs price indexes for nails going back to 1695. The price of nails fell significantly relative to prices of an overall basket of consumption goods, with the real price of nails falling by a factor of about 10 from the late 1700s to the middle of the 20th century. A growth-accounting type of decomposition into a weighted average of input costs and multifactor productivity highlights that while falling material prices for iron and steel were important during the two bursts of rapid price drops (1821–1860 and 1881–1930), the most important factor contributing to these rapid price declines was increases in multifactor productivity during the period.

The mid-20th century represented the end of what could be called the "golden age of nail manufacturing" that started in the late 18th century. Since then, real prices of domestically produced nails have risen considerably. These increases reflect higher materials prices as well as, likely at least in part, difficulties in tracking prices of basic nails amidst a shift in US production to more specialized and higher-priced products. This shift in the composition of US production occurred in the wake of foreign competition in basic nails as advances in the technology of producing basic nails dwindled. Pushing in the opposite direction, however, the introduction of nail guns in the 1980s offset a part of these price increases by significantly lowering the "all-in" price of an installed nail.

The study of nail prices highlights some of the issues involved in price measurement: the details of what the available data are measuring can be quite important, along with issues of evolving quality, import substitution, and complementary input technologies. Nails also serve as a reflection of some of the key patterns of US manufacturing industries: for example, the idea that technology-affecting factors other than automation were of crucial importance in the 19th century, and how the growth of imports for basic products where the technological frontier was not advancing much have altered the face of US manufacturing since the middle of the 20th century. All told, a focus on the price of nails provides a useful window into economic changes over the past 300 years.

■ *I would like to thank Ernst Berndt, Daniel Fetter, Robert Gordon, Eric Hilt, Robert Margo, Sam Nelson, Stephen Oliner, Timothy Taylor, and participants at the summer 2011 CRIW/NBER Workshop, the summer 2021 Development of the American Economy NBER Summer Institute, a Wellesley College seminar, and other seminars for helpful suggestions and ideas. In addition, the editors of this journal provided extremely valuable comments and guidance. For assistance on data sources, I thank Gerben Bakker, Joe Davis, Claudia Goldin, Doug Irwin, and Ben Mandel. Peter Chen, Katharine Liang, and Sophie Sun provided excellent research assistance. I also thank Lorna Condon (Senior Curator) for providing access to Historic New England's collection of ephemera and Roelif Loveland (President of Maze Nails) for insights about current manufacturing of nails. Gary Anderson, a fifth-generation nailer at Tremont Nails, kindly took me on a tour of Tremont Nail Company's Mansfield, Massachusetts manufacturing facility. Robert Carver, Brent Neiman, and David Rubin provided assistance in connecting with current nail manufacturers. In addition, a number of librarians were extremely helpful, including Yin Zhu at the Federal Reserve Board and Maureen Booth at the Department of the Interior; Peter Burtch and Deb Fenwick at Northwestern University provided valuable assistance in accessing their library's collection of Sears catalogues.*

References

Adams, William Hampton. 2002. "Machine Cut Nails and Wire Nails: American Production and Use for Dating 19th-Century and Early-20th-Century Sites." *Historical Archaeology* 36 (4): 66–88.

Atack, Jeremy, and Fred Bateman. 1999. "Nineteenth-Century US Industrial Development through the Eyes of the Census of Manufactures." *Historical Methods* 32 (4): 177–88.

Atack, Jeremy, Robert A. Margo, and Paul W. Rhode. 2019. "'Automation' of Manufacturing in the Late Nineteenth Century: The Hand and Machine Labor Study." *Journal of Economic Perspectives* 33 (2): 51–70.

Atack, Jeremy, Robert A. Margo, and Paul W. Rhode. 2020. "'Mechanization Takes Command': Inanimate Power and Labor Productivity in Late Nineteenth Century American Manufacturing." NBER Working Paper 27436.

Beveridge, William. 1939. *Prices and Wages in England from the Twelfth to the Nineteenth Century, Vol. 1: Price Tables: Mercantile Era*. London: Longmans, Green, and Co.

Chamberlain, Joseph. 1866. "Manufacture of Iron Wood Screws." In *Birmingham and the Midland Hardware District, A Series of Reports,* edited by *Samuel Timmins*. London: Robert Hardwicke.

Cole, Arthur Harrison. 1938. *Wholesale Commodity Prices in the United States 1700–1861.* Cambridge, MA: Harvard University Press.

David, Paul A. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review* 80 (2): 355–61.

French, B.F. 1858. *History of the Rise and Progress of the Iron Trade of the United States from 1621 to 1857.* New York: Wiley and Halstead.

Gallman, Robert E. 1966. "Gross National Product in the United States, 1834–1909." In *Output, Employment, and Productivity in the United States after 1800, Studies in Income and Wealth,* Vol. 30, edited by the Conference on Research in Income and Wealth, National Bureau of Economic Research, 3–90. New York: Columbia University Press.

Gallman, Robert E., and Paul W. Rhode. 2019. *Capital in the Nineteenth Century.* Chicago: University of Chicago Press.

Goldin, Claudia, and Lawrence F. Katz. 1998. "The Origins of Technology-Skill Complementarity." *Quarterly Journal of Economics* 113 (3): 693–732.

Gordon, Robert J. 2009. "Apparel Prices 1914–93 and the Hulten/Bruegel Paradox." In *Price Index Concepts and Measurement,* edited by W. Erwin Diewert, John S. Greenlees, and Charles R. Hulten, 85–128. Chicago: University of Chicago Press.

Gordon, Robert J. 2018. "Why Has Economic Growth Slowed When Innovation Appears to Be Accelerating?" NBER Working Paper 24554.

Jacks, David S. 2013. "From Boom to Bust: A Typology of Real Commodity Prices in the Long Run." NBER Working Paper 18874.

Jackson, Kenneth T. 1987. *Crabgrass Frontier: The Suburbanization of the United States.* Oxford: Oxford University Press.

Jorgenson, Dale W., Kyoji Fukao, and Marcel P. Timmer. 2016. *The World Economy: Growth or Stagnation.* Cambridge: Cambridge University Press.

Lewis, Miles. 1998. "Australian Building: A Cultural Investigation, Section 8: Metals." www.mileslewis.net/australian-building/#eight (accessed June 29, 2011).

Mandel, Michael. 2011. "Mandel on Innovation and Growth: Air Force Certifies the Weakness of Domestic Manufacturing." innovationandgrowth.wordpress.com/2011/0418/air-force-certifies-the-weakness-of-domestic-manufacturing/.

Margo, Robert A. 2015. "Economies of Scale in Nineteenth-Century American Manufacturing Revisited: A Resolution of the Entrepreneurial Labor Input Problem." In *Enterprising America: Businesses, Banks, and Credit Markets in Historical Perspective*, edited by William J. Collins and Robert A. Margo. Chicago: University of Chicago Press.

Margo, Robert A., and Jeremy Atack. 2019. "Gallman Revisited: Blacksmithing and American Manufacturing, 1850–1870." *Cliometrica* 13: 1–23.

Maze Nails Company Online Catalog. 2020. "2020 List Price Guide." https://www.mazenails.com/assets/pdf/maze-nails-price-list-catalog.pdf (accessed 2020).

Measuring Worth. 2011a. "The Annual Consumer Price Index for the United States, 1774 to Present." https://www.measuringworth.com/datasets/uscpi/ (accesssed May and June 2011).

Measuring Worth 2011b. "Dollar-Pound Exchange Rate From 1791." https://www.measuringworth.com/datasets/exchangepound/. (accessed May and June 2011).

Nordhaus, William D. 1997. "Do Real Output and Real-Wage Measures Capture Reality? The History of Lighting Suggests Not." In *The Economics of New Goods*, Vol. 58, edited by Timothy F. Bresnahan and Robert J. Gordon, 27–70. Chicago: University of Chicago Press.

Nordhaus, William D. 2007. "Two Centuries of Productivity Growth in Computing." *Journal of Economic History* 67 (1): 128–59.

Rees, Albert, and Donald P. Jacobs. 1961. *Real Wages in Manufacturing: 1890–1914.* Princeton: Princeton University Press.

Rothenberg, Winifred. 1979. "A Price Index for Rural Massachusetts, 1750–1855." *Journal of Economic History* 39 (4): 975–1001.

Rozhon, Tracie. 1999. "Fred C. Trump, Postwar Master Builder of Housing for Middle Class, Dies at 93." *New York Times,* June 26. https://www.nytimes.com/1999/06/26/nyregion/fred-c-trump-postwar-master-builder-of-housing-for-middle-class-dies-at-93.html.

Rybczynski, Witold. 2000. *One Good Turn: A Natural History of the Screwdriver and the Screw.* New York: Scribner.

Smith, Adam. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations.* London: Strahan.

Temin, Peter. 1964. *Iron and Steel in the Nineteenth Century America.* Cambridge, MA: MIT Press.

United States Bureau of Labor Statistics. 1899. *Hand and Machine Labor, 13th Annual Report of the Commission of Labor 1898*, Vol. 1. Washington, DC: US Government Printing Office.

Warren, George F., and Frank A. Pearson. 1933. *Prices.* New York: John Wiley and Sons.

Wells, Tom. 1998. "Nail Chronology: The Use of Technologically Derived Features." *Historical Archaeology* 32 (2): 78–99.

Wilder, Laura Ingalls. 1935. *Little House on the Prairie.* New York: Harper Collins.

Journal of Economic Perspectives—Volume 36, Number 1—Winter 2022—Pages 151–176

The Puzzle of Falling US Birth Rates since the Great Recession

Melissa S. Kearney, Phillip B. Levine, and Luke Pardue

Between 1980 and 2007, US birth rates generally fluctuated within a narrow range of roughly 65 to 70 births per 1,000 women between ages 15 and 44. Since then, US birth rates have plummeted, reaching 55.8 per 1,000 women in 2020—about a 20 percent decline over 13 years. Figure 1 plots the trend in the US birth rates. The decline began at the onset of the Great Recession and continued during the ensuing recovery, with no signs of reversing.

This paper considers possible suspects behind the falling birth rates. We begin with a detailed look at birth rates by demographic groups defined by age, education, race and ethnicity, marital status, and birth parity. A detailed examination by group might offer some preliminary clues as to what types of factors might be responsible for the aggregate trend. While the decline is concentrated among women in the under-30 age group, the decline is generally widespread across demographic subgroups, which gives reason to suspect that the dominant explanation for the aggregate decline is likely to be multifaceted or society-wide. We see no indication in the data that there is likely to be a reversal of these trends in the near future.

■ *Melissa S. Kearney is the Neil Moskowitz Professor of Economics, University of Maryland, College Park, Maryland. Phillip B. Levine is the Katharine Coman and A. Barton Hepburn Professor of Economics, Wellesley College, Wellesley, Massachusetts. Luke Pardue is a PhD Candidate in Economics, University of Maryland, College Park, Maryland. Kearney and Levine are also Research Associates, National Bureau of Economic Research, Cambridge, Massachusetts. Their email addresses are kearney@umd.edu, plevine@wellesley.edu, and lpardue@umd.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.36.1.151.

Figure 1
Trend in US Birth Rates

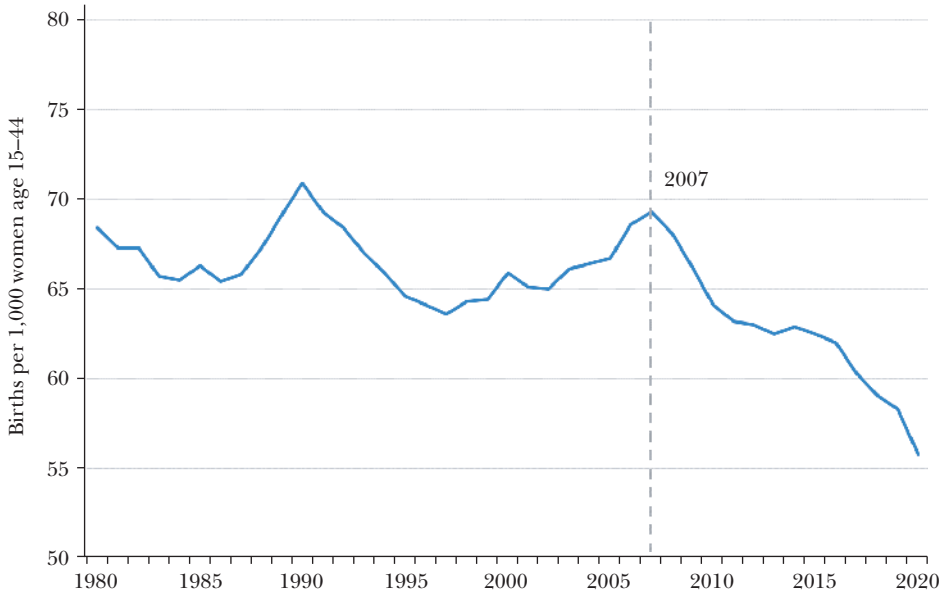

Source: Birth Rates collected from CDC Vital Statistics Births Reports for 2015, 2019 and 2020. See Data Appendix for additional details.

We next turn to an exploration of potential economic, social, and policy factors that might be responsible for the post-2007 decline in US birth rates. We begin with a brief overview of the economics of fertility and the framework that economists typically use to model and study the "demand for children," which is the individual decision underlying aggregate birth rates. We then describe the empirical relationship between annual state-level birth rates and economic and policy factors that vary at the state and year level, including labor market conditions, social policy indicators, and reproductive health policy measures. After that, we consider the impact of a set of slower-moving factors, like women's economic status, changing take-up of contraceptive technology, and the cost of raising children. Aside from the impact of the Great Recession, which contributed to the decline for the first few years of this period, we are unable to identify a strong link between any specific policies or economic factors and the declining birth rates.

We also compare birth trends in the United States to other highly developed countries to examine whether international differences in social, economic, and policy environments hint at a likely cause. The fact that birth rates are also relatively low in other high-income countries supports the notion that localized factors may not explain a significant portion of the decline.

If period- and location-specific factors generally cannot explain declining birth rates, perhaps the cause has to do with changes in the cohorts of women moving through their childbearing years. *Shifting priorities* among more recent birth cohorts—potentially driven by changes in preferences for having children, aspirations for life, and parenting norms—would represent a more universal, harder-to-quantify factor that may be the key driver of the decline in birth rates in the United States (and elsewhere). This line of explanation is potentially related to a concept referred to by demographers as the "second demographic transition." Our conclusion briefly considers the societal consequences for the United States of a declining birth rate—such as reduced productivity growth and instability in the finances of programs to support the elderly like Social Security and Medicare—and what might be done about it.

Recent Trends in US Birth Rates

We begin by providing a descriptive examination of birth rate trends in the United States, using data on the universe of US births from the Vital Statistics system from 1980 through 2020.[1] Our examination focuses on the steady, dramatic decline in birth rates since 2007, but showing data from this longer period helps put the recent decline in context. We describe trends in birth rates for different demographic groups and then evaluate how much of the decline in the total birth rate since 2007 reflects changing population demographics versus changes in birth rates within certain demographic groups.

Birth Rates by Demographic Group

We take advantage of the demographic data included in the publicly available Vital Statistics natality data to document birth rates by maternal age, race and ethnicity, nativity, and marital status.[2] In Figure 2, panel A plots birth rates for six different age groups. Birth rates have declined most dramatically among teenagers, but the downward trend in births among teens began well before the broader decline in births to non-teens. The teen birth rate peaked in 1991 at a level of 61.8 births per 1,000 women aged 15 to 19. The pace of that decline slowed approaching 2007 and then accelerated in the following years. Overall, teen births fell to 41.5 by 2007

[1] At the time of writing, aggregate statistics on births are available from the National Center for Health Statistics through 2020; individual-level microdata is available through 2019. For demographic groups that are not included in the aggregate reports, we tabulate data through 2019 using the microdata.

[2] We combine Vital Statistics birth data with population data from the Surveillance, Epidemiology, and End Results (SEER) program at the National Cancer Institute to generate rates of birth per 1,000 women between the ages of 15 and 44 by race and ethnicity. Since SEER data does not report population by educational attainment, marital status, and nativity, we use data from the Current Population Survey to estimate the number of women in each relevant group in order to construct those birth rates.

Figure 2
Trends in Birth Rates by Population Subgroup

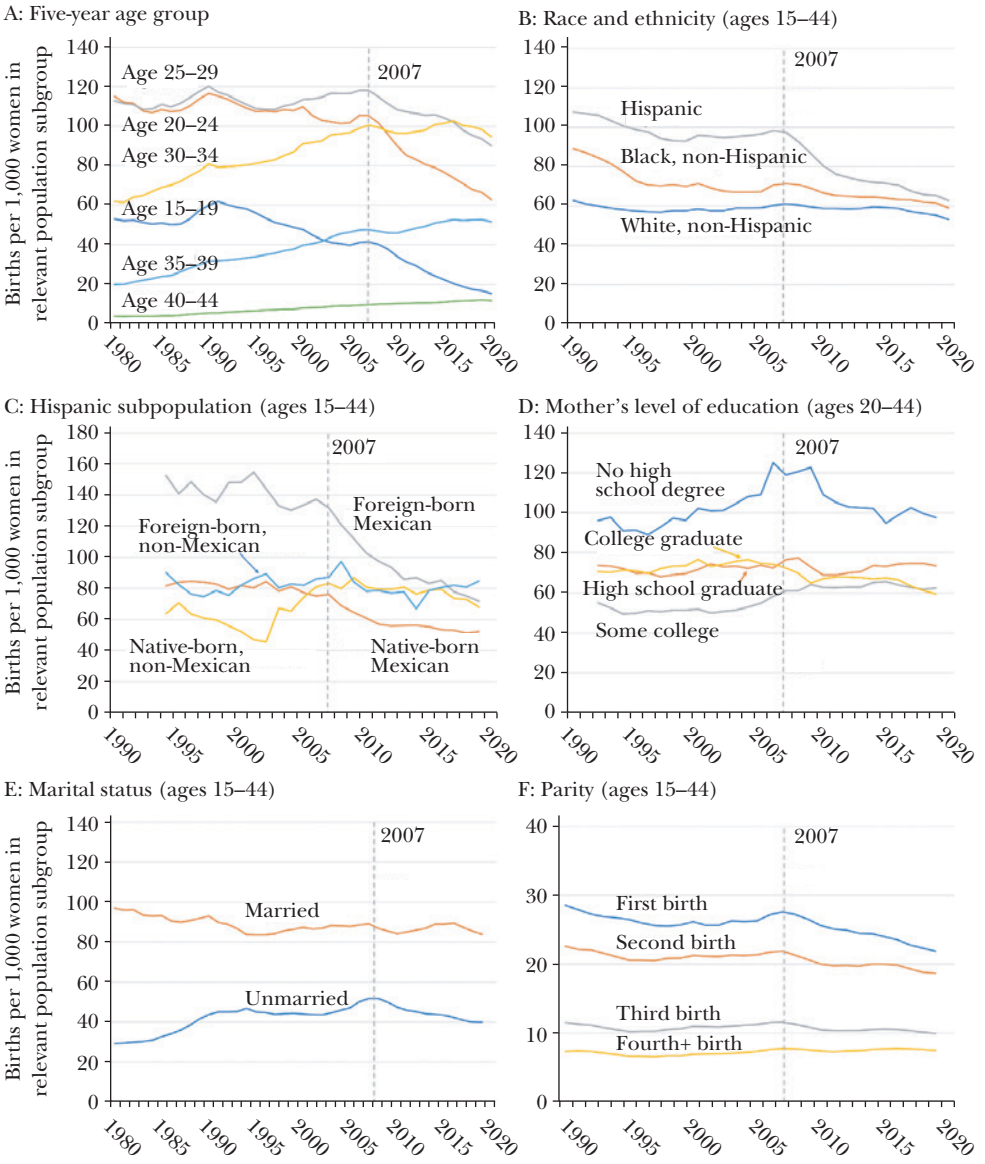

Note: Birth rates by age group, race and ethnicity, and marital status are gathered from CDC Vital Statistics Births Reports. Birth rates by Hispanic subpopulation, mother's education, and by parity are calculated by the author's using NCHS Vital Statistics birth microdata, SEER population data, and the Current Population Survey. The Data Appendix provides detailed information on the specific data sources. .

and then to 15.3 births per 1,000 teen women in that year.[3] The overall decline was 75 percent, representing a massive change in a frequently tracked social outcome (as discussed in this journal by Kearney and Levine 2012).

Other than teens, the main decline in birth rates was women in their 20s. Birth rates among women 20 to 24 fell from 105.4 to 62.8 per 1,000 women through 2020 and birth rates among women 25 to 29 fell from 118.1 to 90.0 per 1,000 women. Births to women at older ages remained constant or rose, but not nearly enough to make up for these large declines at younger ages. These trends are consistent with women having fewer children over their childbearing years, not merely delaying childbearing to older ages (Kearney and Levine 2021).

We next examine birth rates by race and ethnicity, as shown in panel B of Figure 2. Hispanics have experienced the most dramatic recent declines in birth rates. In 2007, the birth rate among Hispanic women was 97.4; it fell to 62.8 by 2020. Birth rates for Black and White non-Hispanic women also fell, but by much smaller amounts. When the Great Recession hit, birth rates differed dramatically by race and ethnicity. By 2020, racial and ethnic differences in birth rate levels remain, but they have become much smaller.

Assimilation offers one possible explanation for the falling birth rate among Hispanic women (Tavernise 2019), if birth rates among Hispanic women converge to those of native non-Hispanic US women over time and generations. The share of Hispanic women of childbearing age who are native-born as opposed to foreign-born has increased from 49.3 percent in 2007 to 61.7 percent in 2018, according to our calculations from the American Community Survey. Parrado and Morgan (2008) consider birth cohorts from the 1835–1839 through the 1960–1964 period, and show that successive generations of Hispanic women, in general, and particularly Mexican women, have birth rates that converge to those of non-Hispanic White women.

In panel C of Figure 4, we look separately at birth rates among native- and foreign-born Hispanics, and further segment Hispanics into country of origin. Among those with Mexican heritage, birth rates are considerably higher among foreign-born women than among native-born women, which is descriptively consistent with assimilation along this dimension. However, birth rates among both groups have fallen rapidly suggesting that something more than just assimilation and an increase in the native-born share of Hispanics is behind the fall in Hispanic birth rates. It is relevant to note that the birth rate in Mexico has fallen dramatically over the past 50 years; it is now only slightly higher than in the United States (World Bank 2021). Birth rates have not changed much over time for Hispanic women not of Mexican origin in the United States, regardless of their nativity.

We next examine birth rates by four different maternal education groups: less than high school degree, high school degree, some college, a four-year college degree or more. We drop teens from this analysis since many will not have

[3] Kearney and Levine (2015b) provide evidence that the introduction of the MTV show *16 and Pregnant* contributed to the more recent decline after it was introduced in 2009.

completed their education. Panel D of Figure 2 shows that recent declines in birth rates are largest for the most and least educated women, those with a four-year college degree (36.3 percent of women aged 20 to 44 in 2018) and those without a high school degree (8.1 percent of those women). Birth rates fell from 72.5 to 59.4 per 1,000 between 2007 and 2019 for college-educated women and from 119 to 97.5 per 1,000 for women without a high school degree. Births to this latter group rose by about the same level in the decade before that, bringing recent levels of births back in line with that 25 years ago. Both women with a high school degree and those who attended some college have had fairly stable birth rates since 2007.

Panel E of Figure 2 separately plots births by maternal marital status. Between 1980 and the early 1990s, birth rates for married women were falling and birth rates for unmarried women were rising. What is relevant for our purposes is that since 2007, birth rates have not trended very differently for married and unmarried women. Births to unmarried women fell somewhat more (perhaps attributable to the large decline in teen births, which are almost exclusively to unmarried women), but births to married women fell as well, albeit not continuously.

Although birth rates in each marital status category have fallen by only a small amount, the percentage of women of childbearing age who are unmarried is growing (58.2 percent in 2007 to 63 percent in 2018). Since unmarried women have lower birth rates, the rise in the share of women unmarried would reduce the overall birth rate purely based on this compositional change. The fact that women now marry at older ages contributes to greater numbers of unmarried women among those of childbearing age. The median age at first marriage has risen continuously over the past 50 years, although perhaps at a faster rate more recently, rising from 22.0 years in 1980 to 25.6 in 2007 to 28.1 in 2020 (US Census Bureau 2021).

Finally, we tabulate trends in births by "parity," referring to the number of children for a given woman. These data reveal that the post-2007 decline in births is driven more by a decline in initial childbearing (first births) than by women not having larger families (third and higher order births). First births declined the most, from a rate of 27.6 per 1,000 women of childbearing age to 21.9 per 1,000 women, a drop of 5.7 births. Second births declined from 21.9 to 18.7, a drop of 3.2 births. The trend lines for third and higher order births are much flatter over this period. These data are consistent with a trend towards childlessness (Stone 2020b).

Decomposing the Decline in Birth Rates into Between- and Within-Demographic Groups

The contribution of any particular demographic group to the overall decline depends on both changes in birth rates for that group and that group's share of the female population of childbearing age. For instance, birth rates fell dramatically among Hispanic women, but they only represent 16.7 percent of the overall population of women of childbearing age. In a mechanical sense, the massive decline in their birth rates is diminished somewhat in explaining the overall decline because they are not a large population subgroup. Larger groups with a more modest decline may have contributed to the aggregate decline just as much or more. Furthermore,

even if birth rates did not change for any of the demographic groups, if the population shifted toward groups that traditionally have lower birth rates, aggregate birth rates would decline.

In this section, we decompose the post-2007 decline in the aggregate birth rate to declines within demographic groups and changes in demographic group's population shares. We also identify which demographic groups have contributed the most to the overall decline, either because they experienced a very large decline in their birth rate or because they are a particular large share of the population. Specifically, we decompose the overall change into the contribution of changing within-group birth rates, the contribution of changes in group population shares, and the interaction of a group's changing rates and changing population shares.[4]

Our decomposition is based on categorizing the population of women of childbearing age into demographic groups defined by the interaction of three race/ethnic groups (White non-Hispanic, Black non-Hispanic, and Hispanic), six five-year age groups (15 to 19 through 40 to 44), and for women over 20, four education levels (no high school degree, high school degree, some college, and college graduate). This breakdown results in 63 subgroups: three race/ethnicity groups, six five-year age groups, and four education levels would total 72, but we omit education levels for teens. We focus on explaining changes in birth rates between 2007 and 2019, the period of rapidly declining fertility.

This exercise makes two points clear. First, changing birth rates within demographic groups is responsible for the declining birth rate since 2007, not changing population shares. From 2007 to 2019, the birth rate declined by 10.8 births per 1,000 women 15 to 44 (from 69.1 to 58.3).[5] Across all groups, had birth rates been constant and only population shares shifted between 2007 and 2019, the birth rate would, in fact, have *risen* by 2.6 births per thousand. On the other hand, if population shares were held constant and only within-group birth rates moved over that period (the change captured by the first term), the overall birth rate would have fallen by 12.8 births per 1,000 women.

Second, this decomposition highlights the importance of the relative size of a demographic group when accounting for the overall decline in the birth rates. Table 1 reports the 8 out of 63 demographic groups that contribute the most to the declining birth rate. These groups account for 34 percent of the population, but changes in their birth rates explain 75 percent of the overall decline. The three

[4] The overall change in the birth rate can be written as:

$$\Delta\left(\frac{B}{P}\right)_{t_o,t_1} = \sum_i s_{i,t_o}\Delta\left(\frac{B}{P}\right)_{i,t_o,t_1} + \sum_i \left(\frac{B}{P}\right)_{i,t_o}\Delta s_{i,t_o,t_1} + \sum_i \Delta s_{i,t_o,t_1}\,\Delta\left(\frac{B}{P}\right)_{i,t_o,t_1}$$

where B is the number of births, P is population, s is the share of overall population, i indexes the 63 groups, and t_0 and t_1 are the beginning and ending years.

[5] Note that these values are slightly different than those reported in Figure 1 because the data used for the decomposition include women whose race is White or Black and omits the small number of births to women categorized in the vital statistics as having a race of "other"; the birth numbers used in Figure 1 include births to women of "other" race, generating a slight discrepancy in birth rates.

Table 1
The Eight Demographic Groups that Contributed the Most to the 2007–2019 Decline in the US Birth Rate

Group	*Relative contribution to declining birth rates*	*2007 share of population*	*2007 birth rate*	*2019 birth rate*	*2007–19 change in birth rate*
Age 15–19, Hispanic	14.0%	3.1%	82.2	24.7	−57.5
Age 15–19, White NH	13.9%	11.0%	27.1	11.0	−16.1
Age 25–29, White NH, College Grad	11.9%	4.2%	101.1	65.1	−36.0
Age 15–19, Black NH	8.8%	2.8%	65.1	25.3	−39.8
Age 20–24, White NH, HS Grad	7.6%	2.8%	139.9	105.3	−34.6
Age 20–24, White NH, Some College	7.1%	5.5%	54.2	37.7	−16.5
Age 30–34, White NH, College Grad	6.4%	4.2%	131.7	112.0	−19.7
Age 20–24, Hispanic, Less than HS	5.3%	0.8%	295.1	206.8	−88.3
Total	75.1%	34.%	76.5	50.5	−25.9

Source: Author's calculations based on Vital Statistics Natality data, as described in the Data Appendix. NH stands for non-Hispanic.
Note: Birth rates are measured as the number of births per 1,000 women in each population subgroup.

teen categories by race/ethnicity explain 37 percent of the overall decline. Hispanic teens contributed the largest share, explaining 14 percent of the overall decline; their birth rate fell dramatically, from 82.2 to 24.7 over the period.

Other demographic groups with smaller declines in their birth rate also contributed extensively to the overall decline because of their relatively large population shares. For instance, the third-largest contributing group is White women between the ages of 25 and 29 with college degrees; their birth rate fell from 101.1 to 65.1, accounting for 11.9 percent of the overall decline. This group represents 4.2 percent of women; with 63 separate categories, the average group comprises 1.6 percent of the female population of childbearing age.

Note that five of the leading groups contributing the most to declining birth rates are subsets of White, non-Hispanic women. Figure 2, though, shows no dramatic decline in births among that broader group. In this decomposition analysis White, non-Hispanic women are broken down into 21 subgroups. The apparent discrepancy in results is attributable to the fact that births were relatively constant or slightly increasing (mostly for older White women) in the other 16 subgroups. This is confirmed by an examination of trends in birth rates for each specific subgroup (not shown here).

How Economists Model Fertility

Before evaluating some evidence concerning the drivers of recent changes in birth rates, we discuss the general economic framework for thinking about fertility. Starting with the seminal work of Gary Becker (1960), economists have viewed

the decision to have a baby (sometimes referred to as the "demand for children") using the framework of constrained utility maximization. This approach recognizes that children bring people utility—perhaps in the form of life satisfaction, general happiness, or pleasurable experiences—but that children also come with associated costs, broadly defined, including both time and money. Becker (1960) also introduced the concept of child "quality," a term he uses to refer to expenditures per child, but which he carefully specifies does not mean "morally better." He explains "a family must determine not only how many children it has but also the amount spent on them—whether they should provide separate bedrooms, send them to nursery school and private colleges, give them dance or music lessons, and so forth" (p. 211).

This approach to modeling the decision to have a child leads to standard predictions of price and income effects. Sometimes the price effect of children is direct, like the costs of housing and childcare. Dettling and Kearney (2014), for example, show that birth rates decrease for renters when housing prices increase, which is consistent with a negative price effect because housing is a large cost associated with having children. Other times the price effect is indirect, like the opportunity cost of a woman's time, which would increase along with women's wages or a greater likelihood of finding employment during an economic expansion. Holding prices and quality constant, an increase in income will lead people to choose to have more children. In the vernacular of consumer demand, children are "normal" goods (jargon that is also unfortunate in this context). This positive relationship between income and births may come as a surprise to some readers, given the negative relationship between birth rates and income or per capita GDP observed over time and place.

There are a few potential explanations for this apparent contradiction within standard economic models of fertility. First, correlational observations are often plagued by a potential conflating of income and price effects. As economic development increases income, it also tends to increase the price of children, in terms of the price of housing, childcare, and the opportunity cost of women's time. Such correlations could also reflect the effect of confounding selection effects. Perhaps birth rates are lower in high-income cities because the cost of living is higher and/or because people who choose to live there desire more adult-centric amenities like restaurants and bars.

Second, smaller families among higher income people—either over time or across place—could reflect a "quantity-quality" trade-off (another unfortunate label that is common in economics jargon). The idea is that as societies become richer, parents may opt to have fewer children and spend more per child, investing in greater "quality," say, through expenditures on education and enrichment (Becker and Lewis 1973).

Economic models also lead to opposite-signed predictions about the effects of male and female wages on birth rates. The seminal work of Butz and Ward (1979) predicted that an increase in male earnings will lead to an increase in the total demand for children, but an increase in female wages will have both positive

income effects and negative price effects on fertility. The "baby boom" of the 1950s is broadly consistent with their predictions of increasing births in response to rising male earnings. The "baby bust" of the 1960s is broadly consistent with their prediction of increases in female earnings leading to fewer births. More recent work by Schaller (2016) considers the period 1980–2009 and documents that exogenously determined improvements in men's labor market conditions lead to increases in birth rates, while exogenously determined improvements in women's labor market conditions lead to small decreases in birth rates.

Apart from the question of how many children to have, parents also face the decision of *when* to have them. There is robust empirical evidence showing that aggregate birth rates tend to be pro-cyclical (Schaller 2016; Dettling and Kearney 2014; Kearney and Levine 2020). This is consistent with the notion that people are more likely to become parents when they have income available to pay for the associated costs of childbearing. If credit markets were perfect, parents could borrow and save to finance the cost of children and optimally choose when to have them. But credit markets are imperfect, and many people are liquidity constrained; couples might thus refrain from having a child at times when their income is low—that is, when the economy is weak.

In the uncertain context of pregnancy and childbearing, economic models also can incorporate the fact that optimized "choices" are not always realized (Buckles, Guldi, and Schmidt 2019). The availability, price, and efficacy of contraception, as well as the degree of access to abortion providers, will all affect the degree to which women are able to achieve their desired level of pregnancy and birth avoidance.[6] There is ample evidence from recent US contexts that expanded access to affordable and efficacious contraception has led to a reduction in births among affected populations (for example, Kearney and Levine 2009; Bailey 2010; Lindo and Packham 2017; Kelly, Lindo, and Packham 2020).

Finally, economic models tend to take preferences as given. However, people's preferences for having children or spending their resources investing in children might change over time. Secular changes in attitudes and aspirations, religiosity and family attachment, and other societal changes could all lead to changes in preferences and the demand for children (Adserà 2013). This is a point to which we will return later.

Potential Explanations for Declining US Birth Rates

Next, we turn to potential empirical explanations for the decline in US birth rates, focusing on economic, policy, and social factors, and distinguishing them by the time horizon over which changing factors occur and how long it would take for behavioral change to be observed. We also compare trends in birth rates in

[6] See Levine (2004) for a detailed presentation of this form of decision-making in the context of changes in abortion policy.

the United States to those in other high-income countries to see how international differences in economic and policy environments may contribute to differences in birth patterns over time.

The Great Recession and Birthrates

An array of empirical evidence from a variety of sources suggests that a recession will cause birth rates to fall for a time. Some of this evidence comes from studies of how changes in income affect the number of children with evidence from a variety of contexts, including the case of individual job loss (Lindo 2010), shocks to area-level earnings and income (Kearney and Wilson 2018; Black et al. 2013), and shocks to the housing market that increase owners' housing wealth and equity (Dettling and Kearney 2014; Lovenheim and Mumford 2013).

Figure 1 reveals a noticeable drop in birth rates after the recessions of the early 1980s and the 1990–91 recession, as well as after the 2007 recession (although there is not much change in birth rates after the mild recession of 2001). The economic stress of the Great Recession surely contributed to the abrupt downturn in birth rates after 2007. Based on the 5-percentage point increase in the unemployment rate from 2007 to 2010 (from 4.6 percent to 9.6 percent), our analysis described below indicates that one could have expected births to fall by 3.5 percent between 2008 and 2011 (approximating a nine-month gestational lag). Over that period, the birth rate fell 7.2 percent, from 68.1 to 63.2 per 1,000 women. Although the recession clearly contributed to that decline, other factors must have also been at play. The lack of any rebound in births and, in fact, their continued decline following the end of the recession further suggests a role for factors beyond the Great Recession.

Beyond the Great Recession

There are a substantial number of economic and policy factors that plausibly may affect birth rates in one way or another. We begin our investigation with a state-level approach to look at possible economic and policy determinants of overall birth rates from 2001 to 2019. This empirical approach relies on the presumed exogeneity of state-level policy changes to interpret the reported relationships as causal. The regression model controls for year fixed effects (to account for changes in birth rates over time that are not state-specific), as well as state fixed effects (to account for persistent differences across states in average birth rates).

We focus on sets of operational factors that fit with the economic approach to modeling fertility described above. One key factor is the unemployment rate; we described those results above. We also consider an extensive set of relevant social policies, many of which have been separately examined in previous studies of birth rates. These include the generosity of welfare benefits (Moffitt 1998; Grogger, Karoly, and Klerman 2002; Lopoo and Raissan 2012; Ziliak 2016), the state minimum wage (Bullinger 2017), and child support enforcement (which affects the opportunity cost of fathering a child, see Aizer and McLanahan 2005). We also include a number of reproductive health policies that potentially affect a

woman's ability to achieve her desired fertility, including abortion restrictions in the form of parental notification laws and waiting periods (Levine 2004), health insurance coverage through Medicaid (DeLeire, Lopoo, and Simon 2011), mandatory coverage of contraception in private insurance plans, and state mandatory sex education and mandatory contraception instruction laws (Paton, Bullivant, and Soto 2020). We hold constant the demographic composition of female adults in each state and year (specifically, the share White, Black, Hispanic, married, and in four different education groups).

Details regarding the exact policies, source data, variable construction, and results are available in the Data Appendix. The key result is that when we sum the estimated coefficients on our ten economic/policy variables with their average change between 2007 and 2018, their combined effect is 6.2 percent of the total decline in the birth rate from 69.1 to 58.3 births per 1,000 women age 15 to 34 between 2007 and 2018. Of course, it is quite possible that some of these factors affect the birth rates of targeted groups of women in some contexts (for example, births among lower-income and teen women might fall when subsidized contraception becomes more readily available, as past research has shown); the point here is simply that any effect they might have is too small—or on too small a group of women—to explain a sizable share of the total change in births.

This analysis—and the finding that specific policy and economic factors have limited explanatory power—is similar in approach to a previous study two of us published on trends in teen births between 1981 and 2010 (Kearney and Levine 2015a). The birth rate for teens began falling in 1991, much sooner than the overall decline in birth rates beginning in 2007. Our earlier analysis found that the only three factors that had a statistically significant relationship with state/year teen birth rates were the unemployment rate, the maximum welfare benefit amount, and the implementation of a Medicaid family planning expansion waiver. However, none of these factors could explain a large share of the overall decline in teen births, either individually or collectively. The weak results of both analyses make it difficult to determine whether the decline for teens since 2007 was due to continuing factors unique to them or factors common to older women of childbearing age.

Returning to the decline in overall births post-2007, we augment the preceding state-year analysis by examining the potential impact of slower-moving forces that might not change birth rates year-to-year but might have a meaningful effect over a longer period of time. The factors we consider have been suggested by observers as possibly important contributors to the recent decline in US birth rates. Some examples include more widespread usage of long-acting reversible contraception, costs associated with raising children (like housing and child care costs), improvements in women's economic position (which would increase the opportunity cost of women's time), and rising student debt burdens (which would reduce adults' level of disposable income). Popular press articles in outlets including the *Washington Post, Wall Street Journal, New York Times, Vox, Business Insider,* and *CNBC*, among others, have suggested that these factors played an important role in the

Figure 3
Change in Birth Rates by State, 2004–2008 to 2015–2019

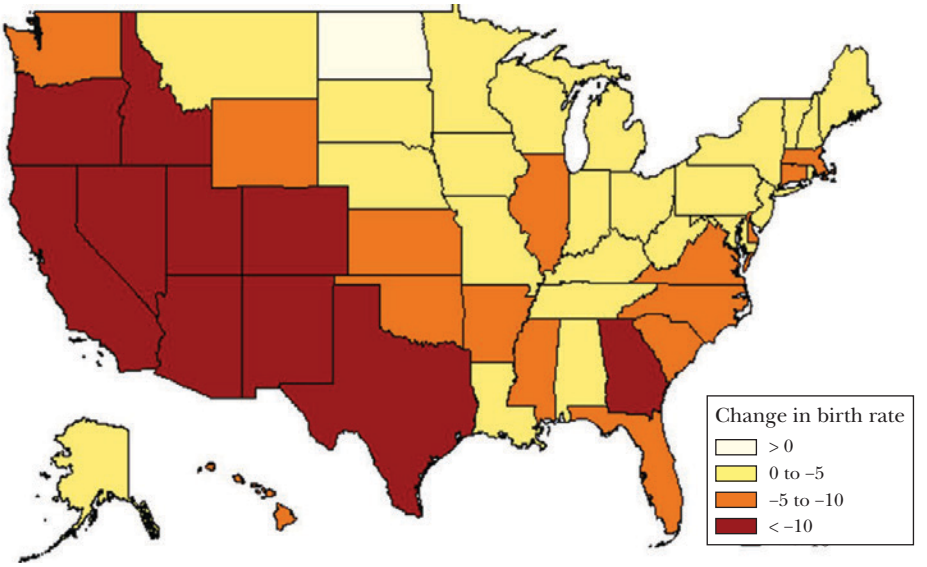

Source: Birth data are from NCHS Vital Statistics. Population data are from CDC Surveillance, Epidemiology, and End Results (SEER) program.
Note: Birth rates are calculated among women aged 15 to 44.

decline.[7] Others have advanced the notion that declining religious observance might affect preferences toward and attitudes about having children (Adserà 2013; Stone 2018; Douthat 2020).[8]

Establishing a causal link between these factors and changes in birth rates would require identifying exogenous variation in factors such as contraceptive take-up and childcare prices, which is beyond the scope of this paper. But basic descriptive evidence suggests that such effects are not likely to explain the large and extended fall in birth rates.

The map in Figure 3 displays changes in state-level birth rates between two five-year periods, 2004–2008 to 2015–2019 (before and after the Great Recession), which we subsequently relate to state-level changes in relevant economic policy and social factors. This grouping of years avoids the confounding influence of state-level variation in the severity of the great recession. Averaging the birth rates over these

[7]Examples of press reports that mention these factors include the following: contraception (Iati 2019; DeBarros and Adamy 2019); the cost of raising children (Miller 2018; Belluz 2020); women's economic advancement (Hoffower 2021; Tavernise 2021), and student debt (Dickler 2018; Snodgrass 2021).
[8]Stone (2018) examines a number of the hypotheses we address here as well. We extend and update that analysis incorporating their ability to explain state-level variation in birth rates over time.

periods reduces the random, year-to-year variation in these data that is present, particularly in smaller states.

The decline in birth rates has been widespread across the country. Birth rates fell in every state over this period, except for North Dakota. One possible explanation for the increase in North Dakota birth rates is the fracking boom that occurred in this state over those years, which has been shown in other research to increase the birth rate (Kearney and Wilson 2018). But as can be readily seen in the map, there is substantial variation in the extent of the decline across places. Births fell the most in the South, in the West, and in the Southwestern and Mountain states. However, the set of states that experienced larger declines is varied, also including some Midwestern and New England states, notably Connecticut, Illinois, and Massachusetts.

Births fell the most in the Southwestern and Western states. The sizable Hispanic population in much of this region is consistent with the particularly large decline in births among Hispanic women, driven by a decline in births among both native and foreign-born Mexicans. The fact that other states with smaller shares of Hispanic residents (like Georgia and Oregon) also experienced large declines, though, further clarifies the broad-based nature of the decline.

Figure 4 presents six scatter plots showing the two-way relationship between the change in each considered factor (on the X-axis) and the change in birth rates (on the Y-axis) over these two five-year periods. Each marker represents a state, labeled with its two-letter abbreviation. The Data Appendix provides details regarding the sources of these data.[9]

The percentage of sexually active women who report using long-acting reversible contraception (LARC) increased from 5.5 percent in 2004 to 10.7 in 2017 and could have contributed to declining birth rates. The simple correlation, though, between the percentage point change in LARC usage in a state and the change in birth rates is wrong-signed (that is, positive), albeit close to zero. This suggests that take-up of LARCs has likely not played an important role in explaining the decline in the aggregate birth rate over this period.[10]

Annual expenditures on childcare for families with children under 12 who report positive childcare spending rose nationwide from $5,020 in 2009 to $7,190 for the average of 2015–2019, based on data from the Current Population Survey (all dollar values throughout are measured in constant 2019 dollars). But in the cross-state correlations, places where childcare expenditures increased more did not experience a noticeable drop in birth rates.

[9] Not all explanatory variables are available for all years. The Data Appendix includes details about these explanatory variables, including the years for which they are available.

[10] This comment does not imply that expanding access to LARCs among low-income teens or young women would not lead to a reduction in births for them. Research by Lindo and Packham (2017) and Kelly, Lindo, and Packham (2020) provide evidence that expanded access to LARCs in Colorado through the Family Planning Initiative led to a reduction in birth rates among teens and low-income women in that state.

Figure 4

Relationship between Changes in Birth Rates and Potential Explanatory Factors: 2004–2008 to 2014–2019

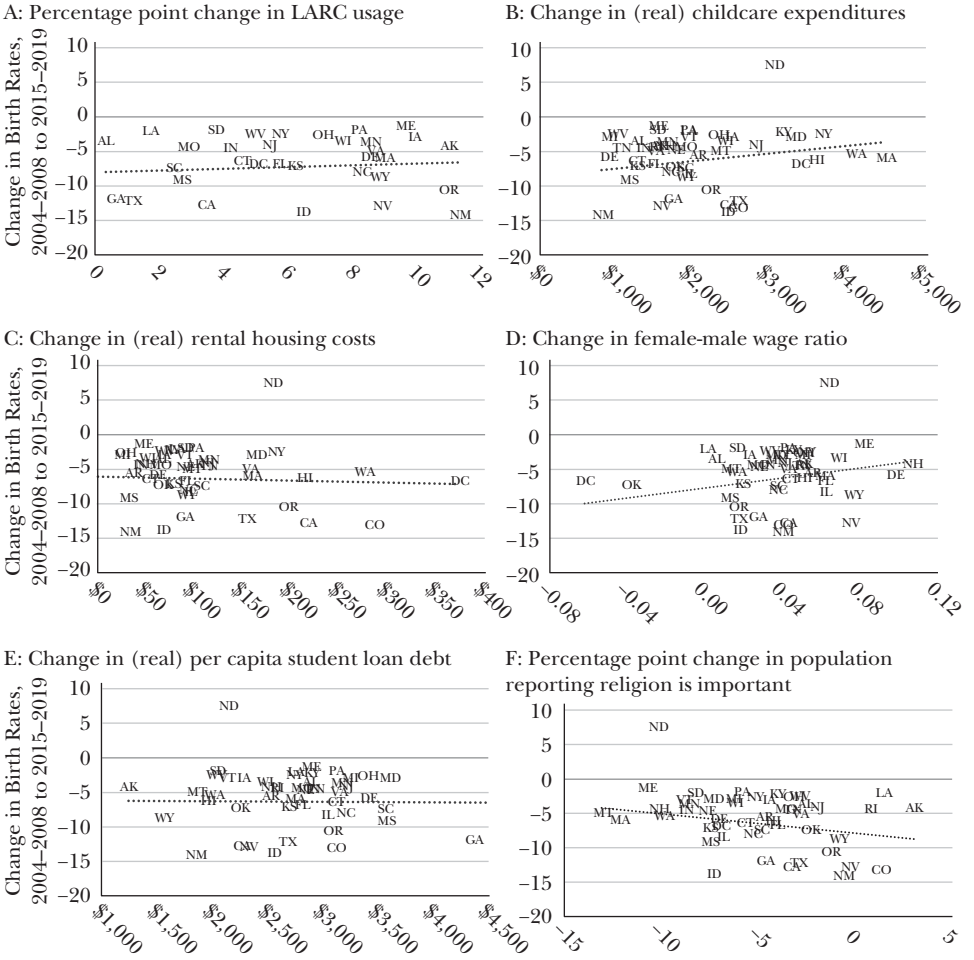

Note: Birth rates are calculated by the authors using NCHS Vital Statistics Natality Data and SEER population data. The Data Appendix provides detailed information on the data sources and variable construction for the six explanatory factors considered in panels A through F.

Average monthly rents for a two- to three-bedroom apartment rose $124 per month (from $930 to $1,060 in 2018 dollars, a 14 percent increase) nationwide over this period. The increase was much larger in some states, like Colorado, Washington, and the District of Columbia. The data, though, do not indicate a negative

association between state-level changes in rents and state-level changes in birth rates.[11]

As one way to measure changes in women's economic opportunities over time, we used the female-to-male median earnings ratio among full-time, year-round workers; our estimates indicate that this rose from 0.80 in the earlier period to 0.84 percent in the later period.[12] States in which the wage ratio rose the most over this sample period did not exhibit a greater decline in the birth rate.

If young adults are saddled with debt, they might not feel like they have sufficient disposable income to have a child or more children. We consider the total level of student debt per capita in a state, which has increased from $2,500 to $5,400 (in 2018 dollars), on average, between the earlier and later periods. The relationship between state-level student debt and the birth rate is generally flat, giving no indication that increases in student debt are related to the aggregate reductions in birth rates.

Finally, we looked at survey data that tracks the percentage of a state's population who report that religion is at least somewhat important to them. This statistic fell from 83 percent to 78 percent between 2007 and 2014. Again, despite the national trend, we see no evidence that states where religiosity declined the most experienced a greater relative decline in birth rates. If anything, the relationship goes the other way.

In short, other than the Great Recession itself as a triggering event for the deeper and more lasting persistent changes in birth rates, it is difficult to find prima facie evidence for other economic or policy factors that offer a plausible explanation. One should probably not be surprised that none of these factors explain much of the decline in the aggregate birth rate based on the timing of the decline. Births clearly dropped beginning in 2007, as shown in Figure 1. For any factor to have explained much of that decline, it would have had to change dramatically around the same time. Aside from the Great Recession, none of these factors exhibit that property.

A Comparison of US Birth Rates to Birth Rates in Other High-Income Countries

The total fertility rate is an estimate of the total number of children the average woman will have over her lifetime, based on age-specific birth rates at a given point in time. A total fertility rate of 2.0 means that a woman is expected to have two children on average; 2.1 is generally regarded as the rate required for population replacement (incorporating small amounts of mortality between birth and reproductive age). Comparable country-level statistics on the total fertility rate are available from the World Bank Database.

[11] When we weight these state-level observations by state population, the specific slopes change, but none provide statistically significant estimated relationships consistent with falling birth rates.

[12] We calculate this statistic among workers between the ages of 25 and 54. We obtained similar results using the female employment-to-population ratio and a measure of occupational prestige.

Through the 1990s and early 2000s, US births were at roughly replacement level throughout this period, but the total fertility rate was typically lower in other high-income countries. For instance, in the year 2000, the total fertility rate was 1.89 births per woman in the United Kingdom, 1.67 in Canada, 1.52 in the European Union, and 1.37 in Japan. Even after the US birth rate fell from 2.12 in 2007 to 1.73 births per soman in 2018 (the most recent year for which the World Bank statistics are currently available), the US total fertility rate is higher than that in the United Kingdom, 1.68; Canada, 1.50; European Union, 1.54; and Japan, 1.42. Even in Scandinavian countries specifically, with their especially generous system of public support, the total fertility rate is lower than in the United States; in 2018 it was 1.56 births per woman in Norway, 1.76 in Sweden, and 1.41 in Finland.

The fact that the US total fertility rate is now closer to other high-income countries, though generally still slightly higher, does not fit with the narrative that if the United States had more supportive government programs—such as subsidized childcare and generous paid work leave—its birth rates would be higher. One cannot prove this counterfactual, of course: perhaps if the United States had a more robust system of child and work supports, as in Scandinavian countries, then perhaps the US birth rate would have stayed elevated. It is always difficult to make comparisons across countries because policy context is not all that differs; societal norms and cultural preferences also tend to differ. Still, the international comparisons combined with the difficulty of finding policy and economic factors to explain the sustained decline in US birth rates suggest that these factors are not driving the changes in US birth rates.

The Role of Cohort Effects in Explaining the Decline in Annual Birth Rates

Falling Birth Rates across Recent Cohorts of Young Adults

Our empirical analyses described above do not uncover a readily identifiable, contemporaneous cause of declining births. That leads us to speculate that perhaps the key explanation for the post-2007 sustained decline in US birth rates is not about some changing policy or cost factor, but rather shifting priorities across cohorts of young adults. In other words, perhaps what we are seeing are cohort effects, rather than period effects. Period effects reflect changes that affect everyone at a point in time, whereas cohort effects reflect changes across people born or raised in different years.

To assess the contribution of cohort effects to the observed declines in birth rates, we examine birth data organized by mother's birth year. The impact of mothers' birth cohort on annual birth rates begins in earnest when those women hit their prime childbearing years, say between the ages of 20 and 24. Women between these ages in 2007 were born between 1983 and 1987. In Figure 5, we track the average number of children ever born by specific ages to women in five-year birth intervals from the 1968–1972 to 1993–1997 birth cohorts. These are the birth

Figure 5
Children Ever Born by Mother's Age and Birth Cohort

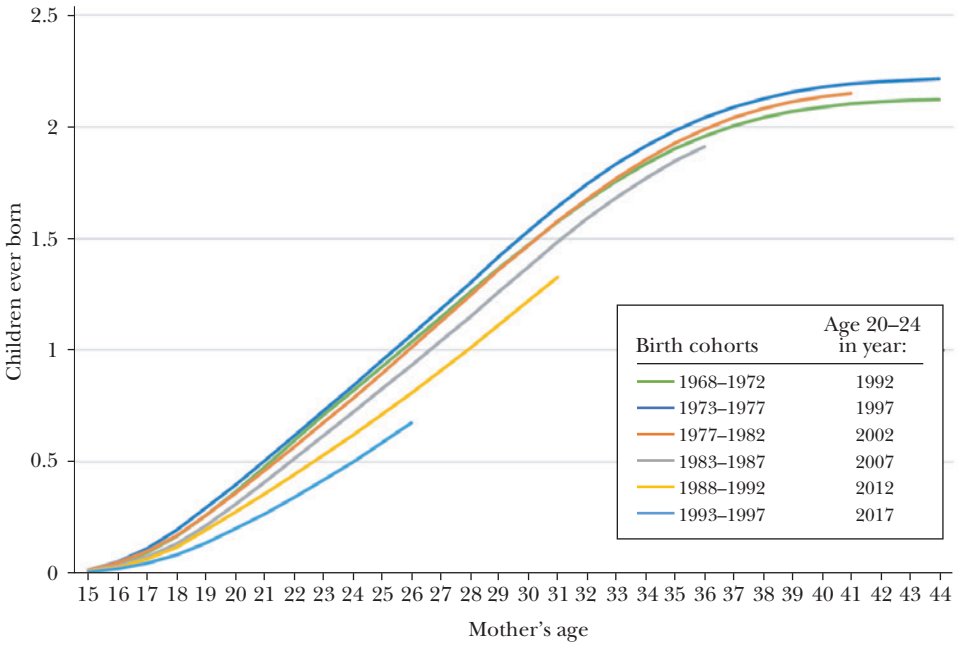

Source: Authors' calculations based on NCHS Vital Statistics birth microdata and CDC SEER population data. See the Data Appendix for specific details.

cohorts comprising women entering their prime childbearing years (ages 20 to 24) in 1992, 1997, 2002, 2007, 2012, and 2017, respectively.

As this figure shows, the three cohorts of women who entered their young adult years in 1992, 1997, and 2002 (born between 1968 and 1982), all had similar childbearing age profiles. Then, the cohort of women who entered young adulthood in 2007 (the 1983–87 birth cohort), had fewer children throughout their 20s and their 30s. The next two cohorts of young adults, the birth cohorts who entered their prime childbearing years in 2012 and 2017 (born between 1988 and 1997), are pulling even further away from earlier cohorts, having fewer children so far. Recall from Figure 2 that the decline in births since 2007 is driven by declines in births to women in their 20s and declines in first and second births. This is consistent with a decline in births mostly to more recent cohorts of women. In a mechanical sense, this divergence across cohorts can explain the sizable decrease in annual birth rates that began in 2007.

This analysis suggests that to understand the factors behind falling annual birth rates, we should be looking for circumstances related to cohorts, as opposed to contemporaneous years. Perhaps the explanation lies more in the way these cohorts were raised or in how they experienced their childhoods, than about a particular

policy or cost factor post-2007. This type of investigation is much harder to undertake with standard data and econometric techniques. The role of factors that affect the decisions of an entire cohort are harder to pin down econometrically than are the role of policy factors that vary across states and years. The remainder of this section is thus speculative.

A Potential Explanation: Shifting Priorities

We propose a general explanation for the decline in births across recent cohorts of US women that focuses on the shifting priorities of cohorts. We introduce this term as a catch-all phrase that encompasses preferences for having children, life aspirations, and the nature of parenting, among other things.

This speculative hypothesis is related to the concept of a "second demographic transition," as proposed in Lesthaeghe and van de Kaa (1986) and Lesthaeghe (2014). The first demographic transition refers to the movement from high to low levels of fertility and death rates historically associated with the industrial revolution. The theory of the second demographic transition highlights instead an overall shift to a greater emphasis on individual autonomy, with a corresponding de-emphasis on marriage and parenthood. The specific manifestations of this shift are taken to include a decoupling of marriage and childbearing, a change in the relationship between education and childbearing, a rise in childlessness, and the establishment of a two-child norm for those having children. Zaidi and Morgan (2017) put this observation into the framework of Becker, arguing that "tastes and preferences have irreversibly changed" (p. 478). These authors point out in their literature review that this explanation has been widely applied to the European context.[13]

Beyond these attitudinal changes, one specific aspect of modern life that may contribute to young adults' views about having children is how the act of "parenting" has evolved over recent decades. Parenting has become more resource- and time-intensive, both in the United States, as well as in many other high-income countries (Bianchi 2011; Kornrich and Furstenberg 2013; Doepke and Zilibotti 2019). Changing norms regarding the intensity of parenting might change people's views toward how many children to have or whether to have them at all. Such changes are particularly relevant in an era where parents, including mothers, work longer hours outside the home, clashing with career aspirations or a desire for more leisure time. This idea incorporates choice in the context of a quality/quantity trade-off, but it also emphasizes external determinants or expectations of what is generally expected or required of parents.

It is unlikely that career aspirations or parenting norms changed exactly in or around 2007. Note, though, that women who grew up in the 1990s were the

[13] With regard to the US context, Bailey, Guldi, and Hershbein (2014) examine demographic changes during the 1960s and 1970s, concluding that it is not clear whether the changes in marriage, family, and childbearing observed during those decades are sufficiently large to constitute a second demographic transition at that time. Lesthaeghe and Neidert (2006) asked whether perhaps the United States was a counterexample of this phenomenon. However, perhaps these dynamics were slower to arrive in the United States. We leave this open question to future research and examination.

daughters of the 1970s generation and women who grew up in the 1970s and 1980s were daughters of the 1950s and 1960s generation. It seems plausible that these more recent cohorts of women were likely to be raised with stronger expectations of having life pursuits outside their roles as wives and mothers. It also seems likely that the cohorts of young adults who grew up primarily in the 1990s or later—and reached prime childbearing years around and post 2007—experienced more intensive parenting from their own parents than those who grew up primarily in the 1970s and 1980s. They would have a different idea about what parenting involves. We speculate that these differences in formed aspirations and childhood experiences could potentially explain why more recent cohorts of young women are having fewer children than previous cohorts.

Related (Limited) Survey Evidence

There is some limited survey evidence that bears on our proposed hypothesis that shifting priorities across cohorts may be a key explanation for declining US birth rates. We readily acknowledge the limitations of such self-reported survey responses. Self-reports about priorities, attitudes, reasons, and so on are often subject to interpretation, sensitive to survey wording or current context, plagued by issues of recall, and other limitations. We also acknowledge the limited amount of survey evidence available, particularly surveys that have been asked repeatedly over time to track trends. Still, they offer some insights, and we describe some of this evidence here.

Some nationally representative surveys ask women about their expectations or desires for childbearing. On this point, the number of children that women report wanting to have has been dropping slightly. Hartnett and Gemmill (2020) report that data from the 2006–2017 National Survey of Family Growth shows that the total number of children women intend to have declined (from 2.26 in 2006–2010 to 2.16 children in 2013–2017) and that the proportion of women intending to remain childless increased slightly. Women also tend to end up having fewer children than they say would be ideal and that gap has been growing(Stone 2021). One interpretation of this discrepancy is that it offers prima facie evidence that constraints or costs are playing a role in depressing birth rates. An alternative interpretation is that women report they want, say, two or three children, but when faced with actual trade-offs associated with having more children, they choose differently.

In a 2018 survey conducted for the *New York Times*, the leading self-reported reasons for why US adults had fewer children than they planned included concerns about the expense of childcare costs, the costs of raising a child, and worries about the economy or their own financial instability (Miller 2018). Other frequently noted reasons include wanting to spend more time with children they already had or wanting more leisure time. The desire to have more leisure time is also reported as the leading reason among adults who said they did not want to have children or were not sure whether they did. We have been unable to find comparable data from an earlier period to see if stated priorities have shifted, but even without that comparison, these responses are potentially illuminating.

We also looked to survey data on young adults' stated attitudes about having children and the importance of various goals and achievements in life. Data from the World Values Survey gives some insight into the expressed priorities of women between the ages of 20 and 44 in the United States in 2005–2009 (Inglehart et al. 2014) and in 2017–2020 (Haerpfer et al. 2020). Between those two survey years, the percentage of women who report that work is very important to them rose from 31.9 percent to 47.4 percent.

Conclusion: Some Thoughts on a Declining US Birth Rate

Why Does a Declining Birth Rate Matter?

A decline in annual birth rates does not necessarily imply a long-term reduction in childbearing. If the recent decline in annual birth rates simply reflects women pushing off having children from their 20s to their 30s, then annual birth rates will eventually rebound and the total number of children the average US woman has over her lifetime will not change. Any long-run implications of the current decline in births would then be modest.

But this pattern of offsetting changes seems unlikely. As we showed in the previous section, the decline in annual birth rates since 2007 is consistent with more recent cohorts of women having fewer births. Those cohorts have not completed their childbearing years yet, but the number of births they would have to have at older ages to catch up to the lifetime childbearing rates of earlier cohorts is so large that it seems unlikely they will do so. Kearney and Levine (2021) project that the total number of children ever born of more recent cohorts of women is likely to fall well below that of previous cohorts, and specifically, below the replacement level of 2.1 births per woman.

A persistent decline in births across age groups will eventually affect population composition and size. If birth rates remain persistently low, and net immigration is not increased to make up for smaller cohorts at younger ages (which seems politically unlikely), then the US population will age and potentially shrink. Some of the consequences that have been discussed include a decline in productivity, instability in financing of old-age programs, and potential for environmental gains.

The potential for interaction between population growth and economic growth has been on the agenda of economists for some time. For example, in the American Economic Association Presidential Address delivered by Alvin Hansen (1939), "Economic Progress and Declining Population Growth," he argues that the rock-bottom birth rates of the Great Depression were one of the reasons leading to less incentive for investment, thus leading to a future of "secular stagnation." Hansen also pointed out that Adam Smith had hypothesized about how population growth expands productivity growth, because a larger population has more opportunities for the division of labor. More recently, there has been an array of arguments about the "demographic dividend," the idea that when a larger share of the workforce is in its early or prime working years, it will tend to stimulate

economic growth. Conversely, an aging population could mean lower per capita GDP if older workers are less likely to work and, conditional on working, be less productive (Maestas, Mullen, and Powell 2016). Growth theorists like Jones (2020) have presented models where lower population growth leads to lower economic growth via a reduced number of new ideas that can become the source for technological progress. These theories all differ in various ways. Our goal here is not to sort them out, but only to suggest that a link from slower population growth to slower economic growth has some plausibility and a pedigree in the research literature.

An aging population also puts pressure on social insurance programs, like Social Security and Medicare (technically, the Old Age, Survivors, Disability and Health Insurance program), because these programs provide benefits to nonworking individuals funded through taxes on workers. The striking decline in birth rates since 2007 means that predictions made at that time about the long-run fiscal sustainability of those programs were overoptimistic (Office of the Chief Actuary, Social Security Administration 2007).

Some contend that a shrinking population would be beneficial for the environment. Ecologists use the "IPAT" equation to describe the impact of human activity on the environment: Impact = Population × Affluence × Technology. Affluence is defined to be consumption per person (which is linked to GDP per capita), while technology represents the amount of resources required to produce a unit of GDP (Chertow 2000). We claim no expertise in this domain, but our review of the relevant evidence suggests that the amount of population decline that would be necessary to reduce the impact of human activity on the environment in a meaningful way is far greater than what will be achieved by the realized reduction in US birth rates. Reducing human impact on the environment through reduced consumption of energy and materials per person, along with more sustainable production processes, is much more likely to have meaningful impacts.

Is There a Role for Pronatalist Policies?

"Pronatalist" policies generally make it easier or more affordable for families to have children. These include steps like subsidized childcare, parental leave policies, and child allowances or tax credits. Many countries are contemplating or implementing pronatalist policies: the United Nations reports that the number of countries with a policy goal of increasing fertility has risen from 19 to 55 between 1986 and 2015 (Sobotka, Matysiak, and Brzozowska 2019).

The evidence about pronatalist policies that have been implemented and evaluated in the United States and in other high-income countries suggests that these types of policies lead to modest increases in birth rates in the short-term, but are unlikely to lead to sustained higher birth rates (Brainerd 2014; Lopoo et al. 2018; Sobotka, Matysiak, and Brzozowska 2019). Stone (2020a) concluded that a pro-natalist policy would cost $200,000 or more per additional baby born; using such policies to close the gap between current fertility in the United States and the replacement level of fertility would cost somewhere between $250 billion and $1 trillion in new spending per year—a daunting sum.

Final Thoughts

The arrival of the Great Recession offers an immediate and obvious reason why a decline in birth rates started in 2007. However, we do not have solid evidence of US-specific policies or economic factors that can explain the depth of that decline and the way it extended through the entire business cycle up through the arrival of the pandemic recession in 2020. We do know that the trend toward lower US birth rates has brought US fertility rates closer to that of other high-income countries. We suspect that this shift reflects broad societal changes that are hard to measure or quantify: possibilities include changing preferences for children, broader career options (and other aspirations) for women, and shifts in the nature of parenting.

In this essay, we have sidestepped any attempt to make an overall judgement on whether the decline in the US birth rate should be viewed as an overall positive or negative development. On one side, if decreasing births are attributable to greater economic opportunities for women, they may be viewed as a positive development. On the other side, if some women would prefer to have children, but do not feel they have the resources to do so, then this suggests viewing the fall in birth rates as a negative development. Because the evidence does not pinpoint strong contributing factors to the decline in US birth rates, it is difficult to reconcile these normative distinctions.

Whatever normative view one takes about declining fertility rates, it is important from an economic policy standpoint to acknowledge that an aging population and shrinking workforce pose challenges for economic growth and the sustainability of social insurance systems. We see no particular reason to believe that a pro-natalist public agenda will have much effect on birth rates (although of course some parts of that agenda may be desirable for other reasons). Thus, the most appropriate way to address declining US birth rates may be to address its two main symptoms directly: that is, a greater emphasis on technological improvements, along with investments in human capital and productivity-enhancing infrastructure, and a greater emphasis on putting the finances of Social Security and Medicare on a secure basis for the long-term. The US economy and political system will need to contend with these issues if the recent, sustained decline in birth rates is not reversed.

■ *The authors are grateful to the editors of the Journal of Economic Perspectives for comments and suggestions.*

References

Adserà, Alícia. 2013. "Fertility, Feminism, and Faith: The Influence of Secularism and Economic Conditions." In *Whither the Child? Causes and Consequences of Low Fertility*, edited by Eric P. Kaufmann and W. Bradford Wilcox. New York: Routledge.

Aizer, Anna, and Sara McLanahan. 2006. "The Impact of Child Support Enforcement on Fertility, Parental Investments, and Child Well-Being." *Journal of Human Resources* 41(1): 28–45.

Bailey, Martha J. 2010. "'Momma's Got the Pill': How Anthony Comstock and Griswold v. Connecticut Shaped US Childbearing." *American Economic Review* 100 (1): 98–129.

Bailey, Martha J., Melanie E. Guldi, and Brad J. Hershbein. 2014. "Is There a Case for a 'Second Demographic Transition'? Three Distinctive Features of the Post-1960 U.S. Fertility Decline." In *Human Capital in History: The American Record*, edited by Leah P. Boustan, Carola Frydman, and Robert A. Margo, 273–312. Chicago, IL: University of Chicago Press.

Becker, Gary S. 1960. "An Economic Analysis of Fertility." In *Demographic and Economic Change in* Developed, edited by George B. Roberts, 209–40. Princeton: Princeton University Press.

Becker, Gary S., and H. Gregg Lewis. 1973. "On the Interaction Between the Quantity and Quality of Children." *Journal of Political Economy* 81 (2): S279–88.

Belluz, Julia. 2020. "The Historically Low Birthrate, Explained in 3 Charts: US Women are Having Fewer and Fewer Babies. In Some Ways, It's a Sign of Progress." *Vox*, January 13. https://www.vox.com/science-and-health/2018/5/22/17376536/fertility-rate-united-states-births-women.

Bianchi, Suzanne M. 2011. "Family Change and Time Allocation in American Families." *Annals of the American Academy of Political and Social Science* 638 (1): 21–44

Black, Dan, Natalia Kolesnikova, Seth Sanders, and Lowell J. Taylor. 2013. "Are Children 'Normal'?" *Review of Economics and Statistics* 95 (1): 21–33.

Brainerd, Elizabeth. 2014. "Can Government Policies Reverse Undesirable Declines in Fertility?" https://wol.iza.org/articles/can-government-policies-reverse-undesirable-declines-in-fertility/long.

Buckles, Kasey, Melanie E. Guldi, and Lucie Schmidt. 2019. "Fertility Trends in the United States, 1980–2017." NBER working paper 25521.

Bullinger, Lindsey Rose. 2017. "The Effect of Minimum Wages on Adolescent Fertility: A Nationwide Analysis." *American Journal of Public Health* 107 (3): 447–52.

Butz, William P., and Michael P. Ward. 1979. "The Emergence of Countercyclical U.S. Fertility." *American Economic Review* 69 (3): 318–28.

Chertow, Marian R. 2000. "The IPAT Equation and Its Variants." *Journal of Industrial Ecology* 4 (4): 13–19.

DeBarros, Anthony, and Janet Adamy. 2019. "US Births Fall to Lowest Level since 1980s: Number of Babies Born Has Declined in 10 of the Last 11 Years as Fertility Rate Slumps." *Wall Street Journal*, May 15. https://www.wsj.com/articles/u-s-births-fall-to-lowest-rates-since-1980s-11557892860.

DeLeire, Thomas, Leonard Lopoo, and Kosali I. Simon. 2011. "Medicaid Expansions and Fertility in the United States." Demography 48 (2): 725–47.

Dettling, Lisa, and Melisa S. Kearney. 2014. "House Prices and Birth Rates: The Impact of the Real Estate Market on the Decision to Have a Baby." *Journal of Public Economics* 110: 82–100.

Dickler, Jessica. 2018. "Student Loan Debt Is a Hurdle for Many Would-Be Mothers." *CNBC*, May 22. https://www.cnbc.com/2018/05/22/student-loan-debt-is-a-hurdle-for-many-would-be-mothers.html.

Doepke, Matthias, and Fabrizio Zilibotti. 2019. *Love, Money, and Parenting: How Economics Explains the Way We Raise Our Kids*. Princeton, New Jersey: Princeton University Press.

Douthat, Ross. 2020. *The Decadent Society: America Before and After the Pandemic*. New York: Simon & Schuster.

Grogger, Jeffrey, Lynn A. Karoly, and Jacob Alex Klerman. 2002. *Consequences of Welfare Reform: A Research Synthesis*. Santa Monica, CA: RAND Corporation.

Haerpfer, C., R. Inglehart, A. Moreno, C. Welzel, K. Kizilova, J. Diez-Medrano, M. Lagos, P. Norris, E. Ponarin & B. Puranen et al. (eds.). 2020. "World Values Survey: Round Seven - Country-Pooled Datafile. Madrid, Spain & Vienna, Austria: JD Systems Institute & WVSA Secretariat." doi.org/10.14281/18241.1 (accessed September 28, 2021).

Hansen, Alvin H. 1939. "Economic Progress and Declining Population Growth." *American Economic Review* 29 (1): 1–15.

Hartnett, Caroline Sten, and Alison Gemmill. 2020. "Recent Trends in U.S. Childbearing Intentions." *Demography* 57 (6): 2035–45.

Hoffower, Hillary. 2021. "The Declining American Birth Rate Could Actually Be Good for the Economy." *Insider*, May 23. https://www.businessinsider.com/birth-rate-decline-pandemic-millennials-change-reshaped-economy-infrastructure-2021-5.

Iati, Marisa. 2019. "Birthrates in the U.S. Are Falling. Abortions Have also Hit an All-Time Low." *Washington Post,* November 27. https://www.washingtonpost.com/health/2019/11/27/birthrates-us-are-falling-abortions-have-also-hit-an-all-time-low/.

Inglehart, R., C. Haerpfer, A. Moreno, C. Welzel, K. Kizilova, J. Diez-Medrano, M. Lagos, P. Norris, E. Ponarin and B. Puranen et al., eds. 2014. *World Values Survey: Round Five - Country-Pooled Datafile* Version: www.worldvaluessurvey.org/WVSDocumentationWV5.jsp. Madrid: JD Systems Institute (accessed September 28, 2021).

Jones, Charles I. 2020. "The End of Economic Growth? Unintended Consequences of a Declining Population." NBER Working Paper 26651.

Kearney, Melissa S., and Phillip B. Levine. 2009. "Subsidized Contraception, Fertility, and Sexual Behavior." *Review of Economics and Statistics* 91 (1):137–51.

Kearney, Melissa S., and Phillip B. Levine. 2012. "Why Is the Teen Birth Rate in the United States So High and Why Does It Matter?" *Journal of Economic Perspectives* 26 (2): 141–66.

Kearney, Melissa S., and Phillip B. Levine. 2015a. "Investigating Recent Trends in the U.S. Teen Birth Rates." *Journal of Health Economics* 41: 15–29.

Kearney, Melissa S., and Phillip B. Levine. 2015b. "Media Influences on Social Outcomes: The Impact of MTV's 16 and Pregnant on Teen Childbearing." *American Economic Review* 105 (12): 3597–3632.

Kearney, Melissa S., and Riley Wilson. 2018. "Male Earnings, Marriageable Men, and Nonmarital Fertility: Evidence from The Fracking Boom." *Review of Economics and Statistics* 100 (4): 678–90.

Kearney, Melissa S., and Phillip B. Levine. 2020. "Half a Million Fewer Children? The Coming COVID Baby Bust." https://www.brookings.edu/research/half-a-million-fewer-children-the-coming-covid-baby-bust/.

Kearney, Melissa S., and Phillip B. Levine. 2021. "Will Births in the US Rebound? Probably Not." https://www.brookings.edu/blog/up-front/2021/05/24/will-births-in-the-us-rebound-probably-not/.

Kelly, Andrea, Jason M. Lindo, and Analisa Packham. 2020. "The Power of the IUD: Effects of Expanding Access to Contraception through Title X Clinics." *Journal of Public Economics* 192.

Kornrich, Sabino, and Frank Furstenberg. 2013."Investing in Children: Changes in Parental Spending on Children, 1972–2007." *Demography* 50 (1): 1–23.

Lesthaeghe, Ron. 2014. "The Second Demographic Transition: A Concise Overview of Its Development." *Proceedings of the National Academy of Sciences* 111 (51): 18112–15.

Lesthaeghe, Ron J., and Lisa Neidert. 2006. "The Second Demographic Transition in the United States: Exception or Textbook Example?" *Population and Development Review* 32 (4): 669–98.

Lesthaeghe, Ron J., and van de Kaa, Dirk. 1986. "Twee Demografische Transities [Two Demographic Transitions]." Lesthaeghe and van de Kaa In *Groei of Krimp [Growth or Shrinkage]?* Van Loghum Slaterus, Deventer: Mens en Maatschappij.

Levine, Phillip B. 2004. *Sex and Consequences: Abortion, Public Policy, and the Economics of Fertility.* Princeton, NJ: Princeton University Press.

Lindo, Jason M. 2010. "Are Children Really Inferior Goods? Evidence from Displacement-Driven Income Shocks." *Journal of Human Resources* 45 (2): 301–27.

Lindo, Jason M., and Analisa Packham. 2017. "How Much Can Expanding Access to Long-Acting Reversible Contraceptives Reduce Teen Birth Rates?" *American Economic Journal: Economic Policy* 9 (3): 348–76.

Lopoo, Leonard M., and Kerri M. Raissian. 2012. "Natalist Policies in the United States." *Journal of Policy Analysis and Management* 31 (4): 905–46.

Lopoo, Leonard M., Kerri M. Raissian, S.L. Averett, L.M. Argys, and S.D. Hoffman. 2018. "Fertility Policy in Developed Countries." *Oxford Handbook of Women and the Economy* 173.

Lovenheim, Michael, and Kevin Mumford. 2013. "Do Family Wealth Shocks Affect Fertility Choices? Evidence from the Housing Market." *Review of Economics and Statistics* 95 (2): 464–75.

Maestas, Nicole, Kathleen J. Mullen, and David Powell. 2016. "The Effect of Population Aging on Economic Growth, the Labor Force and Productivity." NBER Working Paper 22452.

Miller, Claire Cain. 2018. "Americans Are Having Fewer Babies. They Told Us Why." *New York Times*, July 5. https://www.nytimes.com/2018/07/05/upshot/americans-are-having-fewer-babies-they-told-us-why.html.

Moffitt, Robert A. 1998. "The Effect of Welfare on Marriage and Fertility." In *Welfare, the Family and Reproductive Behavior*, edited by Robert A. Moffitt, 50–97. Washington DC: National Academy Press.

Office of the Chief Actuary, Social Security Administration. 2007. *The 2007 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds*. Washington, DC: Government Printing Office.

Parrado, Emilio A., and S. Philip Morgan. 2008. "Intergenerational Fertility among Hispanic Women: New Evidence of Immigrant Assimilation." *Demography* 45 (3): 651–71.

Paton, David, Stephen Bullivant, and Juan Soto. 2020. "The Impact of Sex Education Mandates on Teenage Pregnancy: International Evidence." *Health Economics* 29 (7): 790–807.

Schaller, Jessamyn. 2016."Booms, Busts, and Fertility: Testing the Becker Model Using Gender-Specific Labor Demand." *Journal of Human Resources* 51 (1): 1–29.

Snodgrass, Erin. 2021. "Alexandria Ocasio-Cortez Decries the Student Loan Debt, Low Incomes, and Climate Change Facing 'Entire Generations' as Birth Rates Continue to Fall." *Business Insider*, May 3. https://www.businessinsider.com/aoc-decries-hardships-facing-generations-as-birth-rates-fall-2021-3.

Sobotka, Tomáš, Anna Matysiak, and Zuzanna Brzozowska. 2019. *Policy Responses to Low Fertility: How Effective Are They?* New York: United Nations.

Stone, Lyman. 2018. *Declining Fertility in America*. Washington, DC: American Enterprise Institute.

Stone, Lyman. 2020a. "Pro-Natal Policies Work, but They Come with a Hefty Price Tag." *Institute for Family Studies*, March 5. https://ifstudies.org/blog/pro-natal-policies-work-but-they-come-with-a-hefty-price-tag.

Stone, Lyman. 2020b. "The Rise of Childless America." *Institute for Family Studies*, June 4. https://ifstudies.org/blog/the-rise-of-childless-america.

Stone, Lyman. 2021. "Escaping the Parent Trap: Addressing America's Fertility Crisis Happens to Be What Parents Want." *American Compass*, February 11. https://americancompass.org/essays/escaping-the-parent-trap/.

Tavernise, Sabrina. 2019. "Why Birthrates among Hispanic Americans Have Plummeted." *New York Times*. https://www.nytimes.com/2019/03/07/us/us-birthrate-hispanics-latinos.html.

Tavernise, Sabrina. 2021. "Why American Women Everywhere Are Delaying Motherhood." *New York Times*. https://www.nytimes.com/2021/06/16/us/declining-birthrate-motherhood.html.

US Census Bureau. 2021. "Table MS-2. Estimated Median Age at First Marriage, by Sex: 1890 to the Present." https://www2.census.gov/programs-surveys/demo/tables/families/time-series/marital/ms2.xls (accessed August 31, 2021).

World Bank. 2021. "Fertility Rate, Total (Births per Woman) – Mexico." https://data.worldbank.org/indicator/SP.DYN.TFRT.IN?locations=MX, (accessed June 9, 2021).

Zaidi, Batool and S. Philip Morgan. 2017. "The Second Demographic Transition Theory: A Review and Appraisal." *Annual Review of Sociology* 43: 473–92.

Ziliak, James. P. 2016. "Temporary Assistance for Needy Families." In *Economics of Means-Tested Transfer Programs in the United States,* Vol. 1, edited by Robert Moffitt, 303–93. Chicago: University of Chicago Press.

Journal of Economic Perspectives—Volume 36, Number 1—Winter 2022—Pages 177–190

Isaiah Andrews, 2021 John Bates Clark Medalist

Anna Mikusheva and Jesse M. Shapiro

Isaiah Andrews is an exceptionally warm and caring person who adds to the knowledge and joy of his students, teachers, and colleagues. He is also a brilliant econometrician, and it is through his research in econometrics that he has come to be recognized with the 2021 John Bates Clark Medal. The Clark Medal is the most recent in a list of impressive accomplishments, which include a MacArthur Fellowship, a Fellowship in the Econometric Society, a Sloan Research Fellowship, a Junior Fellowship at the Harvard Society of Fellows, and others.

These and other dazzling intellectual achievements accompany a record of exceptional service to the profession. Isaiah currently serves as a co-editor of the *American Economic Review* and previously served as an associate editor at four different journals. He serves on the American Economic Association Committee on the Status of Minority Groups in the Economics Profession. He has co-organized a recent meeting of the NBER Working Group on Race and Stratification and has served on program committees for both the American Economic Association and the Econometric Society.

In addition to serving the wider profession, Isaiah has made important service contributions to his home department. After moving to Harvard in 2018, Isaiah helped to launch an Econometrics Clinic to which students in all fields can come for research advice. The clinic fits with Isaiah's strongly held view that econometric theory derives value from its ability to influence and improve empirical research on

■ *Anna Mikusheva is Associate Professor of Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts. Jesse M. Shapiro is the George Gund Professor of Economics and Business Administration, Harvard University, Cambridge, Massachusetts.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.36.1.177.

Isaiah Andrews

important real-world problems. The clinic also fits with Isaiah's otherworldly talents as a teacher and communicator, able to distill a complex idea in a way that conveys its meaning to nonspecialists while still preserving many of its subtleties. It is difficult to convey how talented and special Isaiah is in this dimension; we hope that all readers have the chance someday to experience it for themselves.

Isaiah's paternal grandmother, Viola Andrews, graduated from Cheyney State Teachers College, considered the oldest of the nation's historically black colleges and universities (Cheyney University of Pennsylvania 2021). Isaiah's maternal grandparents both completed dissertations involving empirical research (Smith 1966; Smith 1971). Isaiah's mother and father received their PhDs in economics from Yale in 1984 and 1986, respectively (Smith 1984; Andrews 1986). As the child of two Yale-educated economists, Isaiah planned to avoid both Yale and economics, but eventually chose to pursue both. Isaiah received his BA in math and economics from Yale in 2009 and began his PhD studies in economics at MIT in that same year.

It was not inevitable that Isaiah would become a theoretical econometrician. For example, he might have become a theoretical macroeconomist. But at a fateful dinner at a graduate recruiting event in his first year at MIT, Isaiah sat next to Mikusheva. The seat led to a summer research assistantship. The research assistantship led to a coauthorship. And the rest is the intellectual story we are here to tell.

In the remainder of this article, we discuss Isaiah's contributions to econometric theory and empirical practice. A large portion of Isaiah's work concerns instrumental variables and related methods. Accordingly, we begin with a short review of these methods, motivated by the classic case of estimating the elasticity of demand for an agricultural commodity.

To wit, say that we want to estimate an elasticity of demand for wheat using historical data on wheat prices and quantities sold. Tempting though it may be, it is usually a bad idea to estimate the elasticity of demand via ordinary least squares regression of the log of quantity sold on the log of price (Working 1927). Because any observed price-quantity combination is an equilibrium outcome determined both by supply and demand, both price and quantity are influenced by unobserved determinants of demand, and so the ordinary least squares estimate of the demand elasticity is not generally reliable.

A possible solution is to find an *instrumental variable*, say, weather, that induces variation in prices that is unrelated to the unobserved determinants of demand, and to trace the influence of the instrumental variable, through prices, on quantity sold (Wright 1928). Intuitively, if changes in the weather shift the supply curve and are not related to unobserved determinants of demand, then changes in prices and quantity induced by changes in weather will trace out the demand curve.

In this example, the instrumental variables estimator of the elasticity of demand can be written as a ratio of the ordinary least squares estimator of the effect of the weather on the log of quantity sold to the ordinary least squares estimator of the effect of the weather on the log of price. There are two important conditions needed for the validity of this type of estimator. The first is called the *relevance condition*. It states that the instrumental variable is correlated with the variable we wish to instrument for. In our example, this means that the weather is correlated with the log of price. The second is called the *exclusion restriction*. It states that the instrumental variable is uncorrelated with unobserved determinants of the outcome. In our example, this means that the weather is uncorrelated with unobserved determinants of the log of quantity demanded.

Much of Isaiah's research concerns situations in which the relevance condition or the exclusion restriction are not guaranteed to hold convincingly, or even to hold at all, in the economic setting at hand. We discuss the two conditions in turn. In Table 1, we list a selection of Isaiah's published papers, and we refer to those papers by number in the discussion that follows.

The Relevance Condition

If the relevance condition were to fail (say, if the weather were not correlated with the log of price), then the instrumental variable would not help us recover the true value of the elasticity of demand because variation in the weather would tell us nothing about the influence of the price on the quantity demanded. Indeed, if in a particular sample the weather were totally uncorrelated with the log of the price, then the denominator of the instrumental variables estimator would be zero, and the estimator itself would be undefined!

Consider, though, what would happen if the effect of weather on the log of the price were very small but not exactly zero. In this case, the elasticity could, in principle, be estimated, but the instrumental variables estimator would be very sensitive

Table 1
Selected Publications of Isaiah Andrews

1	"Weak Identification in Maximum Likelihood: A Question of Information" (with Anna Mikusheva). 2014 *American Economic Reviews: Papers and Proceedings* 104 (5): 195–99.
2	"Maximum Likelihood Inference in Weakly Identified DSGE Models" (with Anna Mikusheva). 2015. *Quantitative Economics* 6 (1):123–52.
3	"A Geometric Approach to Nonlinear Econometric Models" (with Anna Mikusheva). 2016. *Econometrica* 84 (3): 1249–64.
4	"Conditional Inference with a Functional Nuisance Parameter" (with Anna Mikusheva). 2016. *Econometrica* 84 (4): 1571–1612.
5	"The Allocation of Future Business: Dynamic Relational Contracts with Multiple Agents" (with Daniel Barron). 2016. *American Economic Review* 106 (9): 2742–59.
6	"Conditional Linear Combination Tests for Weakly Identified Models." 2016. *Econometrica* 84 (6): 2155–82.
7	"Unbiased Instrumental Variables Estimation Under Known First-Stage Sign" (with Timothy B. Armstrong). 2017. *Quantitative Economics* 8 (2): 479–503.
8	"Measuring the Sensitivity of Parameter Estimates to Estimation Moments" (with Matthew Gentzkow and Jesse M. Shapiro). 2017. *Quarterly Journal of Economics* 132 (4): 1553–92.
9	"Valid Two-Step Identification-Robust Confidence Sets for GMM." 2018. *Review of Economics and Statistics* 100 (2): 337–48.
10	"On the Structure of IV Estimands." 2019. *Journal of Econometrics* 211 (1): 294–307.
11	"Identification of and Correction for Publication Bias" (with Maximilian Kasy). 2019. *American Economic Review* 109 (8): 2766–94.
12	"Weak Instruments in IV Regression: Theory and Practice" (with James Stock and Liyang Sun). 2019. *Annual Review of Economics* 11: 727–53.
13	"A Simple Approximation for Evaluating External Validity Bias" (with Emily Oster). 2019. *Economics Letters* 178: 58–62.
14	"On the Informativeness of Descriptive Statistics for Structural Estimates" (with Matthew Gentzkow and Jesse M. Shapiro, Matthew Gentzkow's Fisher-Schultz Lecture). 2020. *Econometrica* 88 (6): 2231–58.
15	"Transparency in Structural Research" (with Matthew Gentzkow and Jesse M. Shapiro, invited discussion paper). 2020. *Journal of Business and Economic Statistics* 38(4): 711–22.
16	"Inference After Estimation of Breaks" (with Toru Kitagawa and Adam McCloskey). Forthcoming. *Journal of Econometrics.*
17	"A Model of Scientific Communication" (with Jesse M. Shapiro). Forthcoming. *Econometrica.*
18	"Inference for Linear Conditional Moment Inequalities" (with Jonathan Roth and Ariel Pakes).
19	"Inference on Winners" (with Toru Kitagawa and Adam McCloskey).
20	"Optimal Decision Rules for Weak GMM" (with Anna Mikusheva). Forthcoming. *Econometrica*

to small variations in its denominator, because as the denominator approaches zero, the estimator's value would "explode" toward negative or positive infinity.

Figure 1 illustrates this situation in a case where more favorable weather is correlated with lower prices and greater quantity. As the effect of weather on prices

Figure 1
Sensitivity of the Instrumental Variables Estimator to the Denominator

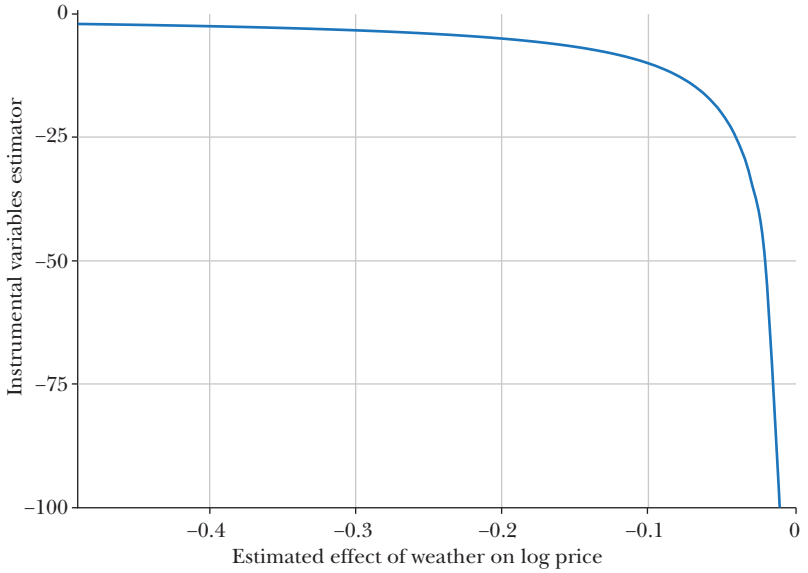

Source: Illustrative calculations by authors.
Note: The figure illustrates the instrumental variables estimator in a hypothetical case (with hypothetical numbers) where the numerator (the ordinary least squares coefficient from a regression of the log of quantity on the weather) is always 1 and we vary the value of the denominator (the ordinary least squares coefficient from a regression of the log of price on the weather). The figure shows that, as the denominator approaches zero, the estimator becomes very sensitive to small variations in the denominator.

approaches zero, the estimated elasticity more and more quickly approaches negative infinity.

The sensitivity of the instrumental variables estimator to its denominator is important because, like any statistical average, the denominator of the instrumental variables estimator is subject to random variation due to the particular sample of data (time periods in our example) included in the analysis. As Figure 1 suggests, when the denominator is close to zero, small sampling variation in the denominator can induce large, even explosive, sampling variation in the estimator.

The resulting statistical variability in the instrumental variables estimator is not captured well by the standard errors and confidence intervals that economists have often used to describe it. The reason is that such tools are often based on assumptions under which the estimator follows a bell curve or normal distribution, at least approximately. The huge sensitivity to the denominator that occurs when the denominator is small can make such approximations very poor.

The situation we are describing is a special case of a broader phenomenon called *weak instruments*, which arises when the instrument has only a small impact

on the variable it is meant to affect. Interest in weak instruments among economists picked up in the 1990s when it was observed in some important research settings (Bound, Jaeger, and Baker 1995; Staiger and Stock 1997). By the time Isaiah began working on this topic in 2010, it was a mature area with a substantial body of theoretical research and practical tools.

The phenomenon of weak instruments is a special case of the more general phenomenon of *weak identification*. Intuitively, one way to think about such a situation is that the data are not sufficiently informative to clearly distinguish a single value of the parameter that is most consistent with the economic model. In our demand equation example, this manifests through the fact that when weather has very little relationship to prices, a large range of elasticities of demand will imply very little correlation between weather and unobserved determinants of demand and will therefore all be nearly consistent with the exclusion restriction.

Although Isaiah has made many important contributions to the study of weak instruments, his earliest work on the topic concerned weak identification in nonlinear models estimated via Generalized Method of Moments (GMM). GMM captures the implications of an economic theory through a set of moment conditions, which specify that some function of the data and parameters should, on average, be zero if the theory is correct. Estimation proceeds by finding the value of the parameter that makes the average value of the moment condition as close to zero as possible in the given dataset. In practice, this is often done by minimizing an objective function defined as a squared deviation of the moment condition from zero.

In our demand estimation example, a key moment condition is that the correlation between the weather and the unobserved determinants of the log of quantity is zero. One way in which instrumental variable regression is special is that this moment condition can be written as linear in the structural parameter (the elasticity of demand). If the weather is totally uncorrelated with the log of price, then the objective function is completely flat and so minimizing it does not yield a unique value of the elasticity of demand. If the weather is only weakly correlated with the log of price, then the objective function is nearly flat close to its minimum, making the solution to the minimization problem, and hence the instrumental variables estimator, very sensitive to sampling variation. Thus, in the case of the linear model, weak identification arises because there is an object (the correlation of the weather and the log of price) that must be nonzero in order for the data to pin down a unique value of the parameter of interest that minimizes the objective function. As this object approaches zero, identification becomes weak, and the data become uninformative about the parameter of interest.

But what about in nonlinear settings? The earliest work on weak identification in GMM settings attempted to generalize directly, and very literally, from the linear model. This work imagined that there is some object (analogous to the correlation of weather and the log of price) that determines whether it is possible to pin down a unique value of the structural parameters, and that the degree of "weakness" can be assessed by the proximity of this object to the value that leads to non-uniqueness (zero in the case of the correlation).

Figure 2
Illustration of Weakly Identified GMM

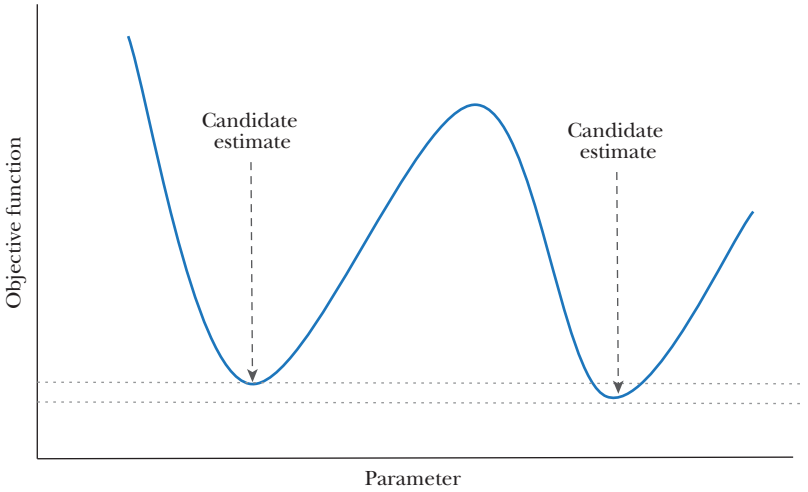

Source: Illustrative drawing by authors.
Note: The figure illustrates a hypothetical GMM objective function with two local minima. Different samples could cause either the right or the left minimum to be the lower one—and thus the parameter estimate, chosen as the minimizer of the objective, may differ substantially across sample realizations.

A central insight and contribution of Andrews and Mikusheva's work was the recognition that this analogy is incomplete. In nonlinear GMM settings, issues of weak identification can arise even if there is always a unique value of the parameter of interest that minimizes the objective function. Consider Figure 2, in which the horizontal axis shows a range of possible values for the parameter being estimated, and the vertical axis shows the value of the GMM objective function at each possible value of the parameter. The objective function exhibits a unique global minimum but two local minima. Because these different local minima achieve similar values of the objective function, small perturbations to the data can lead the global minimum to "jump" between two regions, and as a result, estimates may differ wildly across different datasets.

In addition to showing that the analogy between weak instruments and weak identification was incomplete, Andrews and Mikusheva also found that the analogy was unnecessary in the sense that it was possible to develop tools for statistical inference under weak identification without assuming a structure analogous to that under weak instruments.

In [1] and [2], Andrews and Mikusheva explore the role of weak identification in the estimation of dynamic stochastic general equilibrium models, which are one of the workhorse tools of modern macroeconomics. In these papers, Andrews and Mikusheva propose a new method of diagnosing identification weakness (based on Fisher's information matrix), and [2] proposes a method of hypothesis testing

that is valid under weak identification. It is noteworthy that [2] contains an extensive discussion of practical details of implementation designed to help applied researchers use the ideas in the paper. A concern with applicability is a hallmark of much of Isaiah's research.

In [3] and [4], Andrews and Mikusheva turn their attention to broad classes of nonlinear Generalized Method of Moments (GMM) models that encompass many economic applications, such as dynamic stochastic general equilibrium models, Euler equation models, and quantile instrumental variables models. In [3], Andrews and Mikusheva develop an elegant geometric representation of the behavior of some key statistics in these models. Using this representation, the paper develops new methods for testing statistical hypotheses that are asymptotically valid even under weak identification.

In turn, [4] delves into the underpinnings of weak identification in GMM. In the case of the instrumental variables estimator, the key source of uncertainty is in the denominator, which is a single number. In the more general case, the key source of uncertainty is the shape of the entire objective function, as suggested by Figure 2. Andrews and Mikusheva show how to characterize this uncertainty as a stochastic process, which is a mathematical tool for describing a random function. Using this characterization, [4] develops a form of hypothesis test that is valid under weak identification and can be seen as generalizing earlier tools (Moreira 2003) to a broad class of nonlinear models. A challenge raised by [4] is how to make precise inferences on parameters of interest in the presence of an unknown function that determines the strength of identification. In [20], Andrews and Mikusheva suggest an approach that is based on Bayesian methods and allows a researcher to perform reliable estimation and inference in settings where information about parameters of interest is limited.

Isaiah has also made extensive contributions to the literature on weak instruments, focusing especially on important open problems relevant to empirical researchers. One such question is how to conduct hypothesis tests in situations where there are more instruments than there are unknown structural parameters. No one best hypothesis test applies in this situation, with each type of test tending to underperform in some settings. One way to handle this situation is to switch between different tests depending on features of the setting, to try to leverage the strengths of each test and avoid their weaknesses. Isaiah pursued such an approach in [6], his first sole-authored publication. Isaiah shows how to use information from one test statistic to determine how much weight to give each of two other test statistics. This leads to an approach to hypothesis testing that avoids being too conservative when identification is strong and avoids being too cavalier when identification is weak. It also clarifies some important connections among different ideas in prior work. In [9], Isaiah develops an approach to confidence interval construction suggested by the approach to hypothesis testing in [6] and includes a helpful user's guide designed to aid empirical researchers implementing the approach. Shapiro can report being a satisfied customer (Hastings, Kessler, and Shapiro 2021).

Although most of Isaiah's work on weak instruments and weak identification concerns issues of inference (such as hypothesis testing and confidence interval construction), along the way he also found time to develop an original approach to estimation in collaboration with Tim Armstrong. Recall Figure 1 and consider what would happen if we were not sure that favorable weather conditions lead to lower agricultural prices. Then the sign of the denominator of the instrumental variables estimator would be unknown. Holding fixed the numerator, a switch in the sign of the denominator would lead to even more explosive behavior of the estimator than what Figure 1 suggests, because it would lead the estimator to go from approaching negative infinity to approaching positive infinity! Fortunately, in many situations such as our wheat example, it is possible to make a reasonable economic assumption about the direction of the effect of the instrument (weather) on the variable it instruments (log price). In [7], Andrews and Armstrong show that in such situations, it is possible to construct an instrumental variables estimator that is asymptotically unbiased even when the instrument is weak.

In addition to advancing the frontiers of knowledge, Isaiah has worked to synthesize current knowledge and make it accessible to a wider audience of economists. In 2018, Isaiah and James Stock gave a series of Methods Lectures at the NBER Summer Institute devoted to weak instruments (2018). In [12], Andrews, Stock, and Sophie Sun review theoretical and practical considerations around weak instruments in linear regression settings. They find that weak instruments arise frequently in practice, showing the importance of continuing to develop tools for applied researchers. For scholars interested in understanding current best practices around weak instruments, the Methods Lectures and subsequent review article provide an important resource.

The Exclusion Restriction

When seeking to estimate the elasticity of demand, the weather is an appealing instrument for agricultural prices in part because it is unlikely to directly affect the quantity demanded—people probably do not get hungrier just because growing conditions are more or less favorable. Imagine, though, that a researcher proposes to use the prevailing agricultural wage as an instrument for the price of wheat. Like the weather, the agricultural wage influences supply conditions and would therefore seem likely to satisfy the relevance condition: all else equal, wheat will be more expensive when agricultural labor is scarce. But agricultural wages are also likely to be higher when the economy is doing well, and a strong economy may itself lead to greater demand for food. Therefore, a change in the agricultural wage may relate to the quantity sold both through an effect on the supply of wheat and through an effect on demand via consumers' disposable income. In this case, the exclusion restriction would be violated, and the instrumental variables estimator would be misleading.

Notice that while the relevance condition concerns the relationship between the instrument and an observed variable (in our example, prices), the exclusion

restriction concerns the relationship between the instrument and an unobserved variable (in our example, unobserved determinants of demand). Correspondingly, while the relevance condition is usually testable, the exclusion restriction is often not testable and therefore must be justified on economic or other a priori grounds. The same applies to many other kinds of moment conditions employed in Generalized Method of Moments (GMM).

The fact that the exclusion restriction can be untestable leaves open the possibility that different researchers will disagree about the plausibility of the exclusion restriction, even in a given setting with a given instrument. For example, one researcher might find agricultural wages a very reasonable instrument for the price of wheat, believing that demand for wheat is mostly insensitive to variation in income. Another researcher might find agricultural wages a very unreasonable instrument for the price of wheat, believing that the changes in the economy signaled by a change in the agricultural wage may bring myriad effects on household demand for different goods. Although it may be possible to bring some data to bear on these questions, in many situations data alone cannot resolve such disagreements. In our experience, disagreements about exclusion restrictions are the subject of some of the liveliest exchanges in academic seminars and conferences.

Although, in general, we cannot tell whether the exclusion restriction is violated, we can often say something about how much a given violation of the exclusion restriction would be expected to distort the instrumental variables estimator. In our example, if the instrument is correlated with unobserved determinants of demand, then the instrumental variables estimator is biased, and the extent of the bias is given by the ratio of two coefficients: one the coefficient from a (hypothetical) regression of the unobserved determinants of demand on the instrument and the other the coefficient from a (feasible) regression of the log of price on the instrument.

Conley, Hansen, and Rossi (2012) characterize this bias in a more general set of linear instrumental variables models and advocate using economic intuitions about the correlation between the instrument and the unobservable to adjust inferences about the parameter of interest (in our example, the price elasticity). Conley, Hansen, and Rossi's approach applies to linear models, but, as we noted earlier, many important settings in economics call for estimation of nonlinear models. Within a couple of years, Gentzkow and Shapiro (2014) had muddled through the outlines of a generalization of Conley, Hansen, and Rossi's (2012) result to nonlinear models. Gentzkow presented this work-in-progress at an MIT seminar at which Andrews was present. These and subsequent interactions made clear that Andrews had a rather superior understanding of the issues involved—to adopt a degree of understatement suitable for these pages.[1] Andrews joined the project.

In [8], Andrews, Gentzkow, and Shapiro show how to quantify the effect of violations of an exclusion restriction, or other important economic assumptions, in a class of nonlinear economic models. The paper advocates that researchers report

[1] In what appears to have been Gentzkow's first email to Shapiro on the subject of Isaiah Andrews, Gentzkow wrote, "Did I mention that [Isaiah] is awesome?" (March 21, 2014).

a statistic called *sensitivity* which, when multiplied by a measure of the degree of violation of an exclusion restriction or other moment condition, yields a prediction of the resulting systematic error or bias in the estimator. This recommendation is shown to be practical, in the sense that in many estimation frameworks, such as GMM, sensitivity can be readily calculated, often based on objects a researcher will have already computed for other purposes.

Say, for example, that a researcher uses both the weather and agricultural wages as instruments for the price of wheat in order to estimate the elasticity of demand. By multiplying the sensitivity of the estimator by, say, a guess about the direct effect of agricultural wages on demand, a reader can arrive at an estimate of the likely bias in the estimator. In [8], Andrews, Gentzkow, and Shapiro show how to use sensitivity to quantify the effect of violations of exclusion restrictions in the context of estimating automobile demand.

While finishing [8], Andrews and Shapiro were also working on what became [17]. In that paper, Andrews and Shapiro consider a situation in which a researcher reports statistics to an audience of individuals who have different prior opinions, for example on the likely value of an unknown parameter. Andrews and Shapiro introduce a notion of the quality of a given statistical report, called *communication risk*, which captures the value to the audience of the information in the researcher's report. Andrews and Shapiro show that viewing scientific research as a communication problem can help explain many common research practices that might otherwise seem puzzling, such as reporting an estimate that violates a known sign constraint on the underlying parameter or refusing to report any estimate at all.

Ongoing work on [17] influenced the argument in [8] that reporting sensitivity is helpful to readers of a research article, because it allows them to interpret the estimator in terms of their own beliefs about the plausibility of the exclusion restriction or other important assumptions. The discussion in [8] presents this feature as an example of increasing transparency, in the sense of making the research useful even to readers who do not share the researcher's assumptions. In [15], Andrews, Gentzkow, and Shapiro formalize this connection by introducing a mathematical notion of the transparency of a statistical report based on the notion of communication risk in [17] and show a formal sense in which reporting sensitivity, as advocated in [8], can improve transparency.

A violation of the exclusion restriction distorts the instrumental variables estimator in an especially intuitive way, but the behavior of some estimators may be less intuitive. In [14], Andrews, Gentzkow, and Shapiro show that in such situations, it may be possible to offer guidance on the quantitative importance of different assumptions by connecting the estimator to other more intuitive statistics. Specifically, [14] envisions a situation in which a researcher estimates a parameter of an economic model based on a moment condition that is potentially violated. For a given degree of potential violation of the moment condition, it is possible to characterize the greatest amount of potential bias in the resulting estimator. Now suppose the researcher presents the estimates of some descriptive statistics that are connected to the parameters of interest through the economic model. If we assume that the model correctly

describes the economic connection between the parameter of interest and the descriptive statistics, this reduces the potential for bias in the estimator. The extent to which the potential bias falls in this case is a measure of what Andrews and coauthors call the *informativeness* of the descriptive statistics for the parameter estimate. In [14], Andrews, Gentzkow, and Shapiro advocate that researchers report the informativeness of descriptive statistics alongside estimates of structural parameters. The framework in [14] provides a possible rationale for the common practice of reporting descriptive or summary statistics alongside formal estimates of economic parameters, a practice that is also studied (and formalized) in [15].

Like sensitivity, informativeness is often straightforward to calculate, making the suggestion to report informativeness a practical one. In [14], Andrews, Gentzkow, and Shapiro illustrate with applications. One is to Gentzkow (2007), which estimates a model of newspaper demand and approaches endogeneity using both exclusion restrictions and panel variation. Andrews, Gentzkow, and Shapiro report that in this application, a statistic related to exclusion restrictions is much less informative for a key parameter than is another statistic related to panel variation.

Sometimes researchers have multiple candidate instruments available, and when this happens, there are many different ways to combine the information from each. In [10], Andrews shows that different ways of achieving this combination can have very different properties when the exclusion restriction or other modeling assumptions are violated. This finding may have important implications for applied economists' choice of estimation methods.

A consistent theme in Isaiah's work on transparency is to try to make research more useful to its intended audience while respecting possible differences in opinions and objectives among the members of that audience. An avoidance of dogma and a positive attitude towards differences of opinion are also central to Isaiah's warm and engaging personality and to his professional identity as a researcher, teacher, and communicator. We therefore see this research agenda as an especially good match between the research subject and the researcher.

Publication Bias and Selection among Estimates

Publication bias refers to a situation in which not all estimates that are calculated are reported in the scientific literature (Kasy 2021). For example, researchers or journals may choose not to publish articles whose estimates are not statistically significantly different from zero, or whose estimates fail to support a particular conclusion. If selective publication is not accounted for, a review of the published literature can lead to misleading scientific conclusions. In [11], Andrews and Maximilian Kasy propose a way to correct for publication bias. Suppose we have a model of the publication process that tells us how the probability of publication depends on some aspect of the estimate, such as its degree of statistical significance. In this case, tools of probability theory can tell us how to account for selection and make inferences about the true range of estimates, including those that were not reported.

The challenge, then, is to uncover a valid model of the publication process. In [11], Andrews and Kasy develop an original approach to this problem that importantly avoids placing strong functional form restrictions on the model of the publication process. The approach is developed in two settings. The first setting is a replication study. Here, a researcher or collection of researchers attempt to replicate exactly a set of published studies on a new sample. If there was no publication bias, then the distribution of estimates (and other statistics such as p-values) would be identical between the published studies and the replications. Departures from this type of symmetry can allow us to pin down a model of the publication process from the replication data.

The second setting studied in [11] is a meta-study. Here, a researcher collects existing estimates from the literature, all estimated from different samples but aiming to uncover the same parameter. Absent publication bias, more variable estimates (say, those based on smaller samples) should have a distribution that looks like what we would get if we added noise to less variable estimates (say, those based on larger samples). Departures from this pattern can allow us to pin down a model of the publication process from the metastudy data.

Andrews and Kasy apply these methods to data from Camerer et al. (2016), who replicated a set of experimental economics papers published in the *American Economic Review* and *Quarterly Journal of Economics* between 2011 and 2014. Andrews and Kasy find that estimates significant at the 5 percent level are over 30 times more likely to be published than are estimates not meeting this level of statistical significance! Reassuringly, the correction for publication bias reaches similar conclusions regardless of whether it is based on the replication study or the metastudy.

Publication bias leads to a kind of predictable disappointment: if only the papers with the "best" estimates are published, then replications of previous studies will systematically underperform relative to the published estimates. A similar phenomenon arises in other contexts. Say, for example, that a researcher uses data on test scores to estimate the effect of teachers on students' performance. The teacher with the largest estimated effect on students' test scores may be the best teacher or they may be someone who was lucky enough to have high-performing students. In [19], Andrews, Kitagawa, and McCloskey tackle the problem of learning about the underlying effects in situations where we are interested in selecting a "best" teacher, policy, or treatment. In [16], Andrews, Kitagawa, and McCloskey apply related statistical ideas to the problem of inference after structural breaks.

Conclusion

Amazingly, the areas we discuss above do not exhaust Isaiah's contributions. He has also managed to make contributions to microeconomic theory [5], to the study of external validity bias [13], and to the problem of inference with moment inequalities [18].

Conversations with Isaiah's coauthors on these and other projects reveal what both of us know from working with him: that he is a generous collaborator who treats his research not only as a means of expanding the frontier of knowledge, but also as an opportunity to teach and learn from others. We look forward to continuing to take advantage of these opportunities whenever we can!

References

Andrews, Isaiah, and James Stock. 2018. "Weak Instruments and What to Do About Them." Lectures, NBER Summer Institute 2018 Methods Lectures, Cambridge, MA, July 22, 2018.

Andrews, Marcellus W. 1986. "Competition and Macroeconomics: Some Implications of Schumpeter's Theory of Development for Macroeconomics." PhD diss. Yale University.

Bound, John, David A. Jaeger, and Regina M. Baker. 1995. "Problems with Instrumental Variables Estimation When the Correlation between the Instruments and the Endogenous Explanatory Variable Is Weak." *Journal of the American Statistical Association* 90 (430): 443–50.

Camerer, Colin F., Anna Dreber, Eskil Forsell, Teck-Hua Ho, Jürgen Huber, Magnus Johannesson, Michael Kirchler et al. 2016. "Evaluating Replicability of Laboratory Experiments in Economics." *Science* (351) 6280: 1433–36.

Cheyney University of Pennsylvania. 2021. Cheyney University. https://cheyney.edu/ (accessed November 16, 2021).

Conley, Timothy G., Christian B. Hansen, and Peter E. Rossi. 2012. "Plausibly Exogenous." *Review of Economics and Statistics* (94) 1: 260–72.

Gentzkow, Matthew. 2007. "Valuing New Goods in a Model with Complementarity: Online Newspapers." *American Economic Review* 97 (3): 713–44.

Gentzkow, Matthew and Jesse M. Shapiro. 2014. "Measuring the Sensitivity of Parameter Estimates to Sample Statistics." NBER Working Paper 20673.

Hastings, Justine, Ryan Kessler, and Jesse M. Shapiro. 2021. "The Effect of SNAP on the Composition of Purchased Foods: Evidence and Implications." *American Economic Journal: Economic Policy* 13 (3): 277–315.

Kasy, Maximilian. 2021. "Of Forking Paths and Tied Hands: Selective Publication of Findings, and What Economists Should Do about It." *Journal of Economic Perspectives* 35 (3): 175–92.

Moreira, Marcelo J. 2003. "A Conditional Likelihood Ratio Test for Structural Models." *Econometrica* 71 (4): 1027–48.

Smith, Cheryl I. 1984. "Theories of Fluctuations in the Monetary Economy: A Study of the Concepts of Money, Capital, and Investment in the Theories of Wicksell, Hayek, and Keynes." PhD diss. Yale University.

Smith, William C. 1966. "A Factor Analytic Study of the 'Test of Sound Reasoning.'" PhD diss. University of Missouri-Kansas City.

Smith, Joyce Annunciata Davis. 1971. "The Effects of Modeling Behaviors on the Communications of Clients in Counseling Situations." PhD Diss. University of Missouri - Kansas City.

Staiger, Douglas, and James H. Stock. 1997. "Instrumental Variables Regression with Weak Instruments." *Econometrica* 65 (3): 557–86.

Working, E.J. 1927. "What Do Statistical 'Demand Curves' Show?" *Quarterly Journal of Economics* 41 (2): 212–35.

Wright, Philip G. 1928. *The Tariff on Animal and Vegetable Oils*. New York, NY: Macmillan.

Journal of Economic Perspectives—Volume 36, Number 1—Winter 2022—Pages 191–198

Recommendations for Further Reading

Timothy Taylor

This section will list readings that may be especially useful to teachers of undergraduate economics, as well as other articles that are of broader cultural interest. In general, with occasional exceptions, the articles chosen will be expository or integrative and not focus on original research. If you write or read an appropriate article, please send a copy of the article (and possibly a few sentences describing it) to Timothy Taylor, preferably by e-mail at taylort@macalester.edu, or c/o *Journal of Economic Perspectives*, Macalester College, 1600 Grand Ave., Saint Paul, MN 55105.

Smorgasbord

The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2021 was awarded to given to David Card "for his empirical contributions to labour economics" and to Joshua D. Angrist and Guido W. Imbens "for their methodological contributions to the analysis of causal relationships." Each year, the prize committee produces a highly readable "Popular science" explanation, this year titled "Natural experiments help answer important questions," along with more specialized and longer "Scientific background" paper, this year titled "Answering Causal Questions Using Observational Data" (October 11, 2021, both available at https://www.nobelprize.org/prizes/economic-sciences/2021/summary/). From the "Popular Science" explanation: "This year's Laureates—David Card, Joshua Angrist and Guido Imbens—have shown that natural experiments can be used to

■ *Timothy Taylor is Managing Editor, Journal of Economic Perspectives, based at Macalester College, Saint Paul, Minnesota.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.36.1.191.

answer central questions for society, such as how minimum wages and immigration affect the labour market. They have also clarified exactly which conclusions about cause and effect can be drawn using this research approach. . . . Their solution is to use natural experiments—situations arising in real life that resemble randomised experiments. These natural experiments may be due to natural random variations, institutional rules or policy changes. In pioneering work from the early 1990s, David Card analysed some central questions in labour economics—such as the effects of a minimum wage, immigration and education—using this approach. . . . Natural experiments differ from clinical trials in one important way . . . [in] a natural experiment, the researcher also has access to data from treatment and control groups but, unlike a clinical trial, the individuals may themselves have chosen whether they want to participate in the intervention being offered. . . . In an innovative study from 1994, Joshua Angrist and Guido Imbens showed what conclusions about causation can be drawn from natural experiments in which people cannot be forced to participate in the programme being studied (nor forbidden from doing so). The framework they created has radically changed how researchers approach empirical questions using data from natural experiments or randomised field experiments."

Daniel Rees and Phurichai Rungcharoenkitkul discuss "Bottlenecks: causes and macroeconomic implications" (Bank of International Settlements, BIS Bulletin #48, November 11, 2021, https://www.bis.org/publ/bisbull48.htm). "Pandemic-induced supply disruptions have clearly been a major cause of bottlenecks, especially in the early stages of the global recovery. Producers who had severed relationships with suppliers early in the pandemic found it hard to re-establish them when demand picked up. Asynchronous lockdowns disrupted shipping, while sporadic virus outbreaks led to further dislocations. But there are also other causes. Unexpected natural events have intensified supply pressures. A lack of investment in the years leading up to the pandemic left some industries with little spare capacity. . . . Several factors have amplified the economic severity of bottlenecks. One is the shift in the composition of demand towards manufactured goods during the Covid recession and recovery. These goods are heavily reliant on inputs from other industries, leading to larger demand spillovers than from a services-led recovery. Manufactured goods (and their inputs) also tend to be relatively capital-intensive, making their short-run supply elasticity low as it takes time to expand productive capacity. . . . A second factor is . . . [that] [a]nticipation of product shortages and precautionary hoarding at different stages of supply chain have aggravated initial shortages (the "bullwhip effect"), leading to further incentives to build buffers. These behavioural changes have the potential to lead to feedback effects that exacerbate bottlenecks. . . . A third important background element is the lean structure of supply chains, which have prioritised efficiency over resilience in recent decades. These intricate networks of production and logistics were a virtue in normal times, but have become a shock propagator during the pandemic."

The Hutchins Center on Fiscal & Monetary Policy at the Brookings Institution and the Initiative on Global Markets at the University of Chicago Booth School of Business created a *Task Force on Financial Stability* report, led by Glenn Hubbard,

Donald Kohn, Laurie Goodman, Kathryn Judge, Anil Kashyap, Ralph Koijen, Blythe Masters, Sandie O'Connor, and Kara Stein (June 2021, https://www.brookings.edu/wp-content/uploads/2021/06/financial-stability_report.pdf). The report identifies five main areas of concern: Treasury markets, mutual funds, insurance companies, housing finance, and central clearing counterparties. "Our decision to focus on nonbank finance was reinforced in March 2020 when fear and uncertainty arising from the onset of the pandemic sparked a huge spike in demand for cash, disrupting a wide variety of credit markets, including markets for U.S. Treasury, corporate, and municipal bonds, mortgage-related securities, and commercial paper. The extraordinary growth in those markets had outgrown the capacity of the private sector to meet this outsized demand to sell securities to get cash. . . . The resulting disruptions threatened to cut off credit to households, businesses, and governments, which would have made an extremely serious economic situation much worse. Only very aggressive central bank intervention in the United States and elsewhere stabilized markets and restored credit flows."

Gaurav Nayyar, Mary Hallward-Driemeier, and Elwyn Davies have written *At Your Service? The Promise of Services-Led Development* (World Bank, September 2021. https://openknowledge.worldbank.org/handle/10986/35599). "Evidence suggests that manufacturing-led development in the past delivered the twin gains of productivity growth and large-scale job creation for the relatively unskilled. Underlying these were economies of scale, access to international markets, innovation, and supply chain linkages with other sectors, combined with the ability to leverage relatively unskilled labor with capital. Although services are labor intensive, they often require simultaneous production and consumption that precludes accessing larger markets. Their more limited ability to use capital to improve labor productivity also limits both scale economies and incentives to innovate. Conventional wisdom is therefore pessimistic about the prospects for services-led development. This book seeks to test that conventional wisdom. . . . The data show that services can deliver productivity growth—in several cases, growth that is higher than that of industry. What matters for the longer-term potential of services-led development is whether the features of industrialization that have enabled scale, innovation, and spillovers along with job creation for unskilled labor—as in East Asia—are increasingly shared by the services sector. . . . It is not necessarily the production of `goods' or `services' per se that matters but how these are produced."

A committee convened by the National Academy of Sciences discusses *Reducing the Impact of Dementia in America: A Decadal Survey of the Behavioral and Social Sciences* (September 2021, available with free registration). "More than 6 million people in the United States are currently living with Alzheimer's disease, a number that will rise to nearly 14 million by 2060 if current demographic trends continue. It is estimated that approximately one-third of older Americans have Alzheimer's or another dementia at death. . . . The primary economic costs of dementia to persons living with dementia and their families are (1) medical and long-term care costs, and (2) the value of unpaid caregiving provided by family (most commonly) and friends. Most estimates of these costs in the literature draw on such nationally

representative data sources as the Health and Retirement Study, the Medicare Current Beneficiary Survey, and Medicare claims data. An estimate of annual per-person costs for 2019, which includes health care and the value of unpaid care provided to persons with Alzheimer's disease, is approximately $81,000 ($31,000 is the value of the unpaid care). This estimate is about four times higher than the costs of the same care provided to similarly aged persons without the disease. . . . When aggregated to the U.S. population, the costs are estimated to have exceeded $500 billion in 2019 and are projected to increase to about $1.5 trillion by 2050."

Naomi R. Lamoreaux and John Joseph Wallis tell the story of "Economic Crisis, General Laws, and the Mid-Nineteenth-Century Transformation of American Political Economy" (*Journal of the Early Republic*, 41 (3), Fall 2021, pp. 403–433, https://muse.jhu.edu/article/803780). "In 1851, in the aftermath of an economic crisis that forced the state into default, Indiana rewrote its constitution to require that laws enacted by the legislature 'be general and of uniform operation throughout the state.' This directive may not seem remarkable from the standpoint of the twenty-first century; we take it for granted that that is what legislatures do. From the perspective of the mid-nineteenth century, however, the provision was groundbreaking. The first such mandate ever enacted, Indiana's innovation spread to almost all the other U.S. states over the next few decades . . . Before Indiana's innovation, the main business of legislatures was to enact special or private bills on behalf of specific individuals, organizations, and localities. The year before its new constitution was ratified, for example, Indiana's general assembly passed 550 acts. . . . About half benefited particular local governments, granting them permission to spend public funds, borrow money, levy taxes, set salaries, fees, duties, and meeting times for administrators and judges, and take a variety of other actions. Almost all the rest (nearly 40 percent) aided particular individuals or organizations. Some involved personal matters such as divorces, name changes, and the administration of decedents' estates, but the vast majority conveyed grants of economically valuable privileges such as corporate charters to people specifically named in the bills."

International Corporate Taxation

Alan Auerbach delivered the 2021 Martin S. Feldstein on "The Taxation of Business Income in the Global Economy" (*NBER Reporter*, September 2021, https://www.nber.org/reporter/2021number3/taxation-business-income-global-economy). "Fifty years ago, the top five companies by market capitalization were IBM, General Motors, AT&T, Standard Oil of New Jersey (Esso, the predecessor of today's ExxonMobil), and Eastman Kodak. . . . These were companies that 'made things' in identifiable locations, to a large extent in the United States. If we shift to today, we see another five familiar names, all giant companies: Apple, Microsoft, Amazon, Alphabet (Google's parent), and Facebook. These companies are worldwide multinationals, relying very heavily on the use of intellectual property in the

goods and services they provide . . . In the last half century, the share of intellectual property measured in US nonfinancial corporate assets more than doubled, according to the Fed's Financial Accounts of the United States. That's probably a conservative estimate, because the measurement of intellectual property is a fairly narrow one here. The share of before-tax US corporate profits coming from overseas operations nearly quintupled, according to data from the Bureau of Economic Analysis. US companies have become much more multinational in character, not just selling things abroad, but making them abroad as well. And the share of cross-border equity ownership has steadily increased, to the point that foreign individuals and companies account for a significant fraction of US companies' share ownership."

Ruud de Mooij, Alexander Klemm, and Victoria Perry have edited a collection of 16 essays about *Corporate Income Taxes Under Pressure: Why Reform is Needed and How it Could be Designed* (International Monetary Fund, 2021, https://www.elibrary.imf.org/view/books/071/28329-9781513511771-en/28329-9781513511771-en-book.xml). For a sample, Narine Nersesyan explains in Chapter 3 "The Current International Tax Architecture: A Short Primer." "The generally applied tax architecture for determining where profits are taxed is now nearly 100 years old—designed for a world in which most trade was in physical goods, trade made a less significant contribution to world GDP, and global value chains were not particularly complex. . . . The current international tax framework is based on the so-called '1920's compromise'. In very basic outline, under the 'compromise' the primary right to tax active business income is assigned where the activity takes place—in the 'source' country—while the right to tax passive income, such as dividends, royalties and interest, is given up to the 'residence' country—where the entity or person that receives and ultimately owns the profit resides. The system has, however, evolved in ways that considerably deviate from this historic 'compromise,' and international tax arrangements currently rest on a fragile and contentious balance of taxing rights between residence and source countries. . . . While domestic laws of each individual country set out the rules . . . the international taxation system is—very importantly—overlain with a network of more than 3,000 bilateral double-taxation treaties. . . . The key role of the international tax architecture is to govern the allocation of taxing rights between the potential tax-claiming jurisdictions to avoid both excessive taxation of a single activity and a nontaxation of a business activity."

Eviction Economics

The Summer 2021 issue of *Evidence Matters*, published by the US Department of Housing and Urban Development, is devoted to the theme of "Evictions," with three readable articles heavily footnoted with references to published studies, written by Dana Goplerud and Craig Pollack (https://www.huduser.gov/portal/periodicals/em/summer21/index.html). From the first article, "Affordable Housing, Eviction, and Health," "Based on data from 2016, the most recent year available, landlords

filed an estimated 3.7 million evictions, with about 8 out of every 100 renter households receiving eviction notices. Hepburn and Rutan place the scale of evictions in context by comparing them with the 2.8 million foreclosure starts at the height of the Great Recession. As staggering as these figures are, many more renters may be forced to exit their housing without a formal filing or between the time of a filing and a judgment. . . . Nonpayment of rent is the primary reason for eviction, which itself can arise from various causes, including rising rents combined with stagnant income growth and persistent poverty, job or income loss, or a sudden economic shock such as a health emergency or a car breakdown. Other reasons include lease violations, which can be technical in nature; property damage; and disruptions, such as police calls. Landlords, for their own reasons, may force tenants to move, either informally or through a legal "no-fault" eviction. Renters often are evicted over relatively small amounts of money—in many cases, less than a full month's rent. . . . The Milwaukee Area Renters Study found higher rates of eviction for African-American, Latinx, and lower-income renters and renters with children. Neighborhood crime and eviction rates, the number of children in a household, and 'network disadvantage'—defined . . . as 'the proportion of one's strong ties to people who are unemployed, addicted to drugs, in abusive relationships, or who have experienced major, poverty-inducing events (e.g., incarceration, teenage pregnancy) to increase his or her propensity for eviction'—are factors associated with an increased likelihood of eviction."

One part of the Coronavirus Aid, Relief, and Economic Security Act (the CARES Act) that became law on March 27, 2020, was a national moratorium on evictions from rental housing, which was later extended by the Centers for Disease Control. The Eviction Lab at Princeton University has been tracking the results. For example, Anne Kat Alexander and Sarah Lee published a "Preliminary Analysis: A Year of Eviction Moratoria" (March 29, 2021, https://evictionlab.org/one-year-of-eviction-moratoria). Jasmine Rangel, Jacob Haas, Emily Lemmerman, Joe Fish, and Peter Hepburn followed up with a "Preliminary Analysis: 11 months of the CDC Moratorium" (August 21, 2021, https://evictionlab.org/eleven-months-cdc/). They show considerable variation in state and local rules as well, but summarize: "In total, we estimate that federal, state, and local policies helped to prevent at least 2.45 million eviction filings since the start of the pandemic (March 15, 2020)."

Elijah de la Campa, Vincent J. Reina, and Christopher Herbert published "How Are Landlords Faring During the COVID-19 Pandemic? Evidence from a National Cross-Site Survey" (Joint Center for Housing Studies of Harvard University, August 2021, https://www.jchs.harvard.edu/sites/default/files/research/files/harvard_jchs_covid_impact_landlords_survey_de_la_campa_2021.pdf). Based on a national survey of landlords, "[t]he share of landlords collecting 90 percent or more of yearly rent fell 30 percent from 2019 to 2020. . . . Ten percent of all landlords collected less than half of their yearly rent in 2020, with smaller landlords (1–5 units) most likely to have tenants deeply behind on rental payments. . . . The share of landlords deferring maintenance and listing their properties for sale also increased in 2020 (5 to 31 percent and 3 to 13 percent, respectively) . . . "

The research group at JP Morgan Chase Institute published "How did landlords fare during COVID?" (October 2021, https://www.jpmorganchase.com/institute/research/household-debt/how-did-landlords-fare-during-covid). "Between the Emergency Rental Assistance Program and the American Rescue Plan Act, $46.5 billion of rental assistance has been made available by the federal government for states and localities to distribute. As of the end of September [2021], less than a quarter of the funds have been distributed. The distribution of these funds has been hampered by onerous paperwork requirements for both tenants and landlords to prove that tenants meet strict requirements to qualify for assistance, including matching information from the renter and the landlord."

Interviews with Economists

Allison Schrager speaks with Luigi Zingales in "Break Up Big Tech? A conversation about the future of the industry" (*City Journal*, September 21, 2021, https://www.city-journal.org/a-conversation-about-the-future-of-big-tech-industry). "I think the problem is that we treat Big Tech as one big issue, and we say we need to break them up. Rather, what we should do depends on what we want to accomplish, and what sector in the industry we're taking about. Let's start with social media. I think the government should have tried to stop Facebook's acquisition of Instagram and WhatsApp, but I am not sure that breaking them up now would make a difference in the long term. If there are big network externalities, separating Facebook from Instagram would be just a temporary measure, because eventually only one of the two will prevail. . . . We should separate the two key functions Facebook performs: sharing of information and editing of information. . . . The problem isn't social media; it's the business model, which is to get people addicted to platforms."

Douglas Clement serves as interlocutor in "Rucker Johnson interview: Powering potential," subtitled "Rucker Johnson on school finance reform, quality pre-K, and integration" (*For All*, Federal Reserve Bank of Minneapolis, Fall 2021, https://www.minneapolisfed.org/article/2021/rucker-johnson-interview-powering-potential). "Today, about 75 percent of per pupil spending disparities are between states (rather than between districts within states). And we've witnessed that inequality in school spending has risen since 2000. After three decades of narrowing—the '70s, '80s, and '90s—primarily due to the state school finance reforms . . . there has been a significant rise in inequality, especially sharply following the Great Recession. What I want to highlight here is the current disparities nationwide in school resources. School districts with the most students of color have about 15 percent less per pupil funding from state and local sources than predominantly White, affluent areas, despite having much greater need due to higher proportions of poverty, special needs, and English language learners. . . . [S]chools with a high level of Black and Latino students have almost two times as many first-year teachers as schools with low minority enrollment. And minority students are more likely to be taught by inexperienced teachers than experienced ones in 33 states across the country. . . . [O]nly

a third of public schools with high Black and Latino enrollment offer calculus. Courses like that are gateways to majoring in STEM in college and having a STEM career."

Discussion Starters

The World Bank has ended its well-known "Doing Business" reports. The reasons are available in a report from the outside law firm of WilmerHale, "Investigation of Data Irregularities in Doing Business 2018 and 2020—Investigation Findings and Report to the Board of Executive Directors" (September 16, 2021, available at https://www.worldbank.org/en/news/statement/2021/09/16/statement-on-release-of-investigation-into-data-irregularities-in-doing-business-2018-and-2020). "To that end, we undertook to understand: (1) how improper changes to the data for China (*Doing Business 2018*) and Saudi Arabia, the United Arab Emirates, and Azerbaijan (*Doing Business 2020*) were effected; (2) who at the Bank directed, implemented, or knew about the changes to the data. . . . and (3) what internal circumstances, whether related to policies, personnel, or culture, allowed for the changes to take place." The Summer 2015 issue of this journal included a two-paper symposium about the Doing Business reports.

Steve Kaczynski and Scott Duke Kominers offer a primer in "How NFTs Create Value" (*Harvard Business Review*, November 10, 2021, https://hbr.org/2021/11/how-nfts-create-value). "As the name 'non-fungible token' suggests, each NFT is a unique, one-of-a-kind digital item. They're stored on public-facing digital ledgers called blockchains, which means it's possible to prove who owns a given NFT at any moment in time and trace the history of prior ownership. Moreover, it's easy to transfer NFTs from one person to another—just as a bank might move money across accounts—and it's very hard to counterfeit them. Because NFT ownership is easy to certify and transfer, we can use them to create markets in a variety of different goods. . . . Because blockchains are programmable, it's possible to endow NFTs with features that enable them to expand their purpose over time, or even to provide direct utility to their holders. In other words, NFTs can do things—or let their owners do things—in both digital spaces and the physical world. In this sense, NFTs can function like membership cards or tickets, providing access to events, exclusive merchandise, and special discounts—as well as serving as digital keys to online spaces where holders can engage with each other. Moreover, because the blockchain is public, it's even possible to send additional products directly to anyone who owns a given token. All of this gives NFT holders value over and above simple ownership—and provides creators with a vector to build a highly engaged community around their brands."

IMF Publications

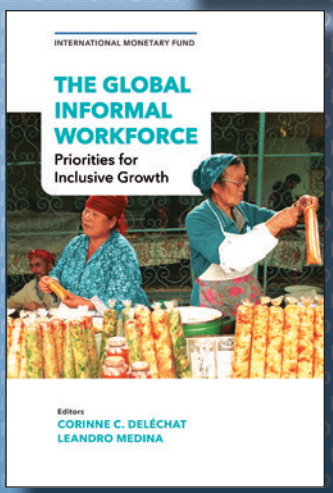

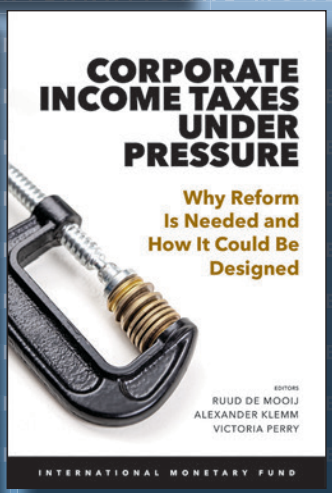

Read the latest macroeconomic research and analysis from the IMF.

IMF.org/pubs

INTERNATIONAL MONETARY FUND

AMERICAN ECONOMIC ASSOCIATION

BEST PRACTICES FOR ECONOMISTS:

BUILDING A MORE DIVERSE, INCLUSIVE, AND PRODUCTIVE PROFESSION

A more diverse profession fosters a more vibrant discipline.

See practical suggestions and supporting research from the *AEA Task Force on Best Practices* regarding actions all economists can and should take. With intention, we CAN make change.

- **CONDUCTING RESEARCH**
- **SERVING AS COLLEAGUES**
- **WORKING WITH STUDENTS**
- **LEADING DEPARTMENTS AND WORKPLACES**

More information at
www.aeaweb.org/resources/bestpractices

AEA Training Opportunity!

EDUCATE Workshop

EXPANDING DIVERSITY in UNDERGRADUATE CLASSES with ADVANCEMENTS in (the) TEACHING (of) ECONOMICS

Overview

This workshop provides opportunities for instructors of undergraduate courses to take part in course design activities and experience pedagogical strategies that will engage all of the students they teach. Attendees will have opportunities to identify learning objectives that focus on the students' ability to "do economics" and to participate in pedagogical practices that enable students to be active participants in economic analysis.

All accepted applicants are expected to fully engage with each of the three phases of the overall program including constructing learning objectives, studying pedagogical practices that are collaborative and inclusive including cooperative learning, engaging lectures, data integration, and classroom experiments, and integrating lessons learned into their own courses. Participants will be provided opportunities to share their work at the 2023 CTREE and ASSA meetings.

Participants will engage with issues of diversity and inclusion throughout the workshop including opportunities to think critically about course goals and learning outcomes, their relationship to pedagogical choices and assessment, and how such decisions might have disparate effects on those of different races, genders, and ethnicities. In addition, attendees will learn how to discuss the sensitive topics that are an important part of the economics classroom.

Eligibility

Applicants must be scheduled to teach during the following fall and spring semesters. Preference will be given to those less than 6 years since PhD. To meet the goals of increasing diversity in the profession, the cohort will be chosen to represent a diverse set of institutions and instructors.

Cost

Accepted applicants must make an electronic payment of $100 to confirm their spot in the workshop. The application portal opens in January 2022 with a rolling acceptance procedure that continues until all workshop slots are filled. Workshop details and the application portal are available at https://www.aeaweb.org/go/educate-workshop.

EDUCATE
Workshop
(A face-to-face and Canvas-supported course)

June 3-4-5, 2022
Chicago, IL

For more information go to
https://www.aeaweb.org/go/educate-workshop

EDUCATE is sponsored by the AEA Outreach Task Force and the AEA Committee on Economic Education.

AEA EconHarmony

aeaweb.org/econharmony

EconHarmony allows AEA members with similar research interests to post information about their papers and join others in submitting a session proposal for the ASSA Annual Meeting.

Proposals for sessions (three to four related papers submitted by a session organizer) have a higher probability of inclusion in the meeting than individual papers. Visit the AEA's EconHarmony website for additional information.

Deadline for ASSA 2023 session proposals:

April 15, 2022

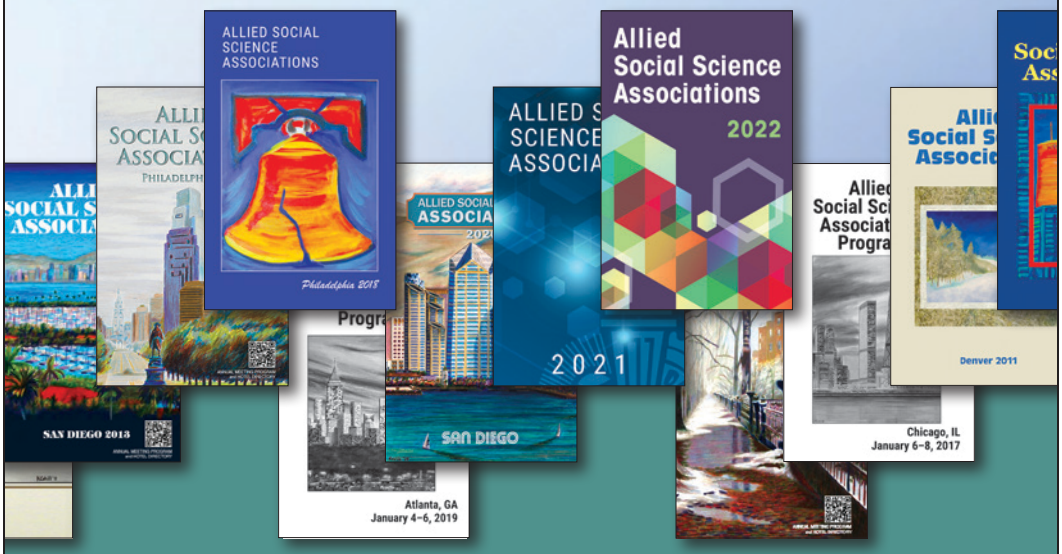

Collaborate on a session for the next ASSA meeting!

AEA INITIATIVES FOR DIVERSITY AND INCLUSION

The American Economic Association is committed to the continued improvement of the professional climate in economics. In cooperation with key committees, the Association has launched several new initiatives to support and promote diversity and inclusion in our profession.

1 AEA Award for Outstanding Achievement in Diversity and Inclusion

This annual award recognizes departments and organizations that demonstrate outstanding achievement in diversity and inclusion practices. Focus will be on those applicants that take productive steps to establish new programs and procedures to create an inclusive environment, and to increase the participation of underrepresented racial/ethnic minorities, women, and LGBTQ+ individuals.

2 Departmental Seed Grants for Innovation in Diversity and Inclusion

These grants, in amounts up to $5,000, will be awarded to economics departments to help establish new bridge programs or training programs for underrepresented minorities (URM). For example, a department might create a mentoring program for URM graduate or undergraduate students, create opportunities for URM students to do meaningful research assistant work, or start a program allowing URM students who need additional preparation for graduate school to take a lighter class load in the first year or to take core economics courses over two years.

3 The Andrew Brimmer Undergraduate Essay Prize

Thanks to the generosity of an anonymous donor, this paper prize has been established in honor of Andrew Brimmer, the first African American to serve on the Board of Governors of the Federal Reserve. The annual award will be presented to an undergraduate student at a US-based institution of higher learning majoring in economics, political science, public policy, or related fields for the best essay on the "economic well-being of Black Americans." The winner will receive a check for $1,000.

4 URM Travel Grants

This award is open to junior economics faculty members from traditionally underrepresented groups in the economics profession. The grants will advance career and professional development by defraying the costs of travel, lodging, and conference registration to attend the annual ASSA Meeting.

5 Small Group Breakfast Meeting for URM

Each year at the ASSA Meeting there will be a breakfast held with scholars from underrepresented minorities and prominent economists in attendance. The goal is to allow URM scholars access to AEA journal editors, executive board members, thought leaders in specific areas of economics, or other economists for the purpose of addressing issues of access to journals, conferences, and networks that are often out of reach for URM scholars.

6 Professional Development Grant for URM

This $2,000 grant was established to help advance the career and professional development of URM in the field of Economics. The award is open to eligible junior economics faculty members. Entrants to the essay competition should detail their research and how it relates to economics education.

For more details and information regarding how to apply for any of these initiatives, please visit

www.aeaweb.org/go/diversity-initiatives

The *Journal of Economic Perspectives*: Proposal Guidelines

Considerations for Those Proposing Topics and Papers for *JEP*

Articles appearing in the journal are primarily solicited by the editors and associate editors. However, we do look at all unsolicited material. Due to the volume of submissions received, proposals that do not meet *JEP*'s editorial criteria will receive only a brief reply. Proposals that appear to have *JEP* potential receive more detailed feedback. Historically, about 10–15 percent of the articles appearing in our pages originate as unsolicited proposals.

Philosophy and Style

The *Journal of Economic Perspectives* attempts to fill part of the gap between refereed economics research journals and the popular press, while falling considerably closer to the former than the latter. **The focus of *JEP* articles should be on understanding the central economic ideas of a question, what is fundamentally at issue, why the question is particularly important, what the latest advances are, and what facets remain to be examined.** In every case, articles should argue for the author's point of view, explain how recent theoretical or empirical work has affected that view, and lay out the points of departure from other views.

We hope that most *JEP* articles will offer a kind of intellectual arbitrage that will be useful for every economist. For many, the articles will present insights and issues from a specialty outside the readers' usual field of work. For specialists, the articles will lead to thoughts about the questions underlying their research, which directions have been most productive, and what the key questions are.

Articles in many other economics journals are addressed to the author's peers in a subspecialty; thus, they use tools and terminology of that specialty and presume that readers know the context and general direction of the inquiry. By contrast, **this journal is aimed at all economists, including those not conversant with recent work in the subspecialty of the author.** The goal is to have articles that can be read by 90 percent or more of the AEA membership, as opposed to articles that can only be mastered with abundant time and energy. Articles should be as complex as they need to be, but not more so. Moreover, the necessary complexity should be explained in terms appropriate to an audience presumed to have an understanding of economics generally, but not a specialized knowledge of the author's methods or previous work in this area.

The *Journal of Economic Perspectives* is intended to be scholarly without relying too heavily on mathematical notation or mathematical insights. In some cases, it will be appropriate for an author to offer a mathematical derivation of an economic relationship, but in most cases it will be more important that an author explain why a key formula makes sense and tie it to economic intuition, while leaving the actual derivation to another publication or to an appendix.

JEP does not publish book reviews or literature reviews. Highly mathematical papers, papers exploring issues specific to one non-U.S. country (like the state of agriculture in Ukraine), and papers that address an economic subspecialty in a manner inaccessible to the general AEA membership are not appropriate for the *Journal of Economic Perspectives.* Our stock in trade is original, opinionated perspectives on economic topics that are grounded in frontier scholarship. If you are not familiar with this journal, it is freely available on-line at www.aeaweb.org/journals/jep.

Guidelines for Preparing *JEP* Proposals

Almost all *JEP* articles begin life as a two- or three-page proposal crafted by the authors. If there is already an existing paper, that paper can be sent to us as a proposal for *JEP.* However, given

the low chances that an unsolicited manuscript will be published in *JEP*, no one should write an unsolicited manuscript intended for the pages of *JEP*. **Indeed, we prefer to receive article proposals rather than completed manuscripts.** The following features of a proposal seek to make the initial review process as productive as possible while minimizing the time burden on prospective authors:

- Outlines should begin with a paragraph or two that precisely states the main thesis of the paper.
- After that overview, an explicit outline structure (I., II., III.) is appreciated.
- The outline should lay out the expository or factual components of the paper and indicate what evidence, models, historical examples, and so on will be used to support the main points of the paper. The more specific this information, the better.
- The outline should provide a conclusion.
- Figures or tables that support the article's main points are often extremely helpful.
- The specifics of fonts, formatting, margins, and so forth do not matter at the proposal stage. (This applies for outlines and unsolicited manuscripts).
- Sample proposals for (subsequently) published *JEP* articles are available on request.
- For proposals and manuscripts whose main purpose is to present an original empirical result, please see the specific guidelines for such papers below.

The proposal provides the editors and authors an opportunity to preview the substance and flow of the article. For proposals that appear promising, the editors provide feedback on the substance, focus, and style of the proposed article. After the editors and author(s) have reached agreement on the shape of the article (which may take one or more iterations), the author(s) are given several months to submit a completed first draft by an agreed date. This draft will receive detailed comments from the editors as well as a full set of suggested edits from *JEP*'s Managing Editor. Articles may undergo more than one round of comment and revision prior to publication.

Readers are also welcome to send e-mails suggesting topics for *JEP* articles and symposia and to propose authors for these topics. If the proposed topic is a good fit for *JEP*, the *JEP* editors will work to solicit paper(s) and author(s).

Correspondence regarding possible future articles for *JEP* may be sent (electronically please) to the assistant managing editor, Alexandra Szczupak at a.szczupak@aeapubs.org. Papers and paper proposals should be sent as Word or pdf e-mail attachments.

Guidelines for Empirical Papers Submitted to *JEP*

JEP is not primarily an outlet for original, frontier empirical contributions; that's what refereed journals are for! Nevertheless, *JEP* occasionally publishes original empirical analyses that appear uniquely suited to the journal. In considering such proposals, the editors apply the following guidelines (in addition to considering the paper's overall suitability):

1) The paper's main topic and question must not already have found fertile soil in refereed journals. *JEP* can serve as a catalyst or incubator for the refereed literature, but it is not a competitor.

2) In addition to being intriguing, the empirical findings must suggest their own explanations. If the hallmark of a weak field journal paper is the juxtaposition of strong claims with weak evidence, a *JEP* paper presenting new empirical findings will combine strong evidence with weak claims. The empirical findings must be robust and thought provoking, but their interpretation should not be portrayed as the definitive word on their subject.

3) The empirical work must meet high standards of transparency. *JEP* strives to only feature new empirical results that are apparent from a scatter plot or a simple table of means. Although *JEP* papers can occasionally include regressions, the main empirical inferences should not be regression-dependent. Findings that are not almost immediately self-evident in tabular or graphic form probably belong in a conventional refereed journal rather than in *JEP*.

Don't miss ...

CTREE 2022

The Eleventh Annual Conference on Teaching and Research in Economic Education (CTREE)

Plenary Speakers

Lisa Cook
Michigan State University

Nancy Rose
Massachusetts Institute of Technology

June 1–3, 2022

Chicago, IL | Palmer House Hilton

Sponsors:
AEA Committee on Economic Education
Journal of Economic Education
Federal Reserve Bank of Chicago

www.aeaweb.org/ctree/2022

Why an AEA Ombudsperson?

- **Informal**
- **Confidential**
- **Impartial**
- **Independent**

If you have experienced or witnessed incidents of harassment or discrimination, the AEA Ombudsperson can help guide you.

The AEA Ombudsperson enables AEA members to speak anonymously with a neutral third party about potential violations of AEA policies and determine if a formal or an informal report is appropriate. Whether the member has been personally impacted by misconduct or was witness to it, the ombudsperson is a confidential resource who can advise the member about additional resources and possible next steps. No action is ever taken without the complainant's authorization.

Need Assistance? Please Reach Out.

Leto Copeley, AEA Ombudsperson

aeaombuds@whiteandstradley.com

919-844-0400

https://copeleylaw.com/aea-ombudsperson-contact/

White & Stradley, PLLC
3105 Charles B Root Wynd
Raleigh, NC 27612

Formal Complaint?

Contact with the AEA Ombudsperson does not constitute official notice to the AEA or trigger a formal complaint. Formal complaints are filed through the AEA and are reviewed and investigated by the AEA Ethics Committee. The AEA Ombudsperson can provide guidance about the formal process and may separately advise the Association about observations of emerging patterns, repeat perpetrators, or systemic indicators.

Find more information about these policies and the AEA Ombudsperson at

www.aeaweb.org/ombudsperson

AMERICAN ECONOMIC ASSOCIATION

The AEA is dedicated to improving the climate of the economics profession by addressing harassment and discrimination, which are in violation of AEA policies and the AEA Code of Conduct.

Join Us in **Shaping The Future of Economics**

Become an AEA Member!

As a member, your partnership with the AEA provides you with distinct career advantages and enables us to elevate the profession for all economists.

- Build lifelong relationships with mentors and peers
- Access essential research and job resources
- Publish and present your research to colleagues
- Gain meaningful support and advice on your work
- Collaborate on solutions to issues you care about

Today our community consists of 20,000+ individuals who are committed to the advancement of economics and its enduring contributions to society. We welcome the opportunity to help you advance your career in economics.

Aim High. Achieve More. Make a Difference.

AEA Journals @ AEA Journals / @ AEA Information

AMERICAN ECONOMIC ASSOCIATION

Join today!

aeaweb.org/membership

The American Economic Association

Correspondence relating to advertising, business matters, permission to quote, or change of address should be sent to the AEA business office: aeainfo@vanderbilt.edu. Street address: American Economic Association, 2014 Broadway, Suite 305, Nashville, TN 37203. For membership, subscriptions, or complimentary access to JEP articles, go to the AEA website: http://www.aeaweb.org. Annual dues for regular membership are $24.00, $34.00, or $44.00, depending on income; for an additional fee, you can receive this journal, or any of the Association's journals, in print. Change of address notice must be received at least six weeks prior to the publication month.

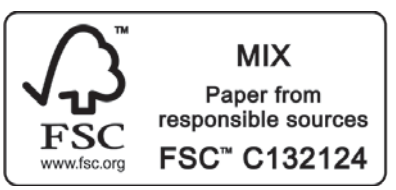

The Journal of

Economic Perspectives

Winter 2022, Volume 36, Number 1

Symposium

Economies of Africa

Articles

Feature

AMERICAN ECONOMIC ASSOCIATION