American Economic Review 2019, 109(3): 1032–1079
https://doi.org/10.1257/aer.20161385

The Arrival of Fast Internet and Employment in Africa†

By Jonas Hjort and Jonas Poulsen*

To show how fast Internet affects employment in Africa, we exploit the gradual arrival of submarine Internet cables on the coast and maps of the terrestrial cable network. Robust difference-in-differences estimates from 3 datasets, covering 12 countries, show large positive effects on employment rates—also for less educated worker groups—with little or no job displacement across space. The sample-wide impact is driven by increased employment in higher-skill occupations, but less-educated workers' employment gain less so. Firm-level data available for some countries indicate that increased firm entry, productivity, and exporting contribute to higher net job creation. Average incomes rise. (*JEL* F14, J23, J24, J63, L86, O15, O33)

Traditional trade theory predicts a decrease in inequality in developing countries during periods of integration in the global economy. The slow economic progress of poor workers in many parts of Africa, Asia, and Latin America during the last few decades, therefore, surprised economists. Two potential explanations were proposed and compared: skill-biased technological change (SBTC) and features of international trade—such as outsourcing (see e.g., Feenstra and Hanson 1996, 1999, 2003) and quality upgrading (see e.g., Verhoogen 2008; Frías, Kaplan, and Verhoogen 2009)—that could alter the logic underlying expectations of job growth and greater equality in unskilled labor-abundant countries post-integration (Feenstra and Hanson 2003; Goldberg and Pavcnik 2007; Harrison, McLaren, and McMillan 2011; Goldberg 2015). Two decades of research led to wide agreement that both explanations play a role, and that they probably interact (Wood 1995; Acemoglu 2003; Attanasio, Goldberg, and Pavcnik 2004; Burstein, Cravino, and

* Hjort: Columbia University, Uris Hall 622, 3022 Broadway, New York, NY 10027, BREAD, CEPR, and NBER (email: hjort@columbia.edu); Poulsen: Uppsala University, Box 514, 751 20 Uppsala, Sweden (email: jonas.poulsen@nek.uu.se). This paper was accepted to the *AER* under the guidance of Marianne Bertrand, Coeditor. We are grateful to five anonymous referees for insightful comments that significantly improved the paper. We also thank Adrian Adermon, Sebastian Axbard, Niklas Bengtsson, Martina Björkman-Nyqvist, Greg Bruich, Esther Duflo, Ray Fisman, Oded Galor, Francois Gerard, Amit Khandelwal, Erik Lindqvist, Mushfiq Mobarak, Kalle Moene, Eva Mörk, Anders Olofsgård, Michel Serafinelli, Kjetil Storesletten, Eric Verhoogen, Tim Waters, Frank Windmeijer, Chris Woodruff, and especially Henrik Sigstad and Matthieu Teachout for many helpful conversations, and seminar participants at the Barcelona Summer Forum, Bocconi, Brown, Columbia, Harvard, IGC Growth Week, NBER Summer Institute, Oslo, Oxford, Stockholm School of Economics, UBC, and Uppsala for comments and suggestions. Sawal Acharya, Patrick Kennedy, and Roxanne Rahnama provided great research assistance. We thank Akamai, Steve Song, and the World Bank for data access. Hjort thanks the Center for Development Economics and Policy at Columbia University for financial support. Any errors are our own. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.
† Go to https://doi.org/10.1257/aer.20161385 to visit the article page for additional materials and author disclosure statement(s).

Vogel 2013; Koren and Csillag 2016; Raveh and Reshef 2016). But this conclusion was built on studies of *trade-induced* technological change. To date, there is no direct evidence on the average and distributional economic effects in poor countries of the spread of the modern information and communication technologies (ICT) that help explain increasing inequality in rich countries' labor markets.

In this paper, we estimate how fast Internet—"the greatest invention of our time" (the *Economist* 2012)—affects poor countries' economies.[1] To do so, we compare individuals and firms in locations in Africa that are on the terrestrial network of Internet cables to those that are not. We compare these two groups during the gradual arrival on the coast of submarine cables from Europe that greatly increase speed and capacity on the terrestrial network. We show how employment rates, occupational employment shares, job inequality across the educational attainment range, and the underlying extensive (Internet take-up) and intensive (Internet speed) margin, respond. We also show evidence on three particular mechanisms through which take-up and speed may affect employment: changes in firm entry, changes in productivity in existing firms, and changes in exporting. Finally, we show how average incomes in locations that see changes in employment patterns with the arrival of fast Internet respond.

It has been difficult to study SBTC directly because, other than in local experiments, ICT technologies are not randomly allocated, but introduced where economic benefits are expected. While this is true everywhere, developing countries additionally tend to lack systematic and detailed labor market and firm-level data, especially in the poorest regions of the world, where the economic environment differs the most from the West (see Katz and Autor 1999, Bond and Van Reenen 2007, and Goldin and Katz 2007 for overviews of the SBTC literature on rich countries). We overcome the first obstacle by interacting time variation generated by the gradual arrival of submarine Internet cables at landing points on Africa's coast in the late 2000s and early 2010s with cross-sectional variation in whether a given location is connected to the terrestrial "backbone" network that starts at the landing point cities.[2] We overcome the second obstacle by combining employment data from representative household surveys (panels at location level) from 12 African countries with a combined population of roughly half a billion people with firm-level datasets (panels at firm or location level) from Ethiopia, South Africa, and a group of 6 African countries.[3] We use the firm-level data to show evidence on three especially important mechanisms—firm entry, productivity, and exporting—through which fast Internet may affect employment.[4] We also use data on Internet speed and

[1] We are not aware of existing causal evidence on this relationship. See World Bank (2016) for an overview of the existing correlational evidence, and more details below.

[2] During this period, each coastal country effectively had its own separate backbone network, as explained in Section I.

[3] One household survey (Afrobarometer) covers Benin, Ghana, Kenya, Madagascar, Mozambique, Nigeria, Senegal, Tanzania, and South Africa; and the other (DHS) Benin, D.R. Congo, Ghana, Kenya, Namibia, Nigeria, Togo, and Tanzania. We refer to these 12 countries jointly as "Africa" for simplicity. We also use a labor force survey from South Africa, and firm data from Ethiopia and Ghana, Kenya, Mauritania, Nigeria, Senegal, and Tanzania.

[4] We study a diverse subset of the world's poorest countries and a transformative technology that may affect employment patterns through many different channels. Data limitations thus prevent us from investigating *all* such potential channels, or determining what share of the identified changes in employment patterns firm entry, productivity, and exporting account for. The literature on information frictions in developing countries, for example, hints at additional mechanisms that may also play a role (Bloom and Van Reenen 2007; Antràs, Garicano, and Rossi-Hansberg,

take-up of the Internet to tie the reduced form estimates to the intensive and extensive margin of use. Finally, we use data on night lights from satellite images to study how fast Internet ultimately affects (a proxy for) average incomes.[5]

Our approach differs from much of the related literature in that employment rates, rather than wages (among the employed), are our primary outcomes of interest. This is partly for data availability reasons, but it is also a sensible choice in a developing region context. *Job inequality* includes inequality in human capital accumulation, future labor market prospects, and income that is due to (i) current (un)employment—a component that focusing on wage inequality would miss (see e.g., Magruder 2012, Hardy and McCasland 2015)—and (ii) the *quality* of the individual's job (if any) (see also Davis and Harrigan 2011; Card, Heining, and Kline 2013; Card, Cardoso, and Kline 2016). Moreover, changes in the probability of a worker being employed in a position belonging to a given type of occupation are informative not only of demand for qualified workers, but also of trends in structural change in developing economies.

Our three main sets of results are as follows.[6] First, we find that the probability that an individual is employed increases by 6.9 and 13.2 percent in the two groups of countries covered by our household survey datasets, and by 3.1 percent in South Africa, when fast Internet becomes available. We show that the increase in employment in connected areas is not due to displacement of jobs in unconnected areas.

Second, in both South Africa and the eight poorer countries covered by a household survey that records occupation information, we find that the probability of being employed in a position belonging to a skilled occupation increases substantially, but the probability of holding an unskilled job is statistically unaffected when fast Internet becomes available. While the impact on overall trends in structural change is likely modest, fast Internet appears to shift employment shares towards higher-productivity occupations.

Third, employment inequality if anything falls when fast Internet arrives in Africa. The percentage point increase in the probability of having a job is, for example, of comparable magnitude for those who only completed primary school and those with secondary or tertiary education in all three of our samples. The estimated increase in employment in a *skilled* occupation is biggest for those with tertiary education in the group of countries covered by a detailed household survey, but is comparable in magnitude across the educational attainment range in South Africa. In both these samples, those with only primary school see increased employment in unskilled occupations.

To compare these results to the existing evidence on recent SBTC in developed countries, we distinguish between the skill level of *jobs* and *workers*. Our findings suggest that fast Internet in Africa affects employers' relative demand for skilled and unskilled positions similarly to "computerization" and broadband Internet in rich

2008; Magruder 2010; Beaman and Magruder 2012; Allen 2014; Eaton et al. 2015; Hardy and McCasland 2015; Atkin et al. 2017; Mitra et al. forthcoming).

[5] See Henderson, Storeygard, and Weil (2012); Bleakley and Lin (2012); Michalopoulos and Papaioannou (2013, 2018); Lowe (2014) on night lights as a proxy for average incomes.

[6] That a given cable reaches different countries at different times and in a geographically determined order, and that we consider ten different cables, a priori lower concerns about nonparallel prior trends in economic outcomes in locations on versus off the backbone network. The collection of datasets we use enables an extensive battery of tests that supports a causal interpretation of our results.

countries (Autor, Katz, and Krueger 1998; Autor, Levy, and Murnane 2003; Autor, Katz, and Kearney 2008; Goos, Manning, and Salomons 2014; Katz and Margo 2014; Akerman, Gaarder, and Mogstad 2015), although the increase in overall employment and employment in skilled occupations is notably bigger in Africa. In contrast, while ICT tends to increase inequality across the educational attainment range in rich countries, fast Internet, if anything, decreases (un)employment inequality in Africa.[7] These results underscore that the factor bias of new technologies varies by context.

The changes in employment patterns observed when submarine Internet cables arrive in Africa occur through a combination of extensive margin (new users) and intensive margin (different use of the Internet by existing users) responses. We find a large and significant increase in net firm entry (in South Africa), notably in sectors that use ICT extensively (e.g., finance), and in the productivity of existing manufacturing firms (in Ethiopia). The latter finding comes from a procedure wherein we first estimate how factor output elasticities change with fast Internet, controlling for a possible simultaneous change in firm-level productivity (see De Loecker 2011) to uncover the technology's (positional) skill bias in Ethiopia. In the last step of the procedure, we impose additional structure to estimate how firm level productivity responds, and find a significant increase. We also use World Bank Enterprise Survey data to show that firms in Ghana, Kenya, Mauritania, Nigeria, Senegal, and Tanzania appear to export more, communicate with clients online more, and train employees more after they get access to fast Internet.

In sum, the evidence we present indicates that greater and cheaper access to information and communication due to availability of fast Internet increases employment rates in Africa, and that in at least some countries, this happens in part due to the technology's impact on firm entry, productivity, and exports. In the final part of the paper, we show that average incomes rise in the areas that see changes in employment when fast Internet arrives.

This paper contributes to the literatures on the relationship between globalization and jobs, poverty, and inequality; structural change; and constraints on firm growth in developing countries. The "new" features of international trade uncovered in the recent body of work on globalization (see Feenstra and Hanson 2003; Goldberg and Pavcnik 2007; Harrison, McLaren, and McMillan 2011; Goldberg 2015 for overviews) are important in part because they alter traditional models' prediction that locally relatively abundant factors necessarily gain the most from global integration.[8] A parallel literature convincingly demonstrates SBTC's role in slowing wage

[7] Interestingly, Atasoy (2013) finds a relatively large correlation between Internet access and employment also in the United States. Specifically, that a county gaining access to broadband services is associated with a 1.8 percentage points higher employment rate, and that the correlation is bigger in rural and isolated areas, among college-educated workers, and in industries and occupations that more-heavily utilize college-educated workers. Acemoglu and Autor (2011) and Michaels, Natraj, and Van Reenen (2014) find that, if three skill levels are considered, ICT technologies substitute most for middle-skill workers in rich countries. Of course, the types of positions that exist within a given skill category in Africa may differ from those in rich countries.

[8] Most existing studies find that trade liberalization tends to increase productivity in developing countries (Goldberg and Pavcnik 2007), with more varied effects on poverty (Topalova 2010; Winters, McCulloch, and McKay 2004) and employment rates (see e.g., Currie and Harrison 1997, Revenga 1997, Harrison and Revenga 1998, Márquez and Páges-Serra 1997, Levinsohn 1999, Moreira and Najberg 2000). Currie and Harrison (1997) is an exception in that they study (trade reform in) Africa (Morocco). Fajgelbaum and Khandelwal (2016) show that trade benefits the poor through another channel, i.e., because their consumption is relatively concentrated in traded goods.

growth and rising unemployment among less-educated workers in rich countries.[9] However, to our knowledge there is no direct, existing evidence on the causal relationship between employment rates, inequality, and incomes in developing countries and the ICT technologies that were shown to adversely affect the relative labor market outcomes of low-skill workers in rich countries (Goldberg and Pavcnik 2007).[10]

To date, research on the factor bias of new technologies in developing countries has largely focused on how technology-driven improvements in agricultural productivity affect the movement of labor in and out of agriculture (see Syrquin 1988 and Foster and Rosenzweig 2008 for overviews, and Bustos, Caprettini, and Ponticelli 2016 for a prominent recent example). Such movement is a form of structural change (Clark 1940; Lewis 1955; Banerjee and Newman 1993; Baumol 2013; Herrendorf, Rogerson, and Ákos Valentinyi 2014), i.e., a persistent change in the relative size of different sectors and occupations. Beyond the role of agricultural productivity and openness to trade, the drivers of structural change are not well understood.[11]

The literature on firms in developing countries has made considerable progress in the last decade and a half. The benefits of importing, exporting, and winning government contracts suggest that the size of the input and output markets that can be accessed is important even conditional on a firm's initial productivity (see e.g., Frías, Kaplan, and Verhoogen 2009; Goldberg et al. 2010a, b; Amiti and Davis 2012; Brambilla, Lederman, and Porto 2012; Atkin, Khandelwal, and Osman 2017; Ferraz, Finan, and Szerman 2015). Greater demand from richer consumers

[9] The relative demand for college graduates increased from the late 1980s onwards with take-up of computers in Europe and the United States (Krueger 1993; Berman, Bound, and Griliches 1994; DiNardo and Pischke 1997; Autor, Katz, and Krueger 1998; Machin and Van Reenen 1998; Autor, Levy, and Murnane 2003; Beaudry and Green 2003, 2005; Beaudry, Doms, and Lewis 2010; Acemoglu and Autor 2011; Goos, Manning, and Salomons 2014; Katz and Margo 2014; Michaels, Natraj, and Van Reenen 2014). The explanation lies not only in "direct" factor complementaries, but also in associated worker sorting and organizational change (Bartel and Sicherman 1999; Caroli and Van Reenen 2001; Bresnahan, Brynjolfsson, and Hitt 2002; Crespi, Criscuolo, and Haskel 2007; Bloom, Sadun, and Van Reenen 2012). Akerman, Gaarder, and Mogstad (2015) document an increase in the relative wages and productivity of high-skill workers when broadband Internet became available in Norway. More generally, SBTC studies that focus on advanced Internet technology in rich countries find positive correlations with local wage levels (see e.g., Czernich. et al. 2011, OECD 2013), and mixed results for the relative wage effects in richer versus poorer US counties (Forman, Goldfarb, and Greenstein 2012; Champion, Kosec, and Stanton 2012; Atasoy 2013). De Stefano, Kneller, and Timmis (2014) find no significant effect of broadband Internet on the performance of British firms.

[10] There are important existing studies of mobile phones, mobile money, and TV in poor countries that focus on price variation across space, risk sharing, and cultural change as outcomes (Jensen 2007; Jensen and Oster 2009; Aker 2010; La Ferrara, Chong,and Duryea 2012; Jack and Suri 2014). Jensen's (2007) innovative study also shows that fishermen's profits increased and consumer prices decreased when mobile phones helped eliminate price dispersion across markets in Kerala. There is also important indirect evidence on SBTC in developing countries from studies that use trade liberalization episodes or exchange rate variation that simultaneously affect trade and technological change for identification, including Harrison and Hanson (1999); Acemoglu (2003); Attanasio, Goldberg, and Pavcnik (2004); Aghion et al. (2005); Amiti and Cameron (2012); Frazer (2013); Raveh and Reshef (2016). Another indirect form of evidence that has been taken to suggest that SBTC has occurred in Latin America and India in recent decades is that the share of skilled workers has increased in most industries there (Goldberg and Pavcnik 2007). Goldberg and Pavcnik (2007) note that the skill premium increased around the same time as trade reform occurred in several Latin American countries and India, but that inequality *decreased* in several Southeast Asian countries and China when they opened up their markets (see also Wood 1999, Wei and Wu 2002). Interestingly, while income inequality has increased in many African countries in recent decades, the picture for Africa as a whole is less clear than for Asia and Latin America (Harrison, McLaren, and McMillan 2011; Dabla-Norris et al. 2015).

[11] Recent work on structural change has emphasized the importance of the manufacturing sector (Gollin, Parente, and Rogerson 2002; Lagakos and Waugh 2013; Gollin, Lagakos, and Waugh 2014; Rodrik Rodrik 2016), improvements in trends in structural change in Africa in the 2000s (McMillan and Harttgen 2014; McMillan, Rodrik, and Verduzco-Gallo 2014), and how trade liberalization can shift workers across sectors and across firms within sectors (see e.g., Attanasio, Goldberg, and Pavcnik 2004; Davis and Harrigan 2011; Young 2014).

abroad has in turn been shown to enable firms to learn and to produce higher quality products that may require more skilled workers (Verhoogen 2008; Frías, Kaplan, and Verhoogen, 2009; Atkin, Khandelwal, and Osman 2017; Hansman et al. 2017). Existing evidence also indicates that firms' financial performance is enhanced by improved coordination with suppliers, access to credit, and good management (Bloom and Van Reenen 2007; de Mel, McKenzie, and Woodruff 2008; Bloom et al. 2013; Casaburi et al. 2013; Macchiavello and Miquel-Florensa 2015). However, we know little about what drives job creation, productivity, and exporting among firms in developing countries. This is especially true when the focus is on the poorest countries and/or the role of specific technologies or inputs. The existing literature reviewed here, the role of ICT in the resurgence of US productivity growth (Draca, Sadun, and Van Reenen 2007; Oliner, Sichel, and Stiroh 2007; Jorgenson, Ho, and Stiroh 2008; Syverson 2011), and a considerable body of important correlational evidence from developing countries all underscore the promise of fast Internet.[12]

We make three main contributions to the literature. First, we use quasi-random variation in access to ICT technology to provide direct evidence on its impact on employment rates, job inequality, and incomes in 12 developing countries. These findings are important because they suggest that the factor bias of modern technologies differs in Africa. This implies that the primary explanation for rising inequality in poor countries may not be SBTC.

Second, we provide evidence on the relationship between structural change and ICT technology. This represents a first step towards understanding what drives structural change beyond the role of agricultural productivity and openness to trade. Our results qualify negative views of (other manifestations of) globalization in that fast Internet appears to increase both the share of skilled jobs and average incomes in Africa, and—at least in the Ethiopian context—productivity and employment in manufacturing.

Finally, we demonstrate how fast Internet affects employment, productivity, and exporting in African firms, expanding the body of evidence on why firms tend to grow slowly, and ways to stimulate job growth, in poor countries. Our findings on fast Internet and exports represent evidence of an interaction between technological change and trade that differs from trade-induced SBTC as analyzed by the existing literature.

The rest of the paper is organized as follows. In Section I we lay out the background on Internet and jobs in Africa and discuss examples of job creation often attributed to the submarine cables. In Section II we present our data, and in Section III the empirical strategy. The paper's main results are in Section IV, and in Section V we analyze *how* fast Internet affects employment in Africa. Section VI explores the ultimate impact on incomes. Section VII concludes.

[12] Qiang and Rossotto (2009) find that, across developing countries, a 10 percent increase in broadband penetration is associated with a 1.38 percentage point higher GDP per capita growth rate. Clarke and Wallsten (2006) find that a 1 percentage point increase in Internet users is associated with 3.8 percentage points higher exports from low-income to high-income countries. Paunov and Rollo (2015) find that use of the Internet correlates positively with firm performance in a range of poor countries. Commander, Harrison, and Menezes-Filho (2011), using more detailed data on Brazilian and Indian manufacturing firms and more extensive controls, find the same for ICT technologies. Their novel results point to much higher rates of return to investment in ICT in Brazil and India than in developed countries.

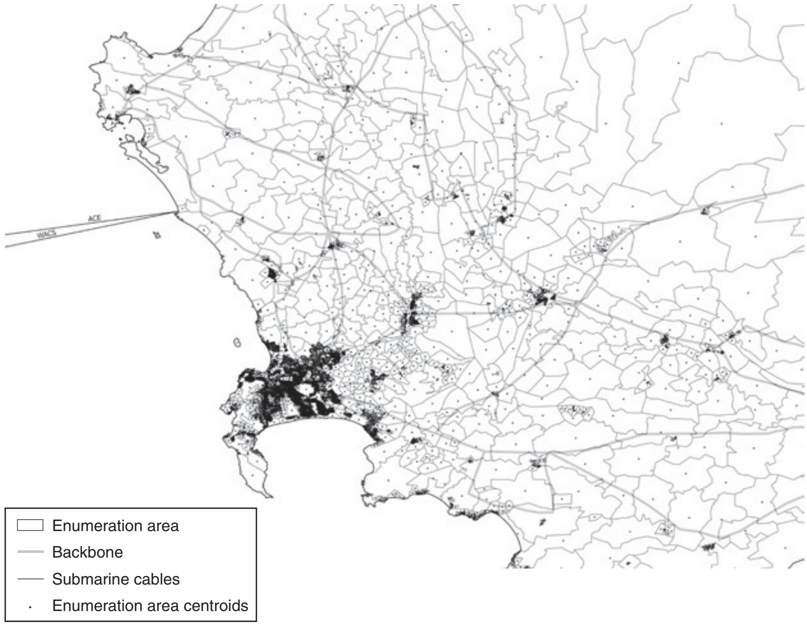

Figure 1. The Terrestrial Backbone Network, Enumeration Areas Used for Location Fixed Effects, and Sampling Clusters (Southwestern South Africa and SA-QLFS Dataset as Example)

Notes: This figure shows submarine Internet cables arriving to Yzerfontein, just north of Cape Town in South Africa, the country's terrestrial backbone network, and centroids of the SA-QLFS enumeration areas. Enumeration areas are used for location fixed effects.

I. Background

A. *Internet Infrastructure and Use in Africa*

In 2000, Africa as a whole had less international Internet bandwidth than the country of Luxembourg (ITU 2000). By 2013, 13 percent of all Africans used the Internet, compared to 36 percent globally (Internet Society 2013), and more than half of urban African adults owned Internet-capable devices (McKinsey Global Institute 2013). The forms of Internet infrastructure that reach users in Africa—the "last mile"—are fiber cables, copper cables, wireless transmission using cell towers, and satellites (de M. Cordeiro et al. 2003, Gallaugher 2012). Prior to the last mile, Internet traffic travels through a national "backbone" of bigger (typically fiber) cables, as depicted in Figure 1 for South Africa. The backbone was built by a national telecom in almost all countries, sometimes with "branches" added by private telecoms. Since Internet traffic was initially transmitted through telephone cables, the majority of the backbone network cables date back many decades (ITU 2013).

In the 2000s, submarine Internet cables from Europe were built by consortia made up of private investors, African governments, and/or multilateral organizations

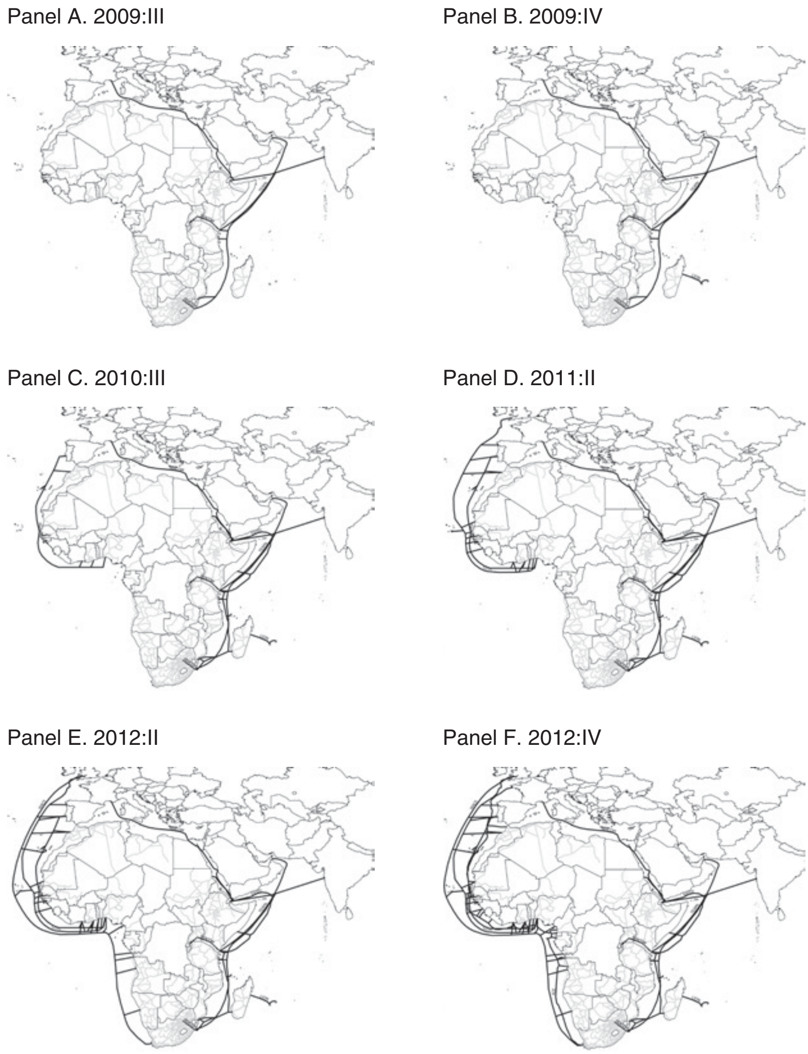

Figure 2. Submarine Internet Cable Arrival in Africa

Notes: This figure shows the arrival of submarine Internet cables in Africa over time. The first two cables during our analysis period arrived in 2009:III and the last in 2012:IV. The submarine cables are Seacom and Teams (2009:III), Lion (2009:IV), Eassy and MainOne (2010:III), Glo1 (2011:II), WACS (2012:II), and ACE (2012:IV).

(OECD 2014).[13] The submarine cables were brought to shore at landing points along the coast, typically one in each country passed by the cable. These were usually located just outside of a big city that was connected to the national backbone. Figure 2 shows the ten submarine cables that arrived in Africa during 2006–2014, as reported by Mahlknecht (2014).

[13] One of the ten submarine cables that arrived in Africa during our data period connected the continent with both Europe and India, and another one with the UAE. We refer to the connection point of the submarine cables outside of Africa as "Europe" for simplicity.

Once plugged in, the submarine cables brought much faster speed and traffic capacities on Internet traffic to and from other continents to locations in Africa connected to the terrestrial network.[14] On a cable network, the technologically feasible increase in speeds and traffic post-submarine cable plug-in decays with cable length to the landing point to a negligible extent. In general, technological bottlenecks therefore arise at the backbone-level primarily where networks owned by different owners connect to each other.

In such cases, the Internet service providers (ISPs) operating on network A will transmit content to (physically connected) network B directly only if the two networks are collaborating, for example through "peering" (ITU 2013). If not, the fees that African networks charge each other for the exchange of traffic ("transit") are such that content stored on network A would likely be sent via other continents to users on network B ("tromboning"). This partly explains the submarine cables' predicted effect on experienced speed and capacity, but a more important contributor is that "in Africa very little Internet content is sourced locally, with the vast majority sourced internationally—including local content that is hosted overseas" (Kende and Rose 2015, p. 15).[15] For example, Chavula et al. (2014) found that on average 75 percent of the traffic originating in Africa that is destined for African universities traverse links outside the continent, and Kende and Rose (2015) report that all of the top 14 commercial websites in Rwanda are hosted in Europe or the United States.

The need for African Internet traffic to travel overseas is important for this paper. In combination with each country being covered by a single backbone network, the lack of spillovers from one coastal country's submarine connection to neighboring countries means that each country has a specific *treatment date*—the date when the first cable has arrived at the country's landing point and is plugged in.[16]

In Table 1 we show the mean and standard deviation of Internet speeds and use of the Internet across locations in Africa before the submarine cables arrived. The average (measured) speed was 430 kbps, with a standard deviation of 419 kbps.[17] These relatively high numbers partly reflect the fact that our speed data measure nonmobile connections. (In Section II we describe the data in detail; some limitations of the speed measure are discussed in Subsection IVA). The proportion of individuals in the countries covered by one of our household survey datasets who used the Internet daily and weekly was 10 and 20 percent on average, with standard deviations of 30 and 40 percent.

[14] Being reached by submarine Internet cables from Europe implies a faster connection also to North America and other continents because of the extensive Internet infrastructure that connects Europe with other continents.

[15] The main reason is cost: "one content developer reported spending US$49.99 per year for up to 150GB capacity overseas, compared to a Rwandan offer of over US$900 for 50GB capacity" (Kende and Rose 2015, p. 3). Africa pays over US$600 million a year for within-Africa traffic exchange that is carried outside the continent (Internet Society 2013).

[16] We exclude landlocked countries from our analysis because the extent to which they get treated (through coastal neighbors) is unclear.

[17] These numbers exclude the four biggest cities in each country (see Table 2).

TABLE 1—INTERNET SPEED AND USE, EMPLOYMENT OUTCOMES, FIRM ENTRY, AND INCOMES BEFORE SUBMARINE CABLE ARRIVAL

	Connected [Standard deviation]	Unconnected [Standard deviation]	All [Standard deviation]	Raw baseline difference [*t*-statistic]
Internet outcomes: location and individual level				
Internet speed, kbps (location, from Akamai)	453.64	423.47	429.60	30.17
	[319.92]	[443.87]	[419.31]	[0.27]
Daily Internet use (individual, from Afrobarometer)	0.08	0.11	0.10	−0.03
	[0.27]	[0.32]	[0.30]	[−2.42]
Weekly Internet use (individual, from Afrobarometer)	0.16	0.21	0.20	−0.05
	[0.37]	[0.41]	[0.40]	[−2.61]
Employment outcomes: individual level				
Benin, D.R. Congo, Ghana, Kenya, Namibia, Nigeria, Togo, Tanzania (DHS)				
Employment	0.67	0.68	0.68	−0.01
	[0.47]	[0.47]	[0.47]	[−1.46]
Skilled	0.57	0.58	0.58	−0.01
	[0.49]	[0.49]	[0.49]	[−1.05]
Benin, Ghana, Kenya, Madagascar, Mozambique, Nigeria, Senegal, South Africa, Tanzania (Afrobarometer)				
Employment	0.56	0.59	0.58	−0.03
	[0.50]	[0.49]	[0.49]	[−1.57]
South Africa (QLFS)				
Employment	0.77	0.71	0.72	0.06
	[0.42]	[0.45]	[0.45]	[7.99]
Skilled	0.55	0.49	0.50	0.07
	[0.50]	[0.50]	[0.50]	[8.13]
Hours worked	45.26	45.38	45.36	−0.11
	[14.20]	[15.03]	[14.92]	[−0.41]
Wants to work more	0.62	0.66	0.66	−0.04
	[0.48]	[0.47]	[0.47]	[−5.82]
Formal employment	0.54	0.47	0.48	0.07
	[0.50]	[0.50]	[0.50]	[7.83]
Informal employment	0.12	0.12	0.12	0.00
	[0.33]	[0.33]	[0.33]	[−0.67]
Employment outcomes: firm level				
Number employees (from Ethiopia LMMIS)	73.90	80.83	75.47	−6.93
	[146.40]	[300.12]	[192.31]	[−0.35]
Number skilled positions (from Ethiopia LMMIS)	24.30	23.85	24.20	0.45
	[73.21]	[132.75]	[90.19]	[0.05]
Net firm entry: location level				
Net zip-code firm entry per quarter (from South Africa CIPC)	3.46	3.31	3.33	0.15
	[5.01]	[5.48]	[5.42]	[1.58]
Average local incomes: location level				
Light density at night (from NOAA)	3.62	1.41	1.54	2.21
	[10.21]	[6.15]	[6.47]	[8.89]

Notes: All measures displayed are from the period before submarine cable arrival. Internet speed data come from Akamai. They provided us with quarterly data on average connection speeds for ∼900 African locations during the 2007–2013 period. These locations are shown in online Appendix Figure A1. Akamai averages the speeds recorded for residential users, educational institutions, government offices, and firms in a given location × quarter, excluding those who connect via mobile networks. We restrict to location × quarters for which the speed measure is based on more than ten unique IP addresses. (We also exclude the 4 biggest cities in each country from the speed data sample in this table; see Table 2). Internet use rates come from Afrobarometer (survey countries and years are listed in online Appendix Table A1). We restrict the (individual level) Afrobarometer sample to observations near (< 20 km) Akamai locations for comparability (see Table 2). Employment rates are from Demographic Health Surveys (DHS), Afrobarometer, and South African Quarterly Labor Force Surveys (QLFS). Occupational skill levels in DHS and QLFS are defined according to ILO ISCO standards. Firm data are from the Ethiopia Large and Medium Scale Manufacturing Industries Survey (LMMIS). In LMMIS, skilled positions are defined as those where earnings are more than 800 Birr per year (roughly the sample salary median). Net firm entry comes from South Africa CIPC and light density at night (a proxy for average local incomes) comes from NOAA. Individuals (locations) are considered connected if they are closer than 0.5 km to the backbone network. Standard deviations are shown in square brackets in columns 1–3, and *t*-statistics are shown in square brackets in column 4.

B. *Jobs and Firms in Africa*

The 2006–2014 period we focus on was a period of high economic growth in many African countries.[18] Given their diversity, we do not attempt to describe an average labor market among the 12 countries in our sample here. Instead, the second panel of Table 1 displays, for the groups of countries covered by our respective datasets, and focusing again on the period before fast Internet became available, the proportion of individuals that have a job, and the proportion that have a job in a skilled occupation. In Benin, D.R. Congo, Ghana, Kenya, Namibia, Nigeria, Togo, and Tanzania (the DHS sample), the employment rate is on average 68 percent, with a standard deviation of 47 percent. In Benin, Ghana, Kenya, Madagascar, Mozambique, Nigeria, Senegal, Tanzania, and South Africa (the Afrobarometer sample), the employment rate is on average 58 percent, with a standard deviation of 49 percent. In South Africa, the employment rate is 72 percent, with a standard deviation of 45 percent.

In the first group of countries, 58 percent have a job that belongs to a skilled occupation as defined by the International Labour Organization (ILO), with a standard deviation of 49 percent.[19] We also observe the type of occupation to which an individual's job belongs (and several additional employment related outcomes) in South Africa. There, 50 percent have a skilled job (SD $=$ 50 percent), average hours worked per week among the employed are 45 (SD $=$ 15 hours), 66 percent "want to work more" (SD $=$ 47 percent), 48 percent have a formal job (SD $=$ 50 percent), and 12 percent have an informal job (SD $=$ 33 percent).

In the third panel of Table 1, we show that large- and medium-sized Ethiopian manufacturing firms have 75 employees on average, with a standard deviation of 192. The number of skilled positions per firm, as proxied by high salary positions, is 24.

We return to the comparison between eventually treated and untreated locations in Section III.

C. *Examples of New Job Creation after the Arrival of Fast Internet*

Many media articles and case studies illustrate new, and new types, of jobs in Africa being created after the arrival of fast Internet. Scruggs (2015) reports that "In 2009, a submarine fiber-optic cable landed in Mombassa. […] Six years later, Nairobi is bursting with technology startups like Shop Soko, a sort of Etsy for Africa that allows shopkeepers to sell handmade goods to consumers worldwide. The Kenyan capital has also emerged as [a] base for high-tech heavyweights such as Google, IBM and Intel. From 2002 to 2010, the value of Kenya's tech exports rose from US$16 million to US$360 million." Nairobi's iHub incubator had helped develop more than 150 new businesses by 2013 (McKinsey Global Institute 2013). Similarly, Harris (2012) reports that "With the landing of new submarine telecom cables off South Africa's coastline starting in 2009, bandwidth prices began to tumble, removing one of the most significant barriers to the global competitiveness

[18] However, some countries in our sample, especially South Africa, were badly affected by the 2008 global financial crisis.

[19] ILO's definition of skilled occupations is fairly broad; in Section IVE we consider each of the underlying subcategories.

of the country's IT industry. That was a catalyst for the explosion of Cape Town's tech scene […] [and] stature as a business process outsourcing [BPO] and offshoring hub." In 2013 there were more than 54,000 jobs in South Africa's new BPO sector, and Morocco's was at similar scale (McKinsey Global Institute 2013). Growth in the technology sector also has add-on benefits in other sectors, e.g., construction.[20]

Nigeria is one of the African countries where "eCommerce" has taken off, driven in part by major online retailers that also operate e.g. in Egypt, Ivory Coast, Kenya, and Morocco (Rice 2013). Online purchases in Nigeria stood at more than US$1 billion in 2014, tripling in 3 years (Atuanya and Augie 2013). Adepetun (2014) of the online news site AllAfrica.com, argues, based on interviews with officials and industry executives, that Nigeria's ICT sector from 2004 to 2014 created 100,000 direct jobs, and 1.1 million jobs indirectly, and that eCommerce and ICT's success in Nigeria is due in part to the arrival of the submarine cables.

Kenya, Nigeria, and South Africa all now have a manufacturing sector producing Internet-capable devices for the African market, such as low-cost cell phones and computers (McKinsey Global Institute 2013).

There are also signs that the arrival of fast Internet helped make supply chain coordination easier. Mozambican moWoza and similar start-ups elsewhere use apps and websites to deploy drivers to deliver parcels from wholesalers to traders, and doing the bureaucracy required to import and export in Ghana online has decreased delays considerably (McKinsey Global Institute 2013). Such supply chain improvement is believed to enhance productivity in agribusiness and manufacturing. For example, the adoption of cloud-based supply-chain management solutions by the Kenyan Tea Development Agency connected around 60 tea factories with the farmers that supply them. This reportedly reduced delays at collection points and fraud, and increased tea factories' productivity and farmers' incomes (Business Daily 2009, GIZ 2014).

Technology start-ups, BPO, eCommerce, new forms of manufacturing, and innovative supply-chain management companies and regulatory agencies that make doing business easier for factories and farmers are examples of the ways in which fast Internet may enable greater job creation. But the technology may also eliminate jobs in some occupations—or conceivably even on average—for example due to automation or increased exposure to Asian competition. In the next section we present the data that we use to investigate the causal impact of fast Internet on employment in Africa.

II. Data

Our outcome data come from the following sources:

Afrobarometer: surveys are nationally representative repeated cross-sections, conducted every two-three years in many African countries. The order in which locations are surveyed is randomly determined. We geo-code the location based on information provided on the respondent's residence. Men and women of voting age are interviewed. The survey asks socioeconomic questions. We use Afrobarometer data from coastal countries that had survey rounds both before and after sub-

[20] Scruggs (2015): "In Nairobi's Kilimani area, where the tech scene is centered, ten-story office buildings are shooting up."

marine cable arrival in the relevant country: Benin, Ghana, Kenya, Madagascar, Mozambique, Nigeria, Senegal, Tanzania, and South Africa.

From Afrobarometer we construct an outcome variable for the individual being employed.[21] We also use variables on educational attainment and Internet use.

Demographic and Health Surveys (**DHS**): are nationally representative repeated cross-sections. The order in which sampling clusters are surveyed is randomly determined. GPS coordinates for sampling clusters are recorded. Women and men between 15 and 49 years old are interviewed.[22] The survey asks questions about labor market participation, health, and demographic background. We use DHS data from coastal countries that had survey rounds both before and after submarine cable arrival in the relevant country: Benin, D.R. Congo, Ghana, Kenya, Namibia, Nigeria, Togo, and Tanzania.

From DHS we construct outcome variables for the individual being employed,[23] and for being employed in a specific type of occupation. We also use educational attainment variables.

The South Africa Quarterly Labor Force Survey (**QLFS**): is a nationally representative repeated cross-section. Unlike in Afrobarometer and DHS, QLFS surveys are carried out every quarter. GPS coordinates for enumeration areas are recorded. The current version of the survey began in 2008.[24]

From QLFS we construct outcome variables for the individual being employed,[25] and for being employed in a specific type of occupation. We also use educational attainment and other employment-related variables.

South African companies are required to register with the *Companies and Intellectual Property Commission* (*CIPC*) *Firm Registry*. We use the resulting zip code $\times$ date level panel registry, which captures entry and exit of formal firms. CIPC provided us with data from 2007:I to 2014:IV. We code up each firm's sector when its name contains sufficient information to do so.[26]

The Ethiopia Large and Medium Scale Manufacturing Industries Survey (**LMMIS**): is an annual survey of all Ethiopian manufacturing establishments that engage ten or more persons and use power-driven machines. We use the 2006 to 2013 rounds. The survey collects information on employees, inputs, production, sales, and assets, and is used to construct the country's national accounts.

[21] The question is "Do you have a job that pays a cash income?"

[22] DHS surveys both women and men, but its primary focus is on women (and children) and fewer men are surveyed. About 30 percent of the DHS sample we use is male. Note also that, for two of the countries in our DHS sample (Tanzania and Togo), the pretreatment survey round we use was conducted in the late 1990s, rather than in the years preceding the arrival of submarine cables as for the other countries (these two countries did not have a survey round in the years preceding the arrival of submarine cables). Our results are very similar if these two countries are excluded. In Afrobarometer, Tanzania was surveyed in 2008.

[23] The question is "Aside from your own housework, have you done any work in the last seven days?"

[24] From 2010:III onwards, the QLFS changed the way observations are linked to enumeration areas and locations. We thus restrict attention to the period prior to then.

[25] The question is "In the last week, did you work for a wage, salary, commission or any payment in kind (including paid domestic work), even if it was only one hour?"

[26] The procedure is described in the online Appendix. We were able to assign a sector to 67 percent of the firms based on their names.

From LMMIS we construct an outcome variable for the number of employees per firm. As proxies for skilled and unskilled positions, we use high-salary and low-salary positions.[27] When estimating production functions, we also use measures of output (value added), capital (total book value), and intermediate inputs.

The World Bank Enterprise Survey (**WBES**): is a nationally representative sample of formal firms from all sectors with five or more employees. The survey asks about the business environment, operations, output, and input use. We use WBES data from coastal countries that had survey rounds both before and after submarine cable arrival in the relevant country: Ghana, Kenya, Mauritania, Nigeria, Senegal, and Tanzania. The surveys for these countries were carried out in 2006, 2007, 2013, and 2014.

From WBES we use an outcome variable for the number of employees per firm. We also use measures of sales per unit of labor costs; direct exports; communication with clients through a website and email; and whether the firm provides training to its employees.

We also use Internet infrastructure and speed data. We use *Mahlknecht's map of submarine cables* to measure landing points and times (Mahlknecht 2014), and www.africabandwidthmaps.com and *AfTerFibre's* (AfTerFibre 2014) maps of terrestrial backbone networks to measure locations' connectivity.[28]

Our data on Internet speeds come from the content delivery network *Akamai Technologies, Inc.*, which owns servers worldwide and serves 15–30 percent of all Internet traffic. Akamai averages the speeds recorded for residential users, educational institutions, government offices, and firms in a given location $\times$ quarter, excluding those who connect via mobile networks. (We discuss a limitation of this measurement method in Subsection IVA). Akamai provided us with quarterly data on average connection speeds for $\sim$900 African locations during the 2007–2014 period. These locations are shown in Appendix Figure A1.

III. Empirical Strategy

We analyze the relationship between employment patterns in a given location and time period on the one hand and whether or not the location is connected to submarine Internet cables from Europe via the terrestrial backbone network on the other. We run

$$(1)\quad y_{ij(i)c(i)t} = \alpha + \beta SubmarineCables_{c(i)t} \times Connected_i + \delta_{j(i)} \times Connected_i + \gamma_{c(i)t} + \epsilon_{ij(i)c(i)t},$$

where $y_{ij(i)c(i)t}$ is an outcome for individual i in grid-cell $j(i)$, country $c(i)$, and time period t. $SubmarineCables_{c(i)t}$ is a dummy variable equal to one if the backbone network in country $c(i)$ has been connected to at least one submarine cable at t, and

[27] LMMIS does not contain information on occupational categories. Skilled (high-salary)/unskilled (low-salary) positions are defined as those where salary is higher/lower than 800 Birr per year, approximately the sample salary median.

[28] We consider Ethiopia "treated" because it is well documented that its backbone became internationally connected via the submarine cable landing in Djibouti, which was planned and built to also cover Ethiopia (Giorgis 2010, Oxford Business Group 2016).

$Connected_i$ is a dummy variable equal to one if individual i's location is connected to the backbone network. In some parts of our analysis, i represents a firm or geographical location rather than an individual.

The interaction between 0.1×0.1 degree ($\sim 10 \times 10$ km) grid-cell fixed effects— $\delta_{j(i)}$—and the $Connected_i$ indicator controls for any time-invariant differences in employment outcomes that may be correlated with access to fast Internet. These can be included because all our datasets are panels at "$\delta_{j(i)} \times Connected_i$ level." Some of our datasets are sufficiently balanced across time at the lowest geographical level at which i's location is reported (e.g., an enumeration area) that $\delta_{j(i)} \times Connected_i$ can be replaced with location fixed effects defined at that level.[29]

Country-specific time period (quarter or year) fixed effects, $\gamma_{c(i)t}$, control for any within-country-location-invariant changes in employment outcomes that may be correlated with getting access to fast Internet.[30] The effect of fast Internet is thus identified off of the comparison between the change in outcomes for locations that gain access to fast Internet in a given quarter or year and the change in outcomes for other locations in the same country that do not gain access at the same point in time.

We cluster the standard errors at location level; i.e., $j(i)$. Most of our outcome variables are $0/1$; the ones that are not are highly skewed. We transform these using the inverse hyperbolic sine (asinh).[31]

Since we lack information on last mile infrastructure at the local level,[32] we define i's location as *connected* if it is near infrastructure that makes availability of fast Internet possible, i.e., the country's backbone network. (We refer to a location as *treated* at t if additionally at least one submarine cable has arrived in the country at t). We use maps of Africa's backbone networks prior to the arrival of the ten submarine cables to measure such connectivity.[33] Specifically, we define as connected

[29] In such cases there is no need to include interactions between location fixed effects and $Connected_i$ because the lowest geographical level at which i's location is reported is either connected or not in these datasets. On the other hand, in DHS and Afrobarometer, information on an individual's village or neighborhood is reported, but during our data period specific villages/neighborhoods rarely appear in two different survey rounds of which one is conducted before and one after submarine cable arrival. The connected and unconnected parts of 10×10 km grid-cells, however, often do. Note that $\delta_{j(i)} \times Connected_i$ is shorthand in (1) in the sense that we interact $\delta_{j(i)}$ both with the $Connected_i$ dummy and the converse dummy for *not* being connected.

[30] Online Appendix Table A1 lists when the countries in our sample were surveyed and when they were reached by submarine cables. We use the quarter in which the survey was conducted (or to which the observation belongs) to designate a given observation as pre- versus post-submarine cable arrival in all our outcome datasets as this is the time level at which Mahlknecht's (2014) map of submarine cables reports arrival times in the various landing point cities along the coast.

[31] Our results are robust to instead clustering the standard errors at the level of administrative units and to computing standard errors using methods designed to account for spatial correlation, as discussed in Section IV. The asinh function closely parallels the natural logarithm function, but is well defined at zero (see Card and DellaVigna forthcoming).

[32] In Subsection IA we discussed how the technologically feasible increase in traffic and speeds post-submarine cable arrival decays with cable length *along* backbone networks to a negligible extent. Connectivity *is* lower further away from than close to the backbone network, but the connectivity reach beyond the backbone network depends on the last mile infrastructure in place in a given area (Commonwealth Telecommunications Organisation 2012, Banerji and Chowdhury 2013).

[33] To construct our map of the initial backbone network, we start with AfTerFibre's 2013 map (which is publicly and freely available, and for which corresponding GIS shape files are provided). We then use a map of backbone networks in Africa from www.africabandwidthmaps.com that is available (for purchase, and without shape files) both for 2009:II and 2013:II to identify the (few) backbone segments built during that period. Finally, we "remove" these new segments from the AfTerFibre map. We calculate the distance between an individual, firm or location in the sample, and the nearest point on the country's backbone network. (For QLFS, we observe the location of the ($\sim$80,000) enumeration areas individuals belong to. In QLFS we thus define the location of the individual as the GPS coordinates of the centroid of his or her enumeration area).

those locations that are less than 500 meters from the backbone network. Dividing the sample into two groups facilitates easy inspection of possible differences in pretrends in the outcomes across connected versus unconnected locations, and this approach also simplifies interpretation of the estimates. In Section IV, we show that our results are robust to varying the radius used to define connectivity, and to relaxing the binary definition of connectivity.[34]

We exclude locations that are further than 10 km from the backbone network. The identifying assumption is that locations close to and somewhat further away from the terrestrial backbone network were on parallel trends in employment outcomes prior to the arrival of submarine Internet cables in Africa and did not experience systematically different idiosyncratic shocks after the cables arrived. To illustrate the geographical variation that we exploit, Figures 1 and 3 display points where we observe individuals' location, the backbone networks, and the areas we use to define location fixed effects for two specific areas and datasets (southwestern South Africa for the QLFS dataset and southern Ghana for the Afrobarometer dataset).

Table 1 shows, in addition to the overall employment rates by groups of countries covered by our respective datasets, the breakdown by connected versus unconnected areas. Differences in employment rates are small in most countries; in the DHS and Afrobarometer countries the employment rate is respectively 1 and 3 percentage points higher in unconnected areas, while in South Africa the employment rate is 6 percentage points higher in connected areas. The rate of employment in skilled positions is 1 percentage point higher in unconnected areas in the DHS countries, and 7 percentage points higher in connected areas in South Africa. Firms on average employ 74/81 and 24/24 workers overall and in skilled positions in connected/unconnected areas in Ethiopia. Internet speeds are slightly higher in connected areas, whereas take-up rates are slightly higher in unconnected areas.[35]

In Subsection IVC we investigate possible violations of the identifying assumption of parallel trends. We show that our results are robust to varying the radius around the backbone network used to define connectivity status; to varying the size of the grid-cells used to define location fixed effects; to defining the backbone network as the intersection of cables reported by two different data sources; to excluding landing point locations; to including placebo treatments that interact $\text{SubmarineCables}_{it}$ with proximity to roads, electricity networks or 3G coverage; to controlling for location-specific linear and nonlinear trends in the outcomes; to including leads and lags of $\text{SubmarineCables}_{it}$; and to alternative ways to compute standard errors. We

[34] Given that back-haul networks (such as metropolitan loops) were generally lacking in sub-Saharan Africa during our data period, most telephone and Internet exchange points were likely located along the national backbones. Technological considerations indicate that 500 meters is a reasonable proxy for potential fast Internet reach beyond the backbone cables for copper-cable last mile technologies. (For last mile transmission via microwaves, the distance-connectivity relationship beyond the backbone is less clear-cut. We thus choose a conservative radius based on copper-cable technologies.) Our empirical strategy may underestimate the true effect of fast Internet in treated locations since locations further than 500 meters from the backbone network may also benefit from the arrival of submarine cables, even if they do so to a lesser extent. It is also possible that neighboring locations suffer (or benefit) from the greater increase in access to fast Internet in connected locations. In Section IV we vary the assumed connectivity radius and also compare locations at varying distances to the backbone network.

[35] Column 4 in Table 1 displays the raw baseline differences between connected—that is, eventually treated—and unconnected locations/firms/individuals in our samples. Recall that the generalized difference-in-differences approach used in our empirical analysis nets out any level differences.

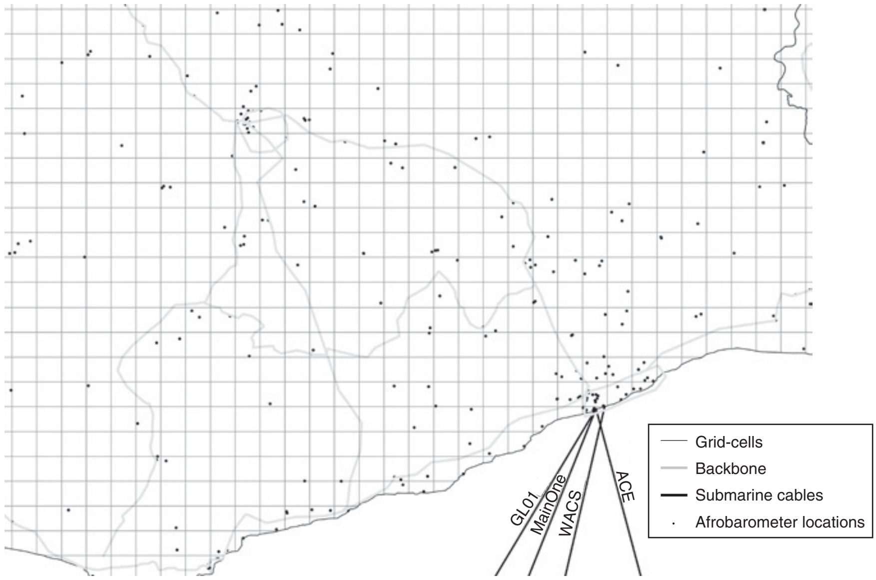

Figure 3. The Terrestrial Backbone Network, Grid-Cells Used for Location Fixed Effects, and Sampling Clusters (Southern Ghana and Afrobarometer Dataset as Example)

Notes: This figure shows submarine Internet cables arriving to the Accra region in Ghana, the country's terrestrial backbone network, and locations from the Afrobarometer. Grid-cells are 0.1 $\times$ 0.1 decimal degrees, which is roughly 10 $\times$ 10 km, and are used for location fixed effects.

also show direct evidence of parallel pretrends, and that our estimates remain significant if we use a nonparametric permutation test for inference.

Figure 2 shows the submarine cables that had arrived in different landing point locations along the coast at various times during our data period. The figure illustrates two important aspects of the identifying variation we exploit. First, submarine cables arrive at many different points in time and at different points in time in different countries. This means that we compare connected and unconnected locations across many different points in time rather than a single date. Second, the order in which different countries are reached by a given submarine cable is geographically determined. It is, thus, a priori unlikely that arrival times correlate with temporal variation across countries in differences between the economic trajectories of connected and unconnected areas.

IV. Results

A. *Submarine Cable Arrival and Internet Speed and Use*

Before analyzing how access to fast Internet affects employment in Africa, we document that the arrival of submarine cables increases both average speeds and use of the Internet. Columns 1 to 3 of Table 2 show results from running (1) with the outcome defined as the average Internet speed in a given location $\times$ quarter as measured in Akamai's data. We find that cable arrival increases measured speed in connected locations, relative to unconnected locations, by around 35 percent

Table 2—Submarine Cable Arrival and Internet Speed and Use

				Internet use			
Outcome	Internet speed (asinh)			Daily (0/1)		Weekly (0/1)	
Unit of analysis	Location			Individual			
Sample	Akamai			Afrobarometer			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
SubmarineCables × Connected	0.354 (0.137)	0.362 (0.176)	0.380 (0.191)	0.082 (0.028)	0.124 (0.050)	0.123 (0.032)	0.142 (0.053)
Observations	2,533	1,670	1,670	4,160	4,160	4,160	4,160
Mean of outcome				0.10	0.10	0.20	0.20
Country × time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number unique IPs > 10	Yes	Yes	Yes	No	No	No	No
Connected × time FE	No	No	Yes	No	Yes	No	Yes
Including biggest cities	Yes	No	No	No	No	No	No

Notes: Internet speed (kbps) and use (0/1) data come from Akamai and Afrobarometer respectively (more details in Table 1). Akamai data is for 2007–2014 and the locations are shown in online Appendix Figure A1. Afrobarometer survey countries and years are listed in online Appendix Table A1. Akamai recommends focusing on location × quarters for which the speed measure is based on more than 1,000 unique IP addresses, but this dramatically reduces the sample size. We thus restrict to number unique IP > 10. Individuals (locations) are considered connected if they are closer than 0.5 km to the backbone network. This indicator may incorrectly classify large, connected cities as unconnected when the city is considered as one geographical unit. Since the Akamai data is at location/city × quarter level (e.g., Cape Town × 2009:I), we exclude each country's four largest cities (as defined by UNData) from the Akamai sample (if they are present in the sample to begin with). As seen in column 1, the estimated effect of fast Internet is of similar magnitude and significant also in the full sample. (Note that this issue is not relevant to our other datasets, where we have much more fine-grained information on individuals and firms' location). We restrict the (individual level) Afrobarometer sample to observations near (< 20 km) Akamai locations for comparability. Time FEs are years. Location FEs are the reported location (a village/neighborhood or city) in both samples. Robust standard errors clustered at the level of location FEs in parentheses.

in the full sample, 36 percent when we leave out the biggest cities in each country, and by around 38 percent when we also control for interactions between the Connected_i indicator and the time period fixed effects. (We motivate this control in Subsection IVC).[36]

Akamai informed us that, because only a fraction of their African speed tests were "sent" to servers on other continents during our data period, the coefficients estimated in columns 1 to 3 of Table 2 are likely much smaller than the true effect of the submarine cables on speeds experienced by users.[37] This is in line with numerous media articles and existing analyses reporting large increases in speed with the arrival of submarine cables in Africa (see e.g., BBC 2009, CNN 2009,

[36] We display results excluding the biggest cities because the 500-meter connectivity radius may misclassify the biggest, connected cities in Akamai's sample as unconnected, as explained in more detail in the notes to the table.

[37] The reason is that Akamai's technology normally tests a user's speed of connection to a nearby server. In general, during our data period, the speed recorded was that to a server in another country—typically in or via Europe—only in cases where Akamai did not own a server that was located within the user's ISP's own network or directly upstream. It is, however, primarily speeds on traffic to other continents that are affected by the submarine cables, as discussed in Subsection IA. Almost all Internet traffic from Africa did indeed travel to or via other continents during our data period, as also discussed in Section IA. Despite significant efforts, we have not managed to find Internet speed data covering our data period that explicitly measure speeds between specific locations in Africa and other continents over time, and Akamai would not share a more detailed version of their data (that could allow us to separate out the speed tests that were sent to other continents) with us.

Akamai 2012), and with a large positive jump in *country-level* average broadband speeds when submarine cables arrive.[38]

In the Afrobarometer surveys, respondents are asked if they use the Internet daily or weekly. In columns 4 and 6 of Table 2, we show results from again running (1), except that the outcome variable is now a dummy for the individual reporting that she uses the Internet. We find that submarine cable arrival increases the probability that an individual uses the Internet daily in connected relative to unconnected locations by about 8 percent on average, and the probability that she uses the Internet weekly by about 12 percent. When we control for interactions between the Connected_i indicator and time period fixed effects in columns 5 and 7, the estimated effect on daily and weekly Internet use is respectively 12 and 14 percent.

There are likely two reasons why use of the Internet increases with submarine cable arrival. First, the technology becomes more useful to potential users. Second, the arrival of the submarine cables led to "drastic falls in prices for international capacity" (Kende and Rose 2015, p. 15); a cost decrease that ISPs likely partly pass on to users via lower prices. This is supported by Appendix Figure A1, where we use country-level data from the International Telecommunication Union (ITU). As seen in the figure, there is no discernible trend in broadband subscription or connection charges before submarine cable arrival, whereas we see a large decrease in charges the years after the cables arrive.

Of course, the increase in take-up by *employers* after the arrival of the submarine cables may differ from that for individuals. In Section V, where we use firm-level data to explore channels through which fast Internet may affect employment, we analyze how firms' use of websites and email responds.

We conclude that, while data limitations prevent us from pinning down the exact magnitude of the increase in experienced speeds and employment-related use of the Internet with the arrival of fast Internet in Africa, the evidence suggests that both rise considerably. This highlights that an impact on employment patterns may arise both through inframarginal users increasing and changing their use of the Internet, and through take-up by new users.

B. *Fast Internet and Employment Rates*

In Table 3 we report this paper's first main finding: the estimated effect of the arrival of fast Internet on employment rates in Africa. In the eight countries for which we have DHS data (Benin, D.R. Congo, Ghana, Kenya, Namibia, Nigeria, Togo, and Tanzania) we find a 4.6 percentage point, or 6.9 percent, increase in the probability that an individual is employed when fast Internet arrives. In the nine countries for which we have Afrobarometer data (Benin, Ghana, Kenya, Madagascar, Mozambique, Nigeria, Senegal, Tanzania, and South Africa) we find an even bigger 7.7 percentage point, or 13.2 percent, increase in the employment rate. In South

[38] This can be shown using wider-coverage-but-country-level data from the International Telecommunication Union (ITU) (a UN agency covering issues related to information and communication technologies). The regression includes pre- and post-indicators and country and year fixed effects. The outcome is fixed-broadband speeds measured in Mbit/s. As submarine cables arrive in the countries studied in this paper (with the exception of D.R. Congo, for which data is missing), country-level average broadband speeds rise by more than 100 percent.

Table 3—Fast Internet and Employment

Panel A. Employment

Outcome	Employment (0/1)		
Unit of analysis	Individual		
	DHS	Afrobarometer	SA-QLFS
	(1)	(2)	(3)
SubmarineCables × connected	0.046	0.077	0.022
	(0.014)	(0.037)	(0.008)
Observations	59,914	7,918	280,641
Mean of outcome	0.68	0.58	0.72
Country × time FE	Yes	Yes	No
Grid-cell × connected FE	Yes	Yes	No
Time FE	No	No	Yes
Location FE	No	No	Yes

Panel B. Work-related outcomes from SA-QLFS

Outcome	Hours worked (asinh)	Wants to work more (0/1)	Formal employment (0/1)	Informal employment (0/1)
Unit of analysis	Individual			
	(1)	(2)	(3)	(4)
SubmarineCables × connected	0.101	−0.022	0.017	0.004
	(0.035)	(0.008)	(0.009)	(0.005)
Observations	279,482	457,192	280,641	280,641
Mean of outcome		0.66	0.48	0.12
Time FE	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes

Notes: The DHS sample includes Benin, D.R. Congo, Ghana, Kenya, Namibia, Nigeria, Tanzania, and Togo. The Afrobarometer sample includes Benin, Ghana, Kenya, Madagascar, Mozambique, Nigeria, Senegal, South Africa, and Tanzania. The QLFS survey is from South Africa. Survey years for each DHS and Afrobarometer country are reported in online Appendix Table A1. QLFS data are 2008:I–2010:II. Grid-cells are 0.1 × 0.1 decimal degrees, which is roughly 10 × 10 km. Location FEs are enumeration areas in South Africa QLFS. Time FEs are quarters in QLFS and years in DHS and Afrobarometer. Individuals (locations) are considered connected if they are closer than 0.5 km to the backbone network. Hours worked is defined as zero for unemployed individuals. Robust standard errors clustered at the level of location FEs in parentheses.

Africa—for which we use labor force survey data—we find a 2.2 percentage point or 3.1 percent increase in employment.

Given the large magnitude of these estimates, one may wonder to what extent they reflect real additional economic activity. In panel B of Table 3, we use more detailed work-related questions available in the QLFS dataset to investigate this. In column 1, we show that access to fast Internet increases hours worked by about 10 percent on average in South Africa. This helps rule out, for example, that fast Internet simply allows individuals to spread out their work hours over time (which could affect how they answer employment questions in a survey). The increase in hours worked also helps explain why the technology reduces the probability that an individual "wants to work more" by 2.2 percent, as seen in column 2. Another possibility is that the estimates in the top panel of Table 3 reflect formalization of preexisting informal jobs rather than additional employment. This is unlikely because all the surveys we use ask about employment status in a way that should capture also informal employment (see Section II). The QLFS survey explicitly records both formal and informal employment, however. As seen in columns 3 and 4 of panel B, the estimated increase in formal employment is only slightly smaller than the estimated increase

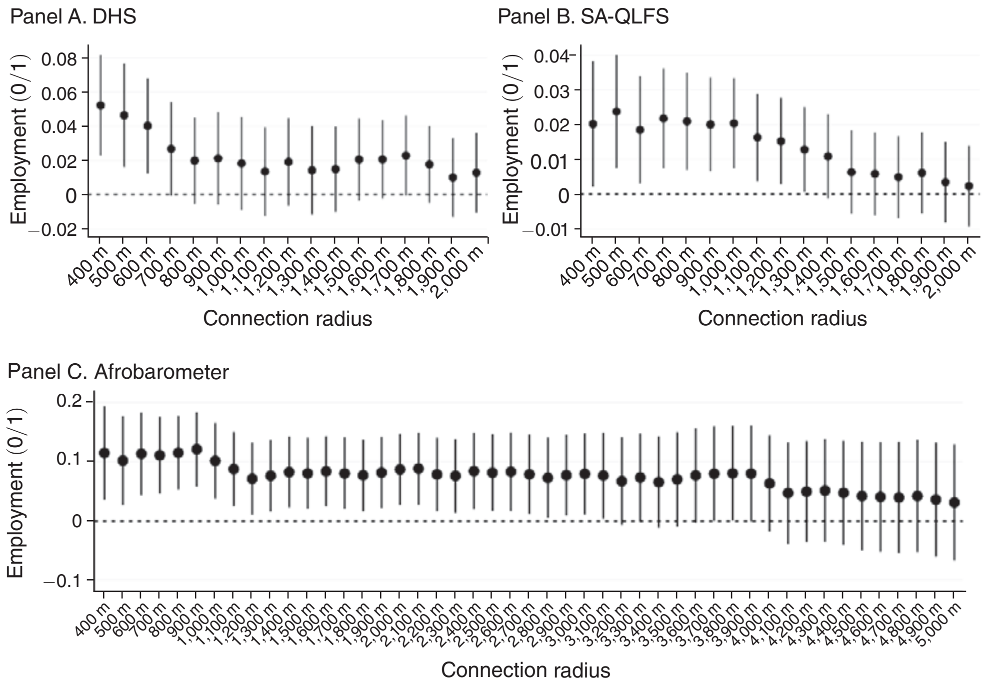

Figure 4. The Estimated Effect of Fast Internet on Employment, Varying the Assumed Connection Radius

Notes: This figure plots the coefficients from running (1) using varying connection radii. The coefficients thus come from separate regressions. We display a longer x-axis for Afrobarometer since the estimated effect becomes insignificant "further out" in that sample.

in any employment, while the estimated effect on informal employment is positive but close to zero and insignificant.

The evidence suggests that real employment in Africa increases substantially when fast Internet becomes available. In Section IVC we probe the identifying assumption underlying our causal interpretation of the estimates in depth.

C. *Robustness*

We start by confirming that the estimated effect of fast Internet is not sensitive to the radius around the backbone network used to define locations' connection status. In Figure 4 we display point estimates and confidence intervals for a wide range of radii, each used to define connectivity in a separate regression. For each of our three main outcome datasets, we display results for several radii beyond the connection radius at which the point estimate becomes insignificant. In all three datasets, the point estimate falls as the connection radius is increased, as one would expect. The decay in the estimate as we increase the assumed connection radius is steepest in the DHS sample and least steep in the Afrobarometer sample. In all three samples, the point estimate remains significant well beyond the 500-meter radius we use to define connectivity in our baseline approach. In Appendix Table A1, we show that the results are also not sensitive to the size of the grid-cells used to define location fixed effects.

In Appendix Table A2, we vary the backbone cables used to define connection status and the sample analyzed in several ways. We first use the intersection of the AfTerFibre and www.africabandwidthmaps.com maps to define connectivity.[39] When we do so, the estimated effect of fast Internet is essentially unchanged in two of our samples, as seen in panel A. In the Afrobarometer sample, the point estimates decreases somewhat and becomes insignificant, but remains large in magnitude when we use only backbone cables reported by both these sources to define connectivity.

In panel B of Appendix Table A2, we exclude from the sample all individuals located less than 20 km from a landing point. The locations that were chosen as landing points are, in addition to being on the coast, typically in or near large cities. If such locations were on a different trend in employment before the arrival of submarine cables, we may incorrectly attribute an estimated treatment effect to the arrival of fast Internet. However, the results are essentially unchanged—if anything the point estimates are slightly larger in magnitude—when we exclude near-landing point locations.[40]

In panel C of Appendix Table A2, we exclude from the sample all observations in locations that are more than 5 km from the backbone network itself. Though there are arguments for including more remote locations in the sample—they are presumably less likely to be indirectly affected by the arrival of fast Internet than unconnected locations closer to the backbone—such locations likely differ more from connected locations. The estimates in panel C make clear that our findings in Table 3 are not driven by the inclusion of more remote, less comparable locations in the analysis sample.

In Table 4 we include additional controls. In most African countries, a part of the backbone network runs parallel to other infrastructure such as roads or electricity cables (see Appendix Figure A2). If locations near such infrastructure saw faster employment growth over time, irrespective of whether they were also connected to the Internet backbone, there is a risk of misattributing employment growth to the arrival of submarine cables. We thus use maps of Africa's road and electricity network to define each location's "road-connectivity" and "electricity-connectivity" status, as we do for Internet backbone connectivity. We interact these with the arrival of submarine Internet cables—analogously to the construction of $SubmarineCables_{it} \times Connected_i$ in (1)—to construct placebo road- and electricity treatments. When these are included, the estimated effect of fast Internet is essentially unchanged and the estimated coefficients on the placebo treatments are small and insignificant, as seen in columns 1, 3, and 5 of Table 4.

In column 5 of Table 4, we also include a placebo treatment that interacts $SubmarineCables_{it}$ with an indicator for the location having 3G mobile coverage at t, similarly to the approach for roads and electricity connectivity (except that 3G coverage varies over time). This is possible when we use the QLFS sample since 3G coverage data is available for South Africa. The coefficient on the treatment variable

[39] The drawback of this approach is that we can only implement it with post-treatment (2013) backbone maps (see Section III). However, few backbone cables were finalized and turned on during our data period.

[40] This finding also implies that an increase in demand due to the building of the submarine cables themselves cannot explain the effect on overall employment rates. Locations near the landing points are presumably places where a lot of the submarine cable-driven increase in construction and related employment would have occurred.

Table 4—Fast Internet and Employment, Including Placebo "Treatments," and Controlling for Trends

Outcome	Employment (0/1)							
Unit of analysis	Individual							
Sample	DHS		Afrobarometer		SA-QLFS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SubmarineCables × connected	0.049	0.108	0.084	0.142	0.021	0.017	0.022	0.018
	(0.014)	(0.040)	(0.037)	(0.071)	(0.008)	(0.010)	(0.009)	(0.010)
SubmarineCables × connected to road network	−0.014		−0.016		0.000			
	(0.009)		(0.024)		(0.006)			
SubmarineCables × connected to electricity grid	0.006		0.004		−0.012			
	(0.009)		(0.032)		(0.017)			
SubmarineCables × connected to 3G					0.007			
					(0.006)			
SubmarineCables × connected, $t-1$							0.000	
							(0.008)	
SubmarineCables × connected, $t+1$								0.005
								(0.007)
Observations	59,914	59,914	7,900	7,900	280,641	280,641	280,641	280,641
Mean of outcome	0.68	0.68	0.58	0.58	0.72	0.72	0.72	0.72
Country × time FE	Yes	Yes	Yes	Yes	No	No	No	No
Grid-cell × connected FE	Yes	Yes	Yes	Yes	No	No	No	No
Time FE	No	No	No	No	Yes	Yes	Yes	Yes
Location FE	No	No	No	No	Yes	Yes	Yes	Yes
Linear grid-cell trend	No	No	No	No	No	Yes	No	No
Connected × time FE	No	Yes	No	Yes	No	No	No	No

Notes: The DHS sample includes Benin, D.R. Congo, Ghana, Kenya, Namibia, Nigeria, Tanzania, and Togo. The Afrobarometer sample includes Benin, Ghana, Kenya, Madagascar, Mozambique, Nigeria, Senegal, South Africa, and Tanzania. The QLFS survey is from South Africa. Survey years for each DHS and Afrobarometer country are reported in online Appendix Table A1. QLFS data are 2008:I–2010:II. Grid-cells are 0.1 × 0.1 decimal degrees, which is roughly 10 × 10 km. Location FEs are enumeration areas in South Africa QLFS. Time FEs are quarters in QLFS and years in DHS and Afrobarometer. Individuals (locations) are considered connected to the backbone, roads, and electricity if they are closer than 0.5 km to the backbone network, the road network, and the electricity grid respectively; and to 3G if the individual (location) is within 3G coverage. The GIS shapefile for African electricity grids comes from The Africa Infrastructure Country Diagnostic (AICD), that for African road networks from the Socioeconomic Data and Applications Center (SEDAC) at the Center for International Earth Science Information Network at Columbia University, and that for 3G data from Collins Bartholomew. Robust standard errors clustered at the level of location FEs in parentheses.

for access to fast Internet remains essentially unchanged; it is thus clear that Internet affects employment rates whether or not the area is covered by the 3G network.

We next control for a nonlinear trend in employment that is specific to the connected locations. Specifically, we include interactions between the Connected indicator and the time fixed effects in columns 2 and 4 of Table 4. This is possible for the multi-country DHS and Afrobarometer samples, where the arrival of fast Internet is staggered across time. Estimating the treatment effect of interest while controlling for Connected × Time FEs is unusually demanding on the data. Remarkably, the estimated coefficient on $SubmarineCables_{it} \times Connected_i$ remains large and significant in both the DHS and the Afrobarometer sample (and in fact the point estimate increases in magnitude in both samples). In column 6, we approximate this multi-country specification in the South Africa sample by including linear grid-cell specific trends. The estimated coefficient on the access-to-fast-Internet indicator is essentially unchanged.

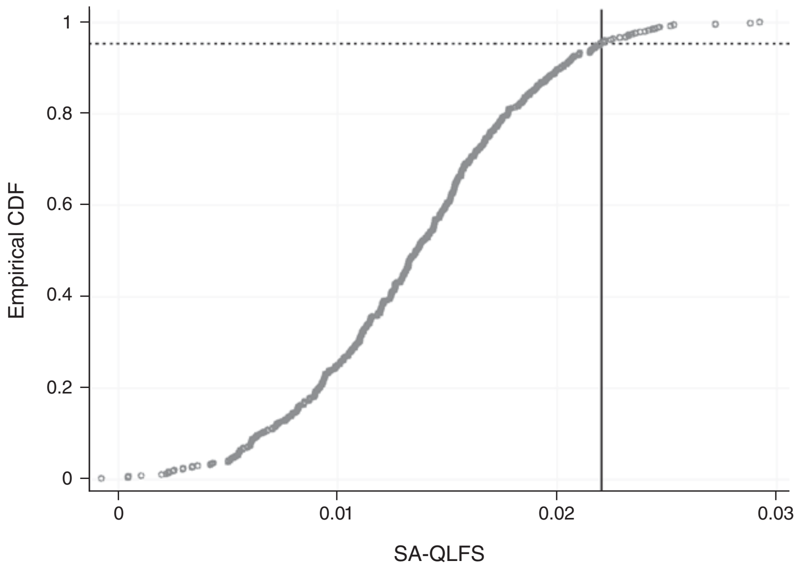

FIGURE 5. DISTRIBUTION OF PLACEBO ESTIMATES: FAST INTERNET AND EMPLOYMENT IN SOUTH AFRICA

Notes: This figure shows a nonparametric permutation test of $\beta = 0$. We sample from all possible submarine cable arrival times, assigning a randomly chosen "fake" arrival time to each location while maintaining each observation's backbone connectivity status. There are as many possible arrival times as there are quarters in the SA-QLFS sample. The figure depicts the empirical cdf of estimates resulting from permuting arrival times 500 times and running (1) on each fake dataset. The vertical line represents the true estimate; where it falls in the empirical cdf of estimates from the dataset with permuted arrival times implies its *p*-value, which is 0.046.

Finally, in columns 6 and 7, we include a lead and a lag of $\text{SubmarineCables}_{it}$. This is possible in the QLFS dataset, wherein data is collected every quarter. Perhaps somewhat surprisingly, the effect of $SubmarineCables_{it} \times Connected_i$ loads on the quarter-of-arrival treatment indicator when a lag is included. More importantly, the estimated coefficient on the lead is near zero and insignificant, supporting the identifying assumption of parallel trends.

Bertrand, Duflo, and Mullainathan (2004) point out that serial correlation can bias standard errors in difference-in-differences analysis. To address this concern, we follow Chetty, Looney, and Kroft (2009) and conduct a nonparametric permutation test of $\beta = 0$. We can do so in the QLFS dataset, where data is collected every quarter. We sample from the set of possible submarine cable arrival times, assigning a randomly chosen "fake" arrival time to each location while maintaining each observation's backbone connectivity status. Figure 5 depicts the empirical cdf of estimates resulting from permuting arrival times 500 times and running (1) on the fake datasets. The vertical line represents the true estimate; where it falls in the empirical cdf of estimates from datasets with permuted arrival times implies its *p*-value. As seen in the figure, the true estimate is near the top of the empirical cdf, with an implied *p*-value of 0.046.

Conley (1999) emphasizes that *spatial* correlation may also require corrections to standard errors and develops a method for implementing such corrections. In panel A of online Appendix Table A3, we present the estimates from Table 3 and standard errors that are calculated using Conley's method. In panel B of the same

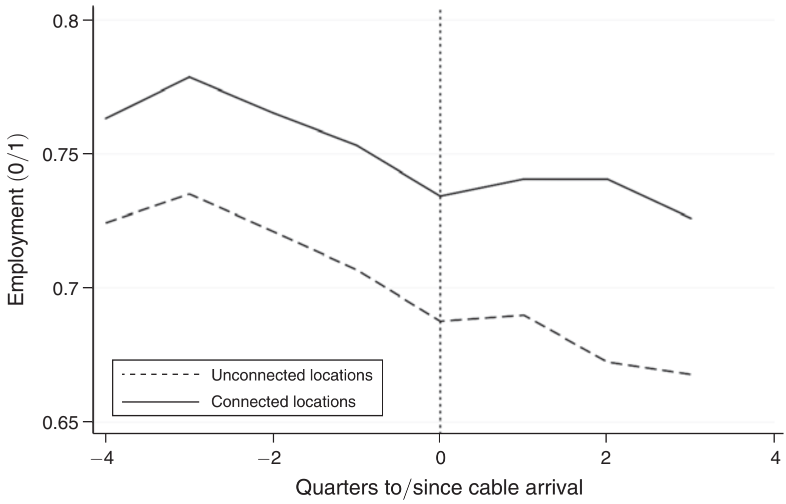

Figure 6. Employment Rate in Connected and Unconnected Locations in South Africa before and after Fast Internet Arrival

Notes: This figure plots the employment rate in connected and unconnected locations in South Africa before and after the first submarine cable during our data period arrived in the country (2009:III).

table, we cluster the standard errors by administrative unit, rather than grid-cells. In both cases, the estimated effect of fast Internet on employment rates remains statistically significant.

Finally, in Figure 6, we again take advantage of the quarterly surveying in QLFS to display the path of the employment rate in connected and unconnected areas before and after the arrival of the first submarine cable in South Africa. This allows us to inspect how the gap between the two areas evolves after fast Internet arrives and, more importantly, to check if the identifying assumption of parallel pretrends appears to hold. Indeed, while the employment rate in both areas declines between 2008 and 2011, in part due to the financial crisis that hit South Africa during that period, the shape of the graph is virtually identical for connected and unconnected areas before the submarine cable arrives in mid-2009. The gap in the employment rate between the connected and unconnected areas starts to increase soon after submarine cable arrival and widens further over time, illustrating the treatment effect estimated in Table 3.

We conclude that the estimated effect of access to fast Internet on employment rates in Africa is robust and likely represents a causal response.

D. *Fast Internet and Employment Rates across Space*

We have established that the arrival of fast Internet in Africa led to a large increase in employment rates in connected areas relative to unconnected areas. This finding would hold even if employment in unconnected areas was also affected. However, it is possible that the impact we estimate in Section IVB does not capture the total effect of fast Internet across space. We would overestimate the total effect if, for example, (existing or newly created) jobs are shifted from unconnected to

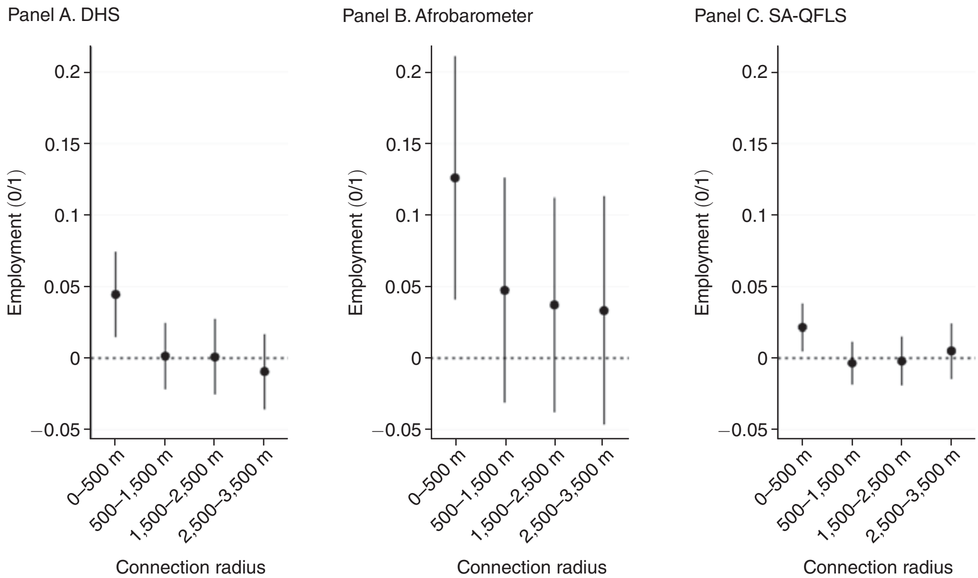

FIGURE 7. THE EFFECT OF FAST INTERNET ON EMPLOYMENT ACROSS SPACE

Notes: This figure plots the coefficients from running (1) using several connection radii. The coefficients thus come from the same regression for each sample. The first radius is the baseline specification in the paper, i.e., 0–500 meters. We then include three "bands" further away from the backbone to display effects across space. These additional treatment groups are each 1 km wide since there are fewer and fewer observations the further out from the backbone.

connected areas. We would underestimate the total effect if, for example, surveyed individuals commute to work in or migrate to connected areas. We now investigate these possibilities.

First, note that we find no effect of access to fast Internet on migration in South Africa and Tanzania (for which the required data is available).[41]

Second, recall that we confirmed in Figure 4 that the estimated positive effect on employment of fast Internet is not sensitive to the radius around the backbone network used to define connection status. This finding implies that our results are unlikely to be driven by a simultaneous decrease (or increase) in employment in areas neighboring connected locations.

Finally, we investigate more directly in Figure 7. Recall that we consider individuals and areas located within 500 meters of the backbone network connected. We now divide those located outside of this connection radius into a maintained control group (those located more than 3,500 meters from the backbone network) and three additional treatment groups that are equally spaced, geographically: those 500–1,500 meters, 1,500–2,500 meters, and 2,500–3,500 meters from the backbone respectively. We compare the four groups closest to the backbone to those furthest away, before and after the arrival of submarine cables on the coast. As seen

[41] None of the surveys we use elicited respondents' migration status or place of birth in both pre and post survey rounds conducted during our data period. (QLFS contains a question about migration, but the variable is missing for the majority of the sample). But in South Africa and Tanzania, it is possible to run (1) with migration status on the left-hand side by using another data source (South Africa) or adding a later survey round conducted by DHS (Tanzania).

in the figure, the estimated coefficients on $SubmarineCables_{c(i)t}$ interacted with indicators for the three additional treatment groups are statistically insignificant and generally near zero. Consequently, the estimated coefficient for those nearest to the backbone network remains essentially unchanged in comparison to the estimates in Subsection IVB (except in the Afrobarometer sample, where the estimate increases in magnitude when individuals intermediate distance from the backbone are excluded from the control group[42]). We conclude that the estimated increase in employment in connected areas is not due to shifting of employment across space.

We now explore how fast Internet affects structural change as measured by occupational employment shares in Africa.

E. *Fast Internet and Employment in Skilled and Unskilled Jobs*

The overall response of employment to the arrival of fast Internet in Africa is made up of underlying changes in job creation and destruction across specific occupations and the sectors associated with those occupations. How technological change affects occupational and sectoral employment shares is especially important in poor countries (Herrendorf, Rogerson, and Valentinyi 2014). To explore this question, we distinguish between *jobs* and *workers*. "The employment rate in occupation X" will here mean the probability of holding a job in occupation X (not the overall employment rate of workers who "permanently" belong to occupation X). We believe that the changes in occupational employment rates we document mostly reflect changes in the size of different sectors.[43] However, readers can alternatively interpret the results in this subsection as reflecting a combination of within- and across-sector changes in employment in skilled and unskilled occupations.

In Table 5 we use the DHS and QLFS datasets, where occupations are recorded and can thus be categorized. In the first two columns of panel A, we define skilled and unskilled employment categories following the ILO's ISCO categorization of occupations' skill level (ILO 2012).[44] In the DHS countries and South Africa, the arrival of fast Internet increases the probability that an individual holds a skilled job by respectively 4.4 and 1.4 percent. The probability of unskilled employment is statistically unaffected in both the first group of countries and South Africa. Our findings thus imply a *positional* skill bias of fast Internet in Africa that is directionally similar to what has been found for computerization and fast Internet in the United States

[42] In the Afrobarometer sample, the point estimate for areas that are intermediate distance from the backbone network suggests that these may also experience employment *gains*, but the estimates are far from significant.

[43] In addition to our expectation that low rates of tertiary education make African workers comparatively likely to switch sectors, this is because we in Section V find that when fast Internet arrives, there are noteworthy changes in firm entry across sectors in South Africa and an expansion of the manufacturing sector in Ethiopia. Note that we cannot estimate how the impact of fast Internet differs for individuals who "permanently" belong to different occupations because none of our individual level datasets include (nonmissing) information on the occupations workers' *past* jobs belong to.

[44] Unskilled jobs (ISCO level 1) "typically involve performance of simple and routine physical or manual tasks" (ILO 2012, p. 12). ILO defines the following DHS occupational categories as skilled work: professional, sales, services, and skilled manual; and the following DHS occupational categories as unskilled work: self-employed agriculture, domestic, and unskilled manual. ILO defines the following QLFS occupational categories as skilled work: legislative, professional, services, skilled manufacturing, and technical; and the following QLFS occupational categories as unskilled work: elementary and domestic.

Table 5—Fast Internet and Employment in Skilled and Unskilled Positions

Outcome	Skilled (0/1)	Unskilled (0/1)	Highly skilled (0/1)	Somewhat skilled (0/1)	Moderately skilled (0/1)	Unskilled (0/1)
Unit of analysis	Individual					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. DHS						
SubmarineCables	0.044	0.003	0.017		0.027	0.003
× connected	(0.018)	(0.015)	(0.011)		(0.020)	(0.015)
Observations	59,966	59,923	59,923		59,957	59,923
Mean of outcome	0.58	0.11	0.09		0.49	0.11
Country × time FE	Yes	Yes	Yes		Yes	Yes
Grid-cell × connected FE	Yes	Yes	Yes		Yes	Yes
Panel B. SA-QLFS						
SubmarineCables	0.014	−0.001	0.001	0.003	0.010	−0.001
× connected	(0.006)	(0.005)	(0.004)	(0.004)	(0.006)	(0.005)
Observations	280,641	280,641	280,641	280,641	280,641	280,641
Mean of outcome	0.50	0.22	0.08	0.08	0.34	0.22
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The DHS sample includes Benin, D.R. Congo, Ghana, Kenya, Namibia, Nigeria, Tanzania, and Togo. The QLFS survey is from South Africa. Survey years for each DHS country are reported in online Appendix Table A1. QLFS data are 2008:I–2010:II. Grid-cells are 0.1 × 0.1 decimal degrees, which is roughly 10 × 10 km. Location FEs are enumeration areas in South Africa QLFS. Time FEs are quarters in QLFS and years in DHS. Individuals (locations) are considered connected if they are closer than 0.5 km to the backbone network. We categorize occupations' skill level following the ILO's ISCO categorization. For DHS, the highly skilled occupation group includes professional; the moderately skilled group clerical, skilled manufacturing, retail and sales, services, and employed agriculture; and the unskilled group unskilled manufacturing, self-employed agriculture, and domestic work. There are no observations in the somewhat skilled occupation group in the DHS sample. For QLFS, the highly skilled occupation group includes legislative work and professional; the somewhat skilled group technical work; the moderately skilled group clerical, skilled agriculture, crafts workers, services, and plant workers; and the unskilled group elementary work and domestic work. The skilled category corresponds to the highly, somewhat, and moderate skilled occupation groups. Robust standard errors clustered at the level of location FEs in parentheses.

and Europe. However, it is noteworthy that the large estimated increase in overall employment in Africa is driven by increased employment in skilled occupations.

In columns 3–6 of Table 5, we break the skilled category (ISCO levels 2–4) into its subcategories as defined by the ILO.[45] We lack power to estimate the impact on each of these separately with precision, but the point estimates are nevertheless worth reporting. In both the DHS countries and South Africa, a relatively large estimated increase in the probability of moderately skilled (ISCO level 2) employment appears to contribute most to the overall increase in skilled employment. The point estimates also point toward a sizable increase in highly skilled (ISCO level 4) employment in the DHS countries, when fast Internet becomes available.

[45] Moderately skilled jobs (ISCO level 2) "typically involve performance of tasks such as operating machinery and electronic equipment; driving vehicles; maintenance and repair of electrical and mechanical equipment; and manipulation, ordering and storage of information." Somewhat skilled jobs (ISCO level 3) "typically involve performance of complex technical and practical tasks that require an extensive body of factual, technical and procedural knowledge in a specialized field." Highly skilled jobs (ISCO level 4) "typically involve performance of tasks that require complex problem-solving, decision-making and creativity based on an extensive body of theoretical and factual knowledge in a specialized field" (ILO 2012, pp. 12–13). There are no observations in the ISCO level 3 categories in the DHS sample.

McMillan, Rodrik, and Verduzco-Gallo (2014) and McMillan and Harttgen (2014) show that the overall trends in structural change in Africa improved after 2000. The estimates in Table 5 suggest that greater and cheaper access to information and communication may be among the changes in the economic environment that helped shift workers toward occupations that usually display higher productivity. We return to this question in Sections V and VI, where we investigate whether firms whose productivity increased, or which started exporting more, also hired more workers when fast Internet became available, and how the technology affects incomes in Africa. In subsection IVF we explore how job inequality in Africa responds to the arrival of fast Internet.

F. *Fast Internet and Job Inequality*

Given the lack of direct evidence on the factor bias of ICT in poor countries, it is a priori unclear if fast Internet affects job inequality across the educational attainment range in Africa in the same way that ICT has been shown to do in rich countries (Katz and Autor 1999; Bond and Van Reenen 2007; Goldin and Katz 2007; Akerman, Gaarder, and Mogstad 2015). We investigate this question in Table 6, where we report results from interacting $SubmarineCables_{it} \times Connected_i$ with educational attainment.

The estimated increase in the employment rate is of comparable magnitude for those with primary school, those with secondary school, and those with tertiary education in all three samples.[46] Those with primary school in fact see a moderately bigger estimated employment gain than those with secondary school in all three samples, though not statistically significantly differentially so. In the Afrobarometer countries (but not in the DHS countries and South Africa) our estimates suggest that fast Internet also increases the employment rate for those who did not complete primary school.

In the DHS countries, those with tertiary education see by far the biggest increase in *skilled* employment. The smaller, but nevertheless noteworthy, estimated increase in skilled employment is of very similar magnitude for those with primary and those with secondary education, but more precisely estimated for the latter group.[47] Employment in unskilled occupations increases significantly for those with primary school in both the DHS countries and South Africa.

The results in Table 6 in combination with those in Table 5 illuminate important similarities and differences in the way modern ICT technologies affect job inequality in Africa versus rich countries. We saw in Table 5 that the skill complementarity of fast Internet as defined by its relative impact on net creation (and/or saving) of high- and low-skill *jobs* in Africa resembles the skill-bias documented in the West. However, the results in Table 6 show that, in the eight DHS countries in our sample, those with almost no education are the only group of *individuals* whose employment outcomes do not benefit from fast Internet. In these countries, while the technology

[46] The initial employment rates that the estimates shown can be compared to are in online Appendix Table A5.

[47] For the South African sample, the increase in skilled employment is imprecisely estimated for all educational attainment groups. (The point estimate for the increase in skilled employment is of roughly comparable magnitude across the educational attainment range, but it is noteworthy that the estimate for those without primary school in South Africa is comparatively large (but statistically insignificant). Recall, however, ILO's fairly broad definition of skilled occupations, e.g., including jobs in services).

Table 6—Fast Internet, Employment, and Employment in Skilled and Unskilled Positions, by Educational Attainment

Outcome	Employed (0/1)	Skilled (0/1)	Unskilled (0/1)
Unit of analysis		Individual	
	(1)	(2)	(3)
DHS			
SubmarineCables × connected			
× Not primary	−0.013	−0.001	−0.024
	(0.032)	(0.036)	(0.023)
× Primary	0.061	0.033	0.035
	(0.023)	(0.028)	(0.019)
× Secondary	0.047	0.033	0.012
	(0.015)	(0.018)	(0.015)
× Higher	0.067	0.119	−0.041
	(0.028)	(0.036)	(0.027)
Observations	59,914	59,966	59,923
Country × time FE	Yes	Yes	Yes
Grid-cell × connected FE	Yes	Yes	Yes
Afrobarometer			
SubmarineCables × connected			
× Not primary	0.109		
	(0.055)		
× Primary	0.077		
	(0.042)		
× Secondary	0.060		
	(0.050)		
× Higher	0.097		
	(0.053)		
Observations	7,902		
Country × time FE	Yes		
Grid-cell × connected FE	Yes		
SA-QLFS			
SubmarineCables × connected			
× Not primary	0.012	0.021	−0.009
	(0.017)	(0.018)	(0.018)
× Primary	0.028	0.007	0.021
	(0.011)	(0.011)	(0.010)
× Secondary	0.022	0.015	0.007
	(0.012)	(0.013)	(0.009)
× Higher	0.019	0.012	0.007
	(0.011)	(0.012)	(0.008)
Observations	277,737	277,737	277,737
Time FE	Yes	Yes	Yes
Location FE	Yes	Yes	Yes

Notes: The DHS sample includes Benin, D.R. Congo, Ghana, Kenya, Namibia, Nigeria, Tanzania, and Togo. The Afrobarometer sample includes Benin, Ghana, Kenya, Madagascar, Mozambique, Nigeria, Senegal, South Africa, and Tanzania. The QLFS survey is from South Africa. Survey years for each Afrobarometer and DHS country are reported in online Appendix Table A1. QLFS data are 2008:I–2010:II. Grid-cells are 0.1 × 0.1 decimal degrees, which is roughly 10 × 10 km. Location FEs are enumeration areas in South Africa QLFS. Time FEs are quarters in QLFS and years in Afrobarometer and DHS. Individuals (locations) are considered connected if they are closer than 0.5 km to the backbone network. We categorize occupations' skill level following the ILO's ISCO categorization (more details in Table 5). Afrobarometer does not record respondents' occupation. Controls for educational attainment (primary school not completed, primary school completed, secondary school completed, and higher education) are included. Robust standard errors clustered at the level of location FEs in parentheses.

increases *skilled* employment among more educated workers the most, fast Internet thus reduces (un)employment inequality across the majority of the adult population. In the Afrobarometer countries, even those who did not complete primary school see significant employment gains with the arrival of fast Internet. These results partially contrast both with existing findings on computers and fast Internet as SBTC in rich countries, and with the adverse estimated effect of another important form of globalization—trade liberalization—on (income and wage) inequality in developing countries (Goldberg and Pavcnik 2007; Harrison, McLaren, and McMillan 2011; Goldberg 2015).

The magnitude of the estimated increase in "any" and skilled employment with the arrival of fast Internet in Africa is surprising, but the overall pattern that emerges across the educational attainment range—particularly in the DHS sample—is arguably less surprising. The especially large increase in skilled employment for workers with tertiary education suggests that fast Internet is in one sense a high education-biased technology in these countries. The productivity of workers with less education may also benefit from fast Internet, however, if for example employers choose to provide targeted on-the-job training to such workers (Green, Dickerson, and Arbache 2001; Frías, Kaplan, and Verhoogen 2009). This may help explain why we observe a considerable increase in skilled employment also for less educated workers in the DHS countries. The increase in employment for workers with primary school in all three samples (and those without primary school in the Afrobarometer countries) may be due, for example, to the emergence of new types of positions that are complementary to jobs wherein more educated workers make more direct use of Internet technology.

In the next section we investigate *how* fast Internet affects employment in Africa.

V. Understanding How Fast Internet Affects Employment in Africa

A. *Firm Entry*

The changes in average speeds and use of the Internet after the arrival of the submarine cables we documented in Subsection IVA suggest that new, and new forms of, employment may arise both through extensive margin (new Internet users) and intensive margin (different use of the Internet by existing users) responses. In this subsection, we analyze how fast Internet affects firm entry and exit; in the next two subsections we explore possible changes in the productivity and exports of existing firms.

We first use a dataset from South Africa's CIPC that records the names and addresses (including zip codes) of firms that register or de-register, and the date of registration/de-registration. We run (1) at the zip code $\times$ quarter level. As seen in Table 7, we find a significant increase in net firm entry per quarter of around 23 percent when fast Internet arrives in South Africa. This overall impact is due both to a large increase in firm entry, and to a decrease in firm exit of similar magnitude. Greater net firm entry may thus help explain the estimated increase in employment when fast Internet arrives in South Africa.

As also seen in Table 7, we find a significant increase in net firm entry in many sectors, but the biggest point estimates are seen in sectors that use ICT extensively (World Bank 2006), such as finance and services.

Table 7—Fast Internet and Firm Entry in South Africa

Unit of analysis	Location								
Outcome	All firms (asinh) (1)	Agri-culture (asinh) (2)	Retail/sales (asinh) (3)	Tech-nology (asinh) (4)	Manu-facturing (asinh) (5)	Services (asinh) (6)	Finan-cial (asinh) (7)	White-collar (asinh) (8)	Blue-collar (asinh) (9)
Panel A. Net firm entry									
SubmarineCables × connected	0.227	0.026	0.086	0.022	0.020	0.120	0.158	0.060	0.076
	(0.079)	(0.014)	(0.043)	(0.028)	(0.013)	(0.052)	(0.063)	(0.033)	(0.034)
Observations	38,108	38,108	38,108	38,108	38,108	38,108	38,108	38,108	38,108
Panel B. Firm entry									
SubmarineCables × connected	0.089	0.007	0.033	0.015	0.012	0.046	0.103	0.029	0.043
	(0.038)	(0.012)	(0.030)	(0.020)	(0.009)	(0.034)	(0.046)	(0.023)	(0.025)
Observations	38,108	38,108	38,108	38,108	38,108	38,108	38,108	38,108	38,108
Panel C. Firm exit									
SubmarineCables × connected	−0.088	−0.021	−0.053	−0.007	−0.008	−0.059	−0.043	−0.031	−0.032
	(0.038)	(0.012)	(0.023)	(0.015)	(0.009)	(0.029)	(0.029)	(0.017)	(0.018)
Observations	38,108	38,108	38,108	38,108	38,108	38,108	38,108	38,108	38,108
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Location × connected FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Data at location × quarter level, 2007–2014. Location FEs are zip codes. Time FEs are quarters. See online Appendix Table A4 for information on how firms are categorized into sectors. Sectors not shown are mining and tourism. The point estimate is positive and insignificant for these sectors. Firms are considered connected if they are closer than 0.5 km to the backbone network. Robust standard errors clustered at the level of location FEs in parentheses.

These results reinforce the view that fast Internet's impact on structural change was at least partly favorable insofar as productivity in the sectors that saw the biggest increase in net firm entry is likely high relative to productivity in other sectors in South Africa.

B. *Firm and Labor Productivity in Existing Firms*

OLS Results.—We have seen that access to fast Internet increases firm entry, which appears to contribute to its impact on employment rates. Does the new technology also affect employment within existing firms, and, if so, why? To investigate these questions, we first use Ethiopia's LMMIS dataset of large and medium-sized manufacturing firms, which is, to our knowledge, the only African dataset with detailed enough information and the geographical and time coverage needed to estimate changes in firms' production function with the arrival of fast Internet. We restrict the sample to firms that are observed both before and after the submarine cable that gets connected to Ethiopia's network arrives on the coast.

In columns 1–6 of Table 8, we continue to use a similar specification and definition of right-hand side variables as in (1), but i now represents a firm and observations are at the firm × year level. The estimated increase in total employment per firm when fast Internet arrives is about 16 percent in column 1, where we control for firm and year fixed effects[48], and about 22 percent in column 2, where

[48] We match firms across years using all available information in LMMIS and cross-check our matches against those of Abebe, McMillan, and Serafinelli (2017) (the most in-depth and authoritative existing work on such

Table 8—Fast Internet, Employment, Output Elasticity of Labor, and Productivity in Ethiopian Firms

Outcome	Employees (asinh)		Skilled employees (asinh)		Unskilled employees (asinh)		Value added (asinh)	Value added (asinh)	Value added (asinh)
Unit of analysis	Firm								
Method	OLS		OLS		OLS		OLS	Reg LP	Adj LP
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
SubmarineCables × connected	0.156	0.224	0.034	0.201	0.115	0.124			
	(0.091)	(0.081)	(0.080)	(0.106)	(0.078)	(0.095)			
Capital							0.276	0.263	0.249
							(0.018)	(0.023)	(0.020)
Unskilled							0.337	0.127	0.135
							(0.064)	(0.043)	(0.045)
Skilled							0.497	0.198	0.198
							(0.043)	(0.025)	(0.026)
SubmarineCables × connected									
× Unskilled							−0.176	−0.048	−0.063
							(0.058)	(0.031)	(0.034)
× Skilled							0.026	0.017	0.016
							(0.033)	(0.027)	(0.027)
Control for productivity							No	Yes	Yes
Control for SubmarineCables × connected × productivity							No	No	Yes
Outcome									Productivity
SubmarineCables × connected									0.127
									(0.058)
Time FE	Yes	No	Yes	No	Yes	No	No	No	No
Grid-cell × connected FE	No	Yes	No	Yes	No	Yes	Yes	Yes	Yes
Industry × time FE	No	Yes	No	Yes	No	Yes	Yes	Yes	Yes
Firm FE	Yes	No	Yes	No	Yes	No	No	No	No
Observations	5,360	5,360	5,360	5,360	5,360	5,360	4,321	4,321	4,321

Notes: Data from the Ethiopian LMMIS manufacturing firm census. 2006–2013. Grid-cells (for location FEs) are 0.1 × 0.1 decimal degrees, which is roughly 10 × 10 km. Time FEs are years. Firms are considered connected if they are closer than 0.5 km to the backbone network. The sample is restricted to firms observed both before and after submarine cable arrival and includes 1,103 firms. Skilled (unskilled) positions are defined as those earning more (less) than 800 Birr per year, approximately the sample salary median. Capital is the average of start-of-year and end-of-year book value. The production function specifications allow fast Internet to directly affect value added via a change in the intercept (not shown). Robust standard errors clustered at grid-cell level in parentheses.

we control instead for grid-cell × connected and industry × year fixed effects. The estimated increase in skilled and unskilled positions per firm (as discussed in Section II, these are proxied using salary bins) is respectively 3.4 and 11.5 percent (but not statistically significant) when we control for firm and year fixed effects and 20 and 12.4 percent when we control instead for grid-cell × connected and industry × year fixed effects (in which case the estimated increase in skilled employment is statistically significant). The firm level estimates of changes in employment when fast Internet arrives in Ethiopia are thus broadly comparable to the individual level employment results for the broader samples of African countries

matching in LMMIS). Our matches are nearly identical to theirs. We are grateful to the authors for allowing this cross-check.

and South Africa in Section IV. The relative increase in unskilled positions may be larger in Ethiopian manufacturing firms.

In the last 3 columns of Table 8, we explore whether the increase in employment in Ethiopian manufacturing firms may be explained by an increase in the output elasticity of labor and/or firm level productivity with the arrival of fast Internet. We start with the following OLS regression:

$$(2)\quad \mathrm{va}_{ij(i)jt} = x'_{ijt}\alpha + \mathit{SubmarineCablesConnected}_{ijt}x'_{ijt}\beta + \delta_{j(i)} \times \mathit{Connected}_i + \psi_{jt} + \epsilon_{ij(i)jt},$$

where va_{ijt} is the value added of firm i, in grid-cell $j(i)$, industry j, and year t; x'_{ijt} is a set of inputs (labor, capital) used by the firm and a constant term; ψ_{jt} is an industry $\times$ year fixed effect, and the other variables are as defined previously.[49] Results from this specification are in column 7 in of Table 8. The coefficients on capital and labor are of similar magnitude to what other studies have found for comparable contexts. The estimated output elasticity of labor in skilled positions increases (insignificantly) from 0.497 to 0.523 with the arrival of fast Internet, but that of labor in unskilled positions falls significantly.

Structural Estimation.—OLS estimates of the share of variation in output attributable to different input factors may partly reflect the fact that some input factors—such as labor—are chosen after a firm's productivity (unobserved to the researcher) is fully or partially known to the firm. Olley and Pakes (1996)—henceforth, OP—and Levinsohn and Petrin (2003)—henceforth, LP—developed practical methods that help overcome such simultaneity bias. The commonly used LP method involves using intermediate inputs to proxy for a firm's unobserved productivity in the production function (see LP for details). We posit the following "structural" model:

$$(3)\quad \mathrm{va}_{ijt} = l_{ijt}\theta + \mathit{SubmarineCablesConnected}_{ijt}l_{ijt}\phi + \kappa k_{ijt} + \omega_{ijt} + \epsilon_{ijt},$$

where l_{ijt} are labor inputs and the productivity term ω_{ijt} subsumes the constant term and the fixed effects. The variable ϵ_{ijt} represents a standard i.i.d. error term capturing unanticipated shocks to productivity and measurement error. We present LP estimates in column 8 of Table 8. As expected, both $\hat{\theta}$ and $\hat{\phi}$ are now much smaller in magnitude than the OLS estimates, and the $\hat{\phi}$ from the interaction of $\mathrm{SubmarineCablesConnected}_{ijt}$ and workers in unskilled positions is no longer significant.

De Loecker (2011) points out an important methodological tension when using the OP/LP methods to investigate how a change in the operating environment affects output elasticities. Suppose that a firm's productivity *itself* is influenced by the change in the operating environment. If the productivity response in turn influences hiring, investment, and value added—as conventional models of firm behavior pre-

[49] To ease comparison with the structural results in columns 8 and 9 of Table 8, we interact only labor and the constant term with $\mathrm{SubmarineCablesConnected}_{ij(i)jt}$.

dict—then changes in the coefficients on labor and capital estimated using methods that do not account for the firm level productivity response will be incorrect.

Inspired by De Loecker (2011), we assume the following law-of-motion for firm productivity:

$$\omega_{ij,t+1} = \alpha\omega_{ijt} + \tau SubmarineCablesConnected_{ijt+1} + \delta_{j(i)} \times Connected_i + \psi_{jt} + \xi_{ij,t+1}, \tag{4}$$

where grid-cell × connected and industry × year fixed effects control for differences across space in, and industry-wide shocks to, productivity. We continue to use the LP estimation procedure, but adjust the method to allow both the output elasticity of labor and firm-level productivity itself to change. We first estimate ϕ while controlling for a possible response in firm level productivity to fast Internet. As in the conventional LP method, we use a flexible polynomial in the other input factors—including intermediate inputs—to proxy for ω_{ijt}. The adjustment we make in this first step is that we include $SubmarineCablesConnected_{ijt}$ among the factors included in the polynomial. We run

$$\mathrm{va}_{ijt} = l_{ijt}\theta + SubmarineCablesConnected_{ijt} l_{ijt}\phi + \Psi\left[m_{ijt}, k_{ijt}, SubmarineCablesConnected_{ijt}, \delta_{\mathbf{j}(i)} \times Connected_i, \psi_{jt}\right] + \epsilon_{ijt}, \tag{5}$$

where $\Psi[m_{ijt}, k_{ijt}, SubmarineCablesConnected_{ijt}, \delta_{j(i)} \times Connected_i, \psi_{jt}]$ is a polynomial of inputs used (m_{ijt}), capital (k_{ijt}), access to fast Internet, and grid-cell × connected and industry × year fixed effects.

The estimated effect of fast Internet on the output elasticity of labor estimated through this procedure is reported in column 9 of Table 8. The estimated decrease in the output elasticity of labor in unskilled positions increases in absolute magnitude to -0.063 and becomes significant. The estimated increase in the output elasticity of labor in skilled positions is very similar to the estimate from the conventional LP method: 0.016.

In the second step of the procedure, we estimate the coefficient on capital by GMM using the moment condition $E[\xi_{ijt}(\kappa)k_{ijt}] = 0$, which is motivated by the assumption that capital cannot be adjusted in response to unobserved shocks to productivity.[50] Note that $\hat{\xi}_{ij,t+1}$ is obtained by taking the OLS residual from (4), where ω_{ijt} and ω_{ijt-1} come from applying (3), that is, by subtracting the labor coefficients estimated in (5) and the coefficient for capital from the predicted value added obtained from (5).

For our purposes the coefficient on capital is needed only as an input into the procedure for estimating how fast Internet affects firm level productivity. With estimates of the coefficients on labor, capital, and the interaction between labor and $SubmarineCablesConnected_{ijt}$ in hand, we can construct $\hat{\omega}_{ijt}$ using (3) and then estimate the law-of-motion for productivity in the third step. The results are reported

[50] We use the OLS estimates as starting values and bootstrap the standard errors.

in column 9 of the bottom panel of Table 8. The estimated increase in firm level productivity when fast Internet becomes available is around 13 percent and statistically significant.[51]

We conclude that an increase in firm level productivity likely contributes to increased hiring in existing Ethiopian manufacturing firms after the arrival of fast Internet and that changes in the relative output elasticity of workers in skilled and unskilled positions *may* also help explain the hiring response.

C. *Firms' Exports*, *On-the-Job Training*, *and Internet Communication*

In the online Appendix, we use data from the World Bank Enterprise Surveys (WBES) to explore how the arrival of the submarine Internet cables changed the behavior and performance of firms in Ghana, Kenya, Mauritania, Nigeria, Senegal, and Tanzania.

We approximate (1) as closely as we can with the WBES dataset. However, almost all the firms in the WBES sample are located in relatively large cities and towns, and firms' location is reported only at city/town level. We thus classify a city/town as connected if the backbone network passes through its perimeter (in contrast to the approach taken in our other samples, where fine-grained geographical information on an individual or firm's location is available). In light of the WBES firms being clustered in space and the coarser connectivity classification required, we view the WBES results as more suggestive than the rest of our analysis.[52] We thus present the full WBES analysis in the online Appendix, and briefly summarize the results here.

Our estimates suggest that access to fast Internet leads African firms in the WBES sample to employ about 14–17 percent more workers per firm. The firm level estimates of changes in employment when fast Internet arrives in the WBES countries are broadly comparable to the individual level employment results in Section IV.

We also find that firms appear more likely to provide on-the-job training to their employees when fast Internet becomes available. This finding may help explain why the technology in general boosts employment not only for highly educated workers, but also less educated workers, in Africa.

We next explore how the composition of firms' sales responds to the arrival of fast Internet. We find evidence of a large increase in direct exports. In light of the existing literature documenting the benefits to firms and employment consequences of exporting (see e.g., Verhoogen 2008; Frías, Kaplan, and Verhoogen 2009; Goldberg et al. 2010b; Atkin, Khandelwal, and Osman 2017), this finding suggests that one way in which fast Internet increases employment in Africa is by making it easier for firms to sell to customers abroad. The increase in exports is also evidence of an interaction between technological change and trade that differs from the trade-induced SBTC analyzed by an existing literature (Wood 1995; Acemoglu 2003; Attanasio, Goldberg, and Pavcnik 2004; Burstein,

[51] There are many potential channels through which fast Internet can boost firm productivity above and beyond the output elasticity of labor. The technology may, for example, allow firms to sell more per unit of marketing cost, give access to information about more efficient production processes, or allow firms to increase the quality of their products. Note that the estimated impact on firm productivity remains statistically significant if we bootstrap also the standard error corresponding to this last step of the procedure, as seen in panel C of online Appendix Table A3.

[52] See the online Appendix for details.

Cravino, and Vogel 2013; Koren and Csillag 2016; Raveh and Reshef 2016). Here, causality runs from technological change to trade, rather than the other way around.

WBES contains information on firms' use of the Internet. We find a significant increase in firms' probability of communicating with clients through email and through a website. Easier online communication with clients may have helped African firms export more when submarine Internet cables reached the continent.

In sum, we have seen evidence in this section indicating that the increase in employment when fast Internet arrived in Africa was driven, in part, by greater firm entry in South Africa; by higher firm level productivity in existing Ethiopian manufacturing firms; and by an increase in exports, on-the-job training, and use of online communication among firms in the six WBES countries in our sample. While the magnitudes of these economic responses is important in their own right, data limitations prevent us from investigating what *share* of the changes in employment patterns they account for. Additional mechanisms likely also played a role.

VI. Fast Internet, Employment, and Incomes

Some would consider employment a means to an end more than an end in itself. In Table 9, we explore how the arrival of fast Internet ultimately affects incomes in Africa. Increasing access to, and lowering the cost of, information and communication may affect incomes also through other channels than employment outcomes. But to the extent that fast Internet affects overall and skilled net job creation, we would a priori also expect such an employment response to ultimately translate into higher incomes.

We follow a growing literature and proxy for average incomes at location level with light density at night as measured by satellites (see e.g., Henderson, Storeygard, and Weil 2012; Bleakley and Lin 2012; Michalopoulos and Papaioannou 2013, 2018; Lowe 2014). In addition to capturing the *aggregate* economic benefits of fast Internet, an advantage of this income proxy is that it is available for all 12 countries in our sample. The National Oceanic and Atmospheric Administration (NOAA) provides pixel-level measures of average night light density from satellite images. We thus construct a grid of such pixels that are 0.1 degree ($\sim$10 km) apart, in the spirit of Michalopoulos and Papaioannou (2013). The estimating equation we use is (1) as throughout the paper; the i subscript now indexes the pixels. In column 1, we see that night light density rises by about 2.4 percent when fast Internet becomes available. Controlling for a nonlinear trend in average incomes that is specific to the connected locations—interactions between the Connected indicator and the time fixed effects—in column 2 increases the estimated impact of fast Internet on night light density to 3.3 percent.

The balanced panel, "high(er)-T" format of the night lights data allows us to trace out how fast Internet affects economic activity over time better than the household surveys we use allow. We do so in Appendix Figure A3 by interacting a location's connectivity status with years-to/since-cable-arrival-dummies, as in an event study. First, we see that year-to-year changes in average incomes in connected locations relative to unconnected locations hover around zero prior to the arrival of submarine cables. Second, relative average incomes in connected locations start to rise the year fast Internet arrives. Third, the rise continues in each of the two following years,

Table 9—Fast Internet and Incomes

Outcome	Light density at night (asinh)	
Unit of analysis	Point	
Sample	NOAA	
	(1)	(2)
SubmarineCables × connected	0.024	0.033
	(0.009)	(0.018)
Observations	80,360	80,360
Country × time FE	Yes	Yes
Grid-cell × connected FE	Yes	Yes
Connected × time FE	No	Yes

Notes: The data is yearly and for 2007–2013. Grid-cells are 0.1 × 0.1 decimal degrees, which is roughly 10 × 10 km. Locations are considered connected if they are closer than 0.5 km to the backbone network. Light density at night proxies for average income at location level. Robust standard errors clustered at grid-cell level in parentheses.

before relative average incomes in connected locations level off in the third "post" year. It thus appears that the impact of fast Internet on incomes in Africa persists over time, but that the *growth* effect may be especially large in the first few years after the submarine cables arrive.

VII. Conclusion

This paper provides evidence on how fast Internet affects employment in Africa. We exploit the gradual arrival of ten submarine Internet cables from Europe in cities on Africa's coast in the late 2000s and early 2010s and interact landing points and times with an indicator for whether a given location is on the terrestrial cable network that connects users with the coast. We first show that both average speeds and use of the technology increase when the submarine cables arrive. We then compare the changes in employment patterns in areas with a bigger versus a smaller increase in access to fast Internet, controlling for location and time effects. In each of 3 different datasets that together cover 12 African countries with a combined population of roughly half a billion people, we find a significant and large relative increase in the employment rate in connected areas when fast Internet becomes available. Extensive prodding of the identifying assumptions that underlie our generalized difference-in-differences approach suggests that these estimates reflect a causal effect of access to fast Internet on employment rates. Employment responses of the magnitude we document indicate that building fast Internet infrastructure may be among the currently feasible policy options with the greatest employment-creating potential in Africa. We also show that the technology's impact is driven by an increase in employment in higher-skill occupations. Finally, fast Internet if anything lowers (un)employment inequality across the educational attainment range in Africa.

The observed changes in average speeds and use of the Internet after the arrival of the submarine cables suggest that new and new types of jobs may have been (net) created both via extensive margin (new Internet users) and intensive margin (different use of the Internet by existing users) responses. We explore these possibilities with more detailed firm-level data available for some countries. In South Africa, firm entry increases—notably in sectors that tend to benefit from ICT—as does the productivity of existing manufacturing firms in Ethiopia, when fast Internet becomes available. We also find more suggestive evidence that fast Internet appears to enable firms in Ghana, Kenya, Mauritania, Nigeria, Senegal, and Tanzania to export more, perhaps in part because online communication with clients became easier.

The impact on job inequality we document indicates that the skill bias of fast Internet in Africa is more nuanced than what has been found for computerization and fast Internet in rich countries. This in turn suggests that the primary explanation for the slow economic progress of poor workers in Africa and other similar contexts during the last few decades is unlikely to be the factor bias of recent technological change. The sectors that ex ante appear to have been most constrained by lack of access to ICT, and that create more "good" jobs when fast Internet becomes available, are broadly speaking sectors associated with high relative productivity in Africa. In at least some of these sectors in some parts of the continent, fast Internet further increases productivity, and appears to enable exporting. This suggests that the technology contributed positively to structural change in Africa during our data period.

Appendix

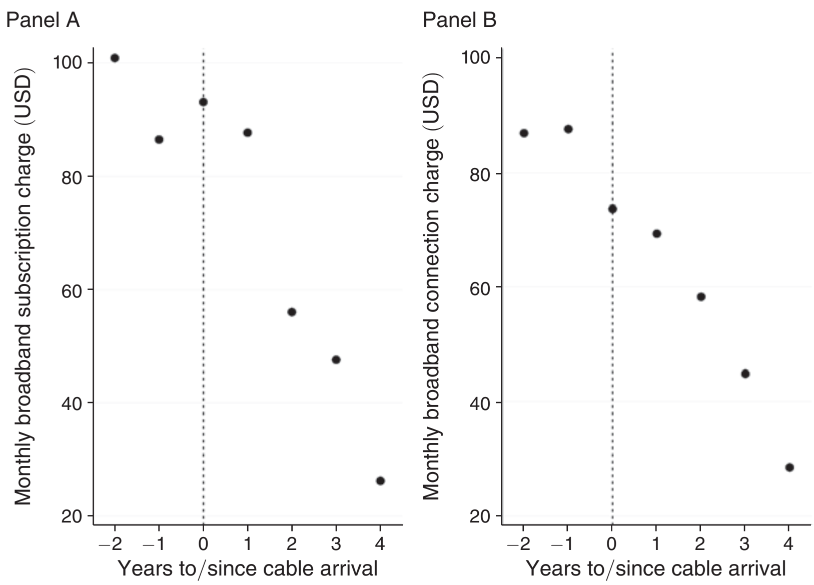

Figure A1. Broadband Subscription and Connection Charges, before and after Submarine Cable Arrival

Notes: This graph plots the coefficients from running a regression with event-time indicators, using a dependent variable from the International Telecommunication Union (ITU). The event time is calculated as the year of outcome measurement net the year of first cable connection. Country and year FEs are included in the regression. The ITU data structure is country × year, and includes observations from all countries used elsewhere in the paper, with the exception of DR Congo (for which data is missing), and is available for every year between 2007 and 2014. The figure plots the effect on monthly fixed-broadband subscription and connection charges measured in USD.

FIGURE A2. ROAD NETWORKS AND ELECTRICITY GRID (SOUTHWESTERN SOUTH AFRICA AS EXAMPLE)

Notes: This graph plots the road networks and electricity grid used in the placebo "treatment" estimations. Road data comes from SEDAC and electricity data from AICD.

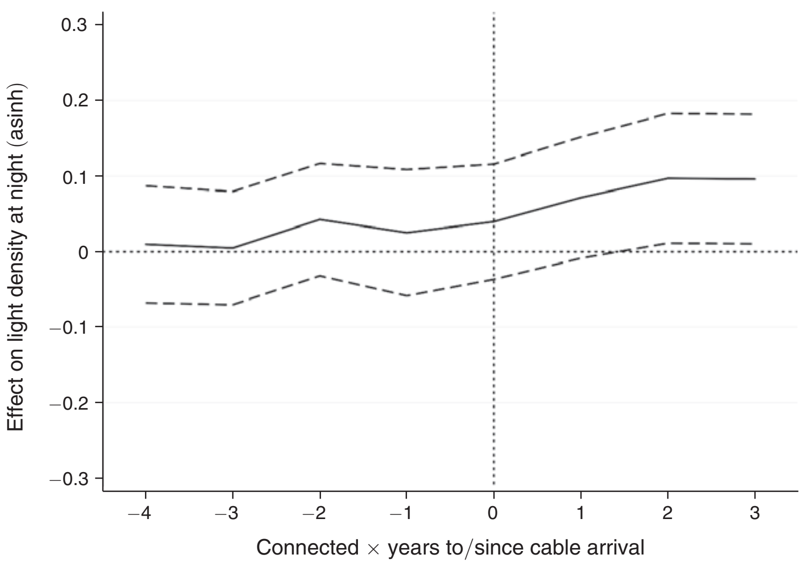

FIGURE A3. FAST INTERNET AND INCOMES OVER TIME

Notes: This graph plots the coefficients from the interaction terms between the connected status and event-time indicators. The event time is calculated as the year of light measurement net the year of first cable connection. A dummy variable is created for each event time and is subsequently interacted with the connected indicator. Country × year FEs are included in the regression.

Table A1—Fast Internet and Employment, Varying Grid-Cell Size

Outcome	Employment (0/1)						
Unit of analysis	Individual						
Grid-cell size	10 km (1)	15 km (2)	20 km (3)	25 km (4)	30 km (5)	35 km (6)	40 km (7)
DHS							
SubmarineCables × connected	0.046	0.047	0.051	0.044	0.052	0.049	0.060
	(0.014)	(0.016)	(0.018)	(0.019)	(0.014)	(0.016)	(0.018)
Observations	59,914	59,914	59,914	59,914	59,914	59,914	59,914
Mean of outcome	0.68	0.68	0.68	0.68	0.68	0.68	0.68
Afrobarometer							
SubmarineCables × connected	0.077	0.076	0.080	0.076	0.069	0.071	0.073
	(0.037)	(0.038)	(0.036)	(0.037)	(0.034)	(0.035)	(0.034)
Observations	7,918	7,918	7,918	7,918	7,918	7,918	7,918
Mean of outcome	0.58	0.58	0.58	0.58	0.58	0.58	0.58
Country × time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Grid-cell × connected FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The DHS sample includes Benin, D.R. Congo, Ghana, Kenya, Namibia, Nigeria, Tanzania, and Togo. The Afrobarometer sample includes Benin, Ghana, Kenya, Madagascar, Mozambique, Nigeria, Senegal, South Africa, and Tanzania. Survey years for each DHS and Afrobarometer country are reported in online Appendix Table A1. Grid-cells range from 10 × 10 km to 40 × 40 km. Time is years in both datasets. Individuals (locations) are considered connected if they are closer than 0.5 km to the backbone network. Robust standard errors clustered at grid-cell level in parentheses, using the same grid-cell size as stated in the column headers.

TABLE A2—FAST INTERNET AND EMPLOYMENT, VARYING THE SAMPLE AND TERRESTRIAL CABLES CONSIDERED

Outcome	Employment (0/1)		
Unit of analysis	Individual		
Sample	DHS (1)	Afrobarometer (2)	SA-QLFS (3)
Panel A. Backbone cables reported in 2 maps			
SubmarineCables × connected	0.048	0.045	0.024
	(0.015)	(0.040)	(0.009)
Observations	59,914	7,900	280,641
Mean of outcome	0.67	0.58	0.71
Panel B. Excluding observations close to landing station			
SubmarineCables × connected	0.054	0.086	0.025
	(0.016)	(0.042)	(0.008)
Observations	51,129	5,871	259,177
Mean of outcome	0.67	0.58	0.71
Panel C. Excluding observations with distance to backbone > 5 km			
SubmarineCables × connected	0.040	0.079	0.022
	(0.015)	(0.038)	(0.008)
Observations	49,982	7,062	231,242
Mean of outcome	0.67	0.58	0.72
Country × time FE	Yes	Yes	No
Grid-cell × connected FE	Yes	Yes	No
Time FE	No	No	Yes
Location FE	No	No	Yes

Notes: The DHS sample includes Benin, D.R. Congo, Ghana, Kenya, Namibia, Nigeria, Tanzania, and Togo. The Afrobarometer sample includes Benin, Ghana, Kenya, Madagascar, Mozambique, Nigeria, Senegal, South Africa, and Tanzania. The QLFS survey is from South Africa. Survey years for each DHS and Afrobarometer country are reported in online Appendix Table A1. QLFS data are 2008:I–2010:II. Grid-cells are 0.1 × 0.1 decimal degrees, which is roughly 10 × 10 km. Location FEs are enumeration areas in South Africa QLFS. Time FEs are quarters in QLFS and years in DHS and Afrobarometer. Individuals (locations) are considered connected if they are closer than 0.5 km to the backbone network. Panel A defines the backbone network as the intersection of AfTerFibre's (2014) map and www.africabandwidthmaps.com map from 2013:II. Panel B excludes observations that are closer than 20 km to a landing station. Robust standard errors clustered at the level of location FEs in parentheses.

REFERENCES

Abebe, Girum, Margaret McMillan, and Michel Serafinelli. 2017. "Foreign Direct Investment and Knowledge Diffusion in Poor Locations: Evidence from Ethiopia." Unpublished.

Acemoglu, Daron. 2003. "Patterns of Skill Premia." *Review of Economic Studies* 70 (2): 199–230.

Acemoglu, Daron, and David H. Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics,* Vol. 4, edited by Orley Ashenfelter and David E. Card, 1043–1171. Amsterdam: Elsevier.

Adepetun, Adeyemi. 2014. "Nigeria's ICT Sector Finds Strength in E-Commerce Growth." *The Guardian*, October 15. https://allafrica.com/stories/201410151110.html.

AfTerFibre. 2014. "African Terrestrial Fibre." http://afterfibre.net (accessed November 2014).

Aghion, Philippe, Robin Burgess, Stephen Redding, and Fabrizio Zilibotti. 2005. "Entry Liberalization and Inequality in Industrial Performance." *Journal of the European Economic Association* 3 (2–3): 291–302.

Akamai. 2012. *The State of the Internet: 4th quarter, 2012 Report* https://www.slideshare.net/AkamaiTechnologies/q4-2012-sotiweb.

Aker, Jenny C. 2010. "Information from Markets near and Far: Mobile Phones and Agricultural Markets in Niger." *American Economic Journal: Applied Economics* 2 (3): 46–59.

Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. 2015. "The Skill Complementarity of Broadband Internet." *Quarterly Journal of Economics* 130 (4): 1781–824.

Akue-Kpakpo, Abossé. 2013. *Study on International Internet Connectivity in Sub-Saharan Africa*. Geneva: International Telecommunication Union.

Allen, Treb. 2014. "Information Frictions in Trade." *Econometrica* 82 (6): 2041–83.
Amiti, Mary, and Lisa Cameron. 2012. "Trade Liberalization and the Wage Skill Premium: Evidence from Indonesia." *Journal of International Economics* 87 (2): 277–87.
Amiti, Mary, and Donald R. Davis. 2012. "Trade, Firms, and Wages: Theory and Evidence." *Review of Economic Studies* 79 (1): 1–36.
Antràs, Pol, Luis Garicano, and Esteban Rossi-Hansberg. 2008. "Organizing Offshoring: Middle Managers and Communication Costs." In *The Organization of Firms in a Global Economy*. edited by Elhanan Helpman, Dalia Marin, and Thierry Verdier, 311–40. Cambridge, MA: Harvard University Press.
Atasoy, Hilal. 2013. "The Effects of Broadband Internet Expansion on Labor Market Outcomes." *Industrial and Labor Relations Review* 66 (2): 315–45.
Atkin, David, Azam Chaudhry, Shamyla Chaudry, Amit K. Khandelwal, and Eric Verhoogen. 2017. "Organizational Barriers to Technology Adoption: Evidence from Soccer Ball Producers in Pakistan." *Quarterly Journal of Economics* 132 (3): 1101–64.
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.
Attanasio, Orazio, Pinelopi K. Goldberg, and Nina Pavcnik. 2004. "Trade Reforms and Wage Inequality in Colombia." *Journal of Development Economics* 74 (2): 331–66.
Atuanya, Patrick, and Bala Augie. 2013. "Online sales boom as Konga, Jumia lure shoppers." *Business Daily*, December 6. https://www.businessdayonline.com/exclusives/article/online-sales-boom-as-konga-jumia-lure-shoppers/.
Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics* 90 (2): 300–23.
Autor, David H., Lawrence F. Katz, and Alan B. Krueger. 1998. "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics* 113 (4): 1169–213.
Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118 (4): 1279–333.
Banerjee, Abhijit V., and Andrew F. Newman. 1993. "Occupational Choice and the Process of Development." *Journal of Political Economy* 101 (2): 274–98.
Banerji, Sourangsu, and Rahul Singha Chowdhury. 2013. "Wi-Fi and WiMAX: A Comparative Study." *Indian Journal of Engineering* 2 (5): 51–54.
Bartel, Ann P., and Nachum Sicherman. 1999. "Technological Change and Wages: An Interindustry Analysis." *Journal of Political Economy* 107 (2): 285–325.
Baumol, William J. 2013. *The Cost Disease.* New Haven: Yale University Press.
BBC. 2009. "East Africa gets high-speed web." July 23. http://news.bbc.co.uk/2/hi/africa/8165077.stm.
Beaman, Lori, and Jeremy Magruder. 2012. "Who Gets the Job Referral? Evidence from a Social Networks Experiment." *American Economic Review* 102 (7): 3574–93.
Beaudry, Paul, Mark Doms, and Ethan Lewis. 2010. "Should the Personal Computer Be Considered a Technological Revolution? Evidence from U.S. Metropolitan Areas." *Journal of Political Economy* 118 (5): 988–1036.
Beaudry, Paul, and David A. Green. 2003. "Wages and Employment in the United States and Germany: What Explains the Differences?" *American Economic Review* 93 (3): 573–602.
Beaudry, Paul, and David A. Green. 2005. "Changes in U.S. Wages, 1976–2000: Ongoing Skill Bias or Major Technological Change?" *Journal of Labor Economics* 23 (3): 609–48.
Berman, Eli, John Bound, and Zvi Griliches. 1994. "Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufactures." *Quarterly Journal of Economics* 109 (2): 367–97.
Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics* 119 (1): 249–75.
Bleakley, Hoyt, and Jeffrey Lin. 2012. "Portage and Path Dependence." *Quarterly Journal of Economics* 127 (2): 587–644.
Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts. 2013. "Does Management Matter? Evidence from India." *Quarterly Journal of Economics* 128 (1): 1–51.
Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. "Americans Do It Better: US Multinationals and the Productivity Miracle." *American Economic Review* 102 (1): 167–201.
Bloom, Nicholas, and John Van Reenen. 2007. "Measuring and Explaining Management Practices across Firms and Countries." *Quarterly Journal of Economics* 122 (4): 1351–408.
Bond, Stephen, and John Van Reenen. 2007. "Microeconometric Models of Investment and Employment." In *Handbook of Econometrics*. Vol. 6, edited by James J. Heckman and Edward E. Leamer, 4417–98. Amsterdam: North-Holland.

Brambilla, Irene, Daniel Lederman, and Guido Porto. 2012. "Exports, Export Destinations, and Skills." *American Economic Review* 102 (7): 3406–38.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *Quarterly Journal of Economics* 117 (1): 339–76.

Burstein, Ariel, Javier Cravino, and Jonathan Vogel. 2013. "Importing Skill-Biased Technology." *American Economic Journal: Macroeconomics* 5 (2): 32–71.

Business Daily. 2009. "'The Edge' names top Kenyan software developers." October 2, 12.

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

Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women." *Quarterly Journal of Economics* 131 (2): 633–86.

Card, David, and Stefano DellaVigna. Forthcoming. "What Do Editors Maximize? Evidence from Four Leading Economics Journals." *Review of Economics and Statistics*.

Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *Quarterly Journal of Economics* 128 (3): 967–1015.

Caroli, Eve, and John Van Reenen. 2001. "Skill-Biased Organizational Change? Evidence from a Panel of British and French Establishments." *Quarterly Journal of Economics* 116 (4): 1449–92.

Casaburi, Lorenzo, Michael Kremer, Sendhil Mullainathan, and Ravindra Ramrattan. 2013. "Harnessing ICT to Increase Agricultural Production: Evidence From Kenya." *Unpublished.*

Champion, Sara, Katrina Kosec, and Christopher Stanton. 2012. "The Effects of Internet Access on Labor-Supply Decisions." *Unpublished.*

Chavula, Josiah, Nick Feamster, Antoine Bagula, and Hussein Suleman. 2014. "Quantifying the Effects of Circuitous Routes on the Latency of Intra-Africa Internet Traffic: A Study of Research and Education Networks." Paper presented at AFRICOMM: 6th International Conference on e-Infrastructure and e-Services for Developing Countries, Kampala, Uganda.

Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and Taxation: Theory and Evidence." *American Economic Review* 99 (4): 1145–77.

Clark, Colin. 1940. *The Conditions of Economic Progress.* London: MacMillan.

Clarke, George R. G., and Scott J. Wallsten. 2006. "Has the Internet Increased Trade? Developed and Developing Country Evidence." *Economic Inquiry* 44 (3): 465–84.

CNN. 2009. "Going Optic." http://edition.cnn.com/2009/TECH/07/22/seacom.on/index.html.

Commander, Simon, Rupert Harrison, and Naercio Menezes-Filho. 2011. "ICT and Productivity in Developing Countries: New Firm-Level Evidence from Brazil and India." *Review of Economics and Statistics* 93 (2): 528–41.

Commonwealth Telecommunications Organisation. 2012. *The Socio-Economic Impact of Broadband in sub-Saharan Africa: The Satellite Advantage.* London: Commonwealth Telecommunications Organisation.

Conley, T. G. 1999. "GMM Estimation with Cross Sectional Dependence." *Journal of Econometrics* 92 (1): 1–45.

Crespi, Gustavo, Chiara Criscuolo, and Jonathan Haskel. 2007. "Information Technology, Organisational Change and Productivity Growth: Evidence from UK Firms." CEPR Discussion Paper 6105.

Currie, Janet, and Ann E. Harrison. 1997. "Sharing the Costs: The Impact of Trade Reform on Capital and Labor in Morocco." *Journal of Labor Economics* 15 (3): S44–71.

Czernich, Nina, Oliver Falck, Tobias Kretschmer, and Ludger Woessmann. 2011. "Broadband Infrastructure and Economic Growth." *Economic Journal* 121 (552): 505–32.

Dabla-Norris, Era, Kalpana Kochhar, Nujin Suphaphiphat, Frantisek Ricka, and Evridiki Tsounta. 2015. *Causes and Consequences of Income Inequality: A Global Perspective.* Washington, DC: IMF.

Davis, Donald R., and James Harrigan. 2011. "Good Jobs, Bad Jobs, and Trade Liberalization." *Journal of International Economics* 84 (1): 26–36.

De Loecker, Jan. 2011. "Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity." *Econometrica* 79 (5): 1407–51.

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 M. Cordeiro, Carlos, Hrishikesh Gossain, Roy L. Ashok, and Dharma P. Agrawal. 2003. "The Last Mile: Wireless Technologies for Broadband and Home Networks." Unpublished.

De Stefano, Timothy, Richard Kneller, and Jonathan Timmis. 2014. "The (Fuzzy) Digital Divide: The Effect of Broadband Internet Use on UK Firm Performance." *Unpublished.*

DiNardo, John E., and Jörn-Steffen Pischke. 1997. "The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too?" *Quarterly Journal of Economics* 112 (1): 291–303.

Draca, Mirko, Raffaella Sadun, and John Van Reenen. 2007. "Productivity and ICTs: A Review of the Evidence." In *The Oxford Handbook of Information and Communication Technologies*, edited by Robin Mansell, Chrisanthi Avgerou, Danny Quah, and Roger Silverstone, 100–147. Oxford: Oxford University Press.

Eaton, Jonathan, Marcela Eslava, C.J. Krizan, Maurice Kugler, and James Tybout. 2015. "A Search and Learning Model of Export Dynamics." *Unpublished.*

The Economist 1843. 2012. "What's the Greatest Invention of All Time?" January/February.

Fajgelbaum, Pablo D., and Amit K. Khandelwal. 2016. "Measuring the Unequal Gains from Trade." *Quarterly Journal of Economics* 131 (3): 1113–80.

Feenstra, Robert C., and Gordon H. Hanson. 1996. "Foreign Investment, Outsourcing, and Relative Wages." In *The Political Economy of Trade Policy: Papers in Honor of Jagdish Bhagwati*, edited by Robert C. Feenstra, Gene M. Grossman, and Douglas A. Irwin, 89–128. Cambridge, MA: MIT Press.

Feenstra, Robert C., and Gordon H. Hanson. 1999. "The Impact of Outsourcing and High-Technology Capital on Wages: Estimates for the United States, 1979–1990." *Quarterly Journal of Economics* 114 (3): 907–40.

Feenstra, Robert C., and Gordon H. Hanson. 2003. "Global Production Sharing and Rising Inequality: A Survey of Trade and Wages." In *Handbook of International Trade.* Vol. 1, edited by E. Kwan Choi and James Harrigan, 146–85. Blackwell Publishing Ltd.

Ferraz, Claudio, Frederico Finan, and Dimitri Szerman. 2015. "Procuring Firm Growth: The Effects of Government Purchases on Firm Dynamics." NBER Working Paper 21219.

Forman, Chris, Avi Goldfarb, and Shane Greenstein. 2012. "The Internet and Local Wages: A Puzzle." *American Economic Review* 102 (1): 556–75.

Foster, Andrew D., and Mark R. Rosenzweig. 2007. "Economic Development and the Decline of Agricultural Employment." In *Handbook of Development Economics*, Vol. 4, edited by T. Schultz and John A. Strauss, 3051–83. Amsterdam: Elsevier.

Frazer, Garth. 2013. "Imports, Import Sources and Skill Utilization." Unpublished.

Frías, Judith A., David S. Kaplan, and Eric Verhoogen. 2009. "Exports and Wage Premia: Evidence from Mexican Employer-Employee Data." *Unpublished.*

Gallaugher, John. 2012. *Getting the Most Out of Information Systems: A Manager's Guide.* Creative Commons.

Giorgis, Tamrat G. 2010. "Ethiopia: Global Firms Vie over Country's Link to the World." *Addis Fortune*, January 18. https://addisfortune.net/.

GIZ. 2014. "Connect the BoP." https://www.giz.de/expertise/downloads/giz2013-de-connect-the-bop.pdf.

Goldberg, Pinelopi K. 2015. *Trade and Inequality.* Camberley: Edward Elgar.

Goldberg, Pinelopi K., Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova. 2010. "Imported Intermediate Inputs and Domestic Product Growth: Evidence from India." *Quarterly Journal of Economics* 125 (4): 1727–67.

Goldberg, Pinelopi K., Amit K. Khandelwal, Nina Pavcnik, and Petia Topalova. 2010. "Multiproduct Firms and Product Turnover in the Developing World: Evidence from India." *Review of Economics and Statistics* 92 (4): 1042–49.

Goldberg, Pinelopi K., and Nina Pavcnik. 2007. "Distributional Effects of Globalization in Developing Countries." *Journal of Economic Literature* 45 (1): 39–82.

Goldin, Claudia, and Lawrence F. Katz. 2007. "Long-Run Changes in the Wage Structure: Narrowing, Widening, Polarizing." *Brookings Papers on Economic Activity* 38 (2007–2): 135–65.

Gollin, Douglas, David Lagakos, and Michael E. Waugh. 2014. "The Agricultural Productivity Gap." *Quarterly Journal of Economics* 129 (2): 939–93.

Gollin, Douglas, Stephen Parente, and Richard Rogerson. 2002. "The Role of Agriculture in Development." *American Economic Review* 92 (2): 160–64.

Goos, Maarten, Alan Manning, and Anna Salomons. 2014. "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring." *American Economic Review* 104 (8): 2509–26.

Green, Francis, Andy Dickerson, and Jorge Saba Arbache. 2001. "A Picture of Wage Inequality and the Allocation of Labor through a Period of Trade Liberalization: The Case of Brazil." *World Development* 29 (11): 1923–39.

Hansman, Christopher, Jonas Hjort, Gianmarco León, and Matthieu Teachout. 2018. "Vertical Integration, Supplier Behavior, and Quality Upgrading among Exporters." NBER Working Paper No. 23949.

Hardy, Morgan, and Jamie McCasland. 2015. "Are Small Firms Labor Constrained? Experimental Evidence from Ghana." *Unpublished.*

Harris, Lance. 2012. "Is South Africa the Best Place to Start Your Start-Up?" *ZDNet*, November 21.

Harrison, Ann, and Ana Revenga. 1998. "Labor Markets, Foreign Investment, and Trade Policy Reform." In *Trade Policy Reform: Lessons and Implications*, edited by John Nash and Wendy Takacs, 247–76. Washington, DC: World Bank.

Harrison, Ann, and Gordon Hanson. 1999. "Who Gains from Trade Reform? Some Remaining Puzzles." *Journal of Development Economics* 59 (1): 125–54.

Harrison, Ann, John McLaren, and Margaret McMillan. 2011. "Recent Perspectives on Trade and Inequality." *Annual Review of Economics* 3 (1): 261–89.

Henderson, J. Vernon, Adam Storeygard, and David N. Weil. 2012. "Measuring Economic Growth from Outer Space." *American Economic Review* 102 (2): 994–1028.

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

International Labour Office (ILO). 2012. *International Standard Classification of Occupations: ISCO–08*. Geneva: International Labour Office.

Hjort, Jonas, and Jonas Poulsen. 2019. "The Arrival of Fast Internet and Employment in Africa: Dataset." *American Economic Review.* https://doi.org/10.1275/20161385.

Internet Society. 2013. *Regional Interconnection Strategy for Africa.* Geneva: Internet Society.

ITU. 2000. *Improving IP Connectivity in the Least Developed Countries.* https://www.itu.int/osg/spu/ni/ipdc/study/Improving%20IP%20Connectivity%20in%20the%20Least%20Developed%20Countries1.pdf.

ITU. 2013. *ITU Study on International Internet Connectivity in Sub-Sarahan Africa.* https://www.itu.int/en/ITU-D/Regulatory-Market/Documents/IIC_Africa_Final-en.pdf.

Jack, William, and Tavneet Suri. 2014. "Risk Sharing and Transactions Costs: Evidence from Kenya's Mobile Money Revolution." *American Economic Review* 104 (1): 183–223.

Jensen, Robert. 2007. "The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector." *Quarterly Journal of Economics* 122 (3): 879–924.

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.

Jorgenson, Dale W., Mun S. Ho, and Kevin J. Stiroh. 2008. "A Retrospective Look at the U.S. Productivity Growth Resurgence." *Journal of Economic Perspectives* 22 (1): 3–24.

Katz, Lawrence F., and Robert A. Margo. 2014. "Technical Change and the Relative Demand for Skilled Labor: The United States in Historical Perspective." In *Human Capital in History: The American Record*, edited by Leah Platt Boustan, Carola Frydman, and Robert A. Margo, 15–57. Chicago: University of Chicago Press.

Katz, Lawrence F., and David H. Autor. 1999. "Changes in the Wage Structure and Earnings Inequality." In *Handbook of Labor Economics*. Vol. 3, edited by Orley C. Ashenfelter and David Card, 1463–555. Amsterdam: Elsevier.

Kende, Michael, and Karen Rose. 2015. *Promoting Local Content Hosting to Develop the Internet Ecosystem.* Geneva: Internet Society.

Koren, Miklós, and Martón Csillag. 2016. "Machines and Machinists: Importing Skill-Biased Technology." *Unpublished.*

Krueger, Alan B. 1993. "How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984–1989." *Quarterly Journal of Economics* 108 (1): 33–60.

La Ferrara, Eliana, Alberto Chong, and Suzanne Duryea. 2012. "Soap Operas and Fertility: Evidence from Brazil." *American Economic Journal: Applied Economics* 4 (4): 1–31.

Lagakos, David, and Michael E. Waugh. 2013. "Selection, Agriculture, and Cross-Country Productivity Differences." *American Economic Review* 103 (2): 948–80.

Levinsohn, James. 1999. "Employment Responses to International Liberalization in Chile." *Journal of International Economics* 47 (2): 321–44.

Levinsohn, James, and Amil Petrin. 2003. "Estimating Production Functions Using Inputs to Control for Unobservables." *Review of Economic Studies* 70 (2): 317–41.

Lewis, W. Arthur. 1955. *The Theory of Economic Growth.* London: Allen and Unwin.

Lowe, Matt. 2014. "Rail Revival in Africa? The Impact of Privatization." *Unpublished.*

Macchiavello, Rocco, and Pepita Miquel-Florensa. 2015. "Supply Assurance, Relational Contracts and Vertical Integration: Evidence from Costa Rica's Coffee Chain." *Unpublished.*

Machin, Stephen, and John Van Reenen. 1998. "Technology and Changes in Skill Structure: Evidence from Seven OECD Countries." *Quarterly Journal of Economics* 113 (4): 1215–44.

Magruder, Jeremy R. 2010. "Intergenerational Networks, Unemployment, and Persistent Inequality in South Africa." *American Economic Journal: Applied Economics* 2 (1): 62–85.

Magruder, Jeremy R. 2012. "High Unemployment yet Few Small Firms: The Role of Centralized Bargaining in South Africa." *American Economic Journal: Applied Economics* 4 (3): 138–66.
Mahlknecht, Greg. 2014. "Greg's Cable Map." https://www.cablemap.info/ (accessed November 16, 2018).
Márquez, Gustavo, and Carmen Páges. 1997. "Trade and Employment: Evidence from Latin America and the Caribbean." *Inter-American Development Bank Working Paper 366.*
McKenzie, David, Suresh de Mel, and Christopher Woodruff. 2008. "Returns to Capital: Results from a Randomized Experiment." *Quarterly Journal of Economics* 123: 1329–72.
McKinsey Global Institute. 2013. "Lions go Digital: The Internet's Transformative Potential in Africa." http://www.mckinsey.com/industries/high-tech/our-insights/lions-go-digital-the-internets-transformative-potential-in-africa.
McMillan, Margaret, Dani Rodrik, and Inigo Verduzco-Gallo. 2014. "Globalization, Structural Change, and Productivity Growth, with an Update on Africa." *World Development* 63: 11–32.
McMillan, Margaret, and Kenneth Harttgen. 2014. "Africa's Quiet Revolution." In *The Oxford Handbook of Africa and Economics*, Vol. 2, edited by Célestin Monga and Justin Yifu Lin, 39–61. Oxford: Oxford University Press.
Michaels, Guy, Ashwini Natraj, and John Van Reenen. 2014. "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years." *Review of Economics and Statistics* 96 (1): 60–77.
Michalopoulos, Stelios, and Elias Papaioannou. 2013. "Pre-Colonial Ethnic Institutions and Contemporary African Development." *Econometrica* 81 (1): 113–52.
Michalopoulos, Stelios, and Elias Papaioannou. 2018. "Spatial Patterns of Development: A Meso Approach." *Annual Review of Economics* 10: 383–410.
Mitra, Sandip, Dilip Mookherjee, Maximo Torero, and Sujata Visaria. Forthcoming. "Middleman Margins and Asymmetric Information: An Experiment with Potato Farmers in West Bengal." *Review of Economics and Statistics.*
Moreira, Mauricio M., and Sheila Najberg. 2000. "Trade liberalisation in Brazil: Creating or exporting jobs?" *Journal of Development Studies* 36 (3): 78–99.
Organisation for Economic Co-operation and Development (OECD). 2013. *Measuring the Internet Economy: A Contribution to the Research Agenda. Paris: OECD.*
Organisation for Economic Co-operation and Development (OECD). 2014. *International Cables, Gateways, Backhaul and International Exchange Points.* Paris: OECD.
Oliner, Stephen D., Daniel E. Sichel, and Kevin J. Stiroh. 2007. "Explaining a Productive Decade." *Brookings Papers on Economic Activity* 2007 (1): 81–137.
Olley, G. Steven, and Ariél Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6): 1263–97.
Oxford Business Group. 2016. "Djibouti to utilise geostrategic location to develop ICT connectivity hub." In *The Report: Djibouti 2016.* https://oxfordbusinessgroup.com/overview/emerging-gateway-country-seeks-leverage-its-geostrategic-location-develop-connectivity-hub.
Paunov, Caroline, and Valentina Rollo. 2015. "Overcoming Obstacles: The Internet's Contribution to Firm Development." *Papers and Proceedings of the Annual Bank Conference on Development Economics: Supplement to The World Bank Economic Review* 29: S192–204.
Qiang, Christine Zhen-Wei, and Carlo M. Rossotto. 2009. *Economic Impacts of Broadband.* Washington, D.C.: World Bank.
Raveh, Ohad, and Ariell Reshef. 2016. "Capital Imports Composition, Complementarities, and the Skill Premium in Developing Countries." *Journal of Development Economics* 118: 183–206.
Revenga, Ana. 1997. "Employment and Wage Effects of Trade Liberalization: The Case of Mexican Manufacturing." *Journal of Labor Economics* 15 (3): S20–43.
Rice, Xan. 2013. "Internet Sales Flourish in Nigeria." *Financial Times*, May 14. https://www.ft.com/content/3f455b7e-b1bb-11e2-9315-00144feabdc0.
Rodrik, Dani. 2016. "Premature Deindustrialization." *Journal of Economic Growth* 21 (1): 1–33.
Sarrocco, Claudia, and Tim Kelly. 2000. *Improving IP Connectivity in the Least Developed Countries*. Geneva: International Telecommunication Union.
Scruggs, Gregory. 2015. "How Nairobi Built a Thriving Tech Community." *Citylab*, April 13. https://www.citylab.com/life/2015/04/how-nairobi-built-a-thriving-tech-community/390387/.
Syrquin, Moshe. 1988. "Patterns of Structural Change." In *Handbook of Development Economics*. Vol. 1, edited by Hollis Chenery and T. Srinivasan, Chapter 7, 203–73. Amsterdam: Elsevier.
Syverson, Chad. 2011. "What Determines Productivity?" *Journal of Economic Literature* 49 (2): 326–65.
Topalova, Petia. 2010. "Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India." *American Economic Journal: Applied Economics* 2 (4): 1–41.

Verhoogen, Eric A. 2008. "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector." *Quarterly Journal of Economics* 123 (2): 489–530.

Wei, Shang-Jin, and Yi Wu. 2002. "Globalization and Inequality Without Differences in Data Definition, Legal System and Other Institutions." Unpublished.

Winters, L. Alan, Neil McCulloch, and Andrew McKay. 2004. "Trade Liberalization and Poverty: The Evidence So Far." *Journal of Economic Literature* 42 (1): 72–115.

Wood, Adrian. 1995. "How Trade Hurt Unskilled Workers." *Journal of Economic Perspectives* 9 (3): 57–80.

Wood, Adrian. 1999. "Openness and Wage Inequality in Developing Countries: The Latin American Challenge to East Asian Conventional Wisdom." In *Market Integration, Regionalism and the Global Economy*, edited by Richard Baldwin, Daniel Cohen, Andre Sapir, and Anthony Venables, 153–84. Cambridge, UK: Cambridge University Press.

World Bank. 2006. *Information and Communications for Development 2006: Global Trends and Policies.* Washington, DC: World Bank.

World Bank. 2016. *World Development Report 2016: Digital Dividends.* Washington, DC: World Bank.

World Bank. 2017. *Firm Level Productivity Estimates.* Washington, DC: World Bank.

Young, Alwyn. 2014. "Structural Transformation, the Mismeasurement of Productivity Growth, and the Cost Disease of Services." A*merican Economic Review* 104 (11): 3635–67.