

The local advantage: Corruption, organized crime, and indigenization in the Nigerian oil sector ^{*}

Jonah Rexer[†]

University of Pennsylvania

December 26, 2019

Abstract

Indigenization is a common policy tool to increase local participation and employment in the natural resource sectors of developing countries. Such policies raise tradeoffs between the technological efficiency of multinationals and the comparative advantages of localness. I study a two-decade indigenization drive in Nigeria's oil sector, during which the share of local production grew substantially. Local takeover of operations considerably increases oilfield output and reduces the share of nonproducing assets, despite increasing incidents of mechanical failure. Local firms increase output by mitigating conflict risk: theft of oil and violent attacks by organized criminal-militant groups both fall following local takeover. A simple bargaining model highlights several mechanisms to explain the differences in performance between local and multinational firms: bargaining costs, partial ownership, corruption costs, and income spillovers. I find evidence that connections to local politicians and the army, as well as limited exposure to home-country anti-corruption laws, give local companies an advantage in reducing criminal activity.

JEL Classifications: F2, L24, Q34, Q35

Keywords: foreign direct investment, hydrocarbons, political risk, organized crime, illegal markets, oil conflict.

^{*}For their outstanding feedback and assistance, I thank Camilo Garcia-Jimeno, Guy Grossman, Santosh Anagol, Arthur van Benthem. This work has been supported by the Amy Morse fellowship and by a grant from the Kleinman Center for Energy Policy, both at the University of Pennsylvania

[†]Wharton School of Business. 3620 Locust Walk, Philadelphia, PA 19104. Email: jorexer@wharton.upenn.edu. <https://bepp.wharton.upenn.edu/profile/jorexer/>

1 Introduction

Global experience and a large body of evidence suggest that foreign direct investment increases productivity in developing countries. Multinational affiliates are more productive than locally-owned firms (Aitken and Harrison 1999, Arnold and Javorcik 2009, Guadalupe et al. 2012, Criscuolo and Martin 2009), and can provide poor countries with much-needed technology (Teece 1977, Guadalupe et al. 2012) and managerial skills (Bloom and Reenen 2010, Bloom et al. 2012). The literature on spillovers suggests a mixed albeit mildly positive effect on local firms,¹ as well as competitive effects that force inefficient firms to exit (Alfaro and Chen 2018), and transmission of human capital to local firms through labor markets (Balsvik 2011, Poole 2013).

Yet in many countries, political economy undermines the textbook benefits of FDI. A growing literature suggests that foreign multinationals may be ill-equipped to deal with the corruption, conflict, and expropriation that often accompanies working in difficult markets (Blair et al. 2019, Burger et al. 2015).² Local companies may possess insider knowledge, political connections, and legal flexibility that allow them to protect investments from expropriation or theft, expedite bureaucratic procedures, or obtain preferential treatment. Evidence shows that corruption encourages joint ventures as multinationals seek partners to navigate local politics (Javorcik and Wei 2009).

Nowhere are these considerations more salient than in the natural resource sectors. Often dominated by foreign multinationals, these industries have historically been bedeviled by nationalizations, corruption, and domestic political conflict over the distribution of gains and losses from the sector. Locally, the effects of resource extraction are mixed: recent evidence suggests both positive local employment spillovers (Aragon and Rud 2013, Lippert 2014, Loayza and Rigolini 2016, Cust and Poelhekke 2015) and sometimes substantial negative environmental and social externalities (Aragon and Rud 2011, Berman et al. 2017, Dube and Vargas 2013). In many countries, foreign multinationals are required to enter joint ventures with local partners, or local firms are given preference in bidding on mineral rights. But whether the local firms empowered by these laws are actually more productive – and generate larger spillover benefits – remains an open question.

I study the benefits of localness in the context of the Nigerian petroleum sector. For the oil supermajors who do business there, the sector has been fraught with controversy over environmental damage, political instability, and corruption. From 2000-2009, the oil-producing Niger Delta witnessed an armed uprising in which local militant groups attacked multinational oil infrastructure in

¹ see Alfaro and Chauvin (2020) for a review of this extensive literature

² Although the classic study of the Angolan civil war by Guidolin and Ferrara (2007) shows that conflict can be beneficial to multinational diamond firms.

order to extract individual transfers and greater regional control over oil revenues. In the aftermath, a 10 billion dollar-a-year black market in stolen oil – tapped directly from the vast unsecured onshore pipeline network – has developed. Over the same period, the sector has experienced significant indigenization – with the local share of onshore output growing from less than 3% to nearly 10%.

Using a unique panel dataset of annual output, revenue, theft, oil spills, violence, and ownership for Nigeria's active oilfields, I first ask whether fields operated by local firms are more productive. Since local assets may differ from those owned by international companies in their underlying capacity, I employ a two-way fixed effects (TWFE) difference-in-differences strategy that exploits domestic takeovers of existing assets. I find that local companies substantially increase output and revenue by nearly 69% of the control group, while reducing the share of shut-in (non-producing) fields.

To measure technical efficiency, I use data on the number of oil spills a field experiences due to equipment failure. I show that the local advantage comes despite local fields operating at a technical disadvantage and incurring greater losses due to spillage. Oil spills – and resultant output losses – are a reasonable measure of technical efficiency, capturing underlying quality and maintenance of physical capital, technology, safety standards, and management.

How do local firms increase output while simultaneously suffering greater inefficiencies? I find that incidents of oil theft, piracy, and militant attacks experience a prolonged and gradually increasing reduction upon the takeover of a local field operator. On average, a locally operated field experiences 3.3 fewer theft incidents per year on average, roughly 30% of the control group mean, and 1.5 fewer militant-related deaths. These effects are primarily – although not exclusively – driven by takeovers by private firms rather than the state oil company.

Next, I study the benefits of localization by asset location. Offshore extraction requires greater technological sophistication and larger upfront investments, but is less exposed to criminality. This implies local companies advantage in production and theft should be concentrated onshore, while their efficiency disadvantage should be relatively greater offshore. Indeed, I find that the local advantage in output, theft, and militancy is driven by onshore fields. In contrast, the local efficiency disadvantage in oil spills is much larger in offshore assets—exactly where high technical requirements suggest that the efficiency gap is likely to be greatest.

Using a simple model of bargaining between a firm and an organized criminal enterprise, I highlight several sources of advantage that could allow local firms to mitigate criminality and subsequently outperform their more efficient international peers. In the model, firms bargain with gangsters of unknown type, setting output quantities and then offering bribes to protect output from theft. Gangs may either accept these bribes or reject and steal a fixed amount. Two sources of black market

inefficiency give rise to the possibility of Coasian side-payments: *i*) that gangsters incur costs of theft, so that the amount lost by the firm is greater than criminal profits, and *ii*) that oil theft leads to oil spillage, a loss internalized by the firm.

I study several frictions that can shrink the bargaining window: *i*) firms face bargaining costs of dealing with gangsters via local middlemen or subcontractors, *ii*) bribery penalties – such as those from anticorruption laws – impose costs on firms undertaking illicit payments, and *iii*) joint ownership, under which the losses to the operating firm from stolen oil are less than the revenue gained by gangsters. The costs of protection enter the firm’s problem as fixed, so firm production decisions on the intensive margin are unaffected by these forces, but extensive margin choices may be. The model yields a rich set of predictions on how firm production decisions and levels of oil theft will respond to variations in these frictions. The key implication of the model is that if local takeovers reduce bargaining frictions and exposure to corruption penalties, or increases ownership share, then theft falls, more assets are operated, and output rises.

To test whether reduced bargaining frictions drive the benefits of local ownership, I study political connections in the sample of oil companies. I first identify politically-connected firms in data scraped from an international business database and company websites on the identities and biographies of board members, shareholders, and senior management. Using differences-in-differences I show that takeovers by politically connected firms produce benefits similar in magnitude to those of local firms. These effects are greatest for connections to the political figures most intimately involved in the local black market – local politicians and the security forces. In contrast, connections to “technocratic” regulatory agencies and the state oil company have limited effects on theft. Local firms are much more likely to have political connections, and the effect of localness on theft becomes insignificant only when connections to local politicians and security forces are controlled for. This suggests that part of the difference in performance between locals and multinationals is explained by political connections that facilitate bargaining with black market participants.

To investigate the effect of bribery penalties, I compare outcomes across multinationals with differing exposure to home-country foreign-bribery statutes. The motivation behind this test is to remain agnostic on the quality of Nigerian anti-corruption statutes – which are often described as weakly enforced and politically manipulated (Aigbovo and Atsegbua 2013) – and to test the implication that if penalties matter to theft, they should also account for differences across multinationals. Again using TWFE models and variation across companies in year of law passage, I find that having an operator covered by a foreign corruption law increases theft on a field by 6.7 incidents per year. Since we observe this effect within multinationals, it is therefore at least plausible – although not di-

rectly testable – that part of the local advantage stems from the weakness of Nigerian anticorruption law.

I also test whether changes to the opportunity cost of crime drive the results. If increased output and greater spillovers due to local ownership improve employment outcomes for young men in local communities, then labor costs in the black market should also rise, reducing theft by channels outside the bargaining process. To test this hypothesis, I match data on employment and consumption from three rounds of a national panel survey to the ownership status of nearby oilfields. Using an individual-level TWFE specification, I find no evidence that local operators have beneficial spillover effects with regards to employment levels or consumption, including among young men – the relevant pool of potential criminal labor.³

The model predicts substantial heterogeneity in the response of theft to local ownership across three dimensions: the local black market structure, costs of theft, and oil prices. Rising costs generally shrink the magnitude of effects, since when costs are high, criminal profits are low and a deal is easy to reach, so firms are less sensitive to bargaining frictions. I find this is true for costs as proxied by the distance to the nearest militant camp. Rising oil prices increase reservation values for both parties, but more so for firms as bargaining costs fall, reducing the likelihood of theft. This implies a negative interaction between localization and prices, for which I also find evidence. Finally, I provide suggestive evidence that theft and violence fell the most where there are numerous militarily weak local gangs. I argue that militarily strong criminal groups – typically those led by powerful ex-militants – can use violence to impose costs that dwarf bargaining frictions. In other words, bargaining costs only matter among weaker groups who are on the margin of being bribed.

All results are robust to controlling for interactions between year dummies and time-invariant field-level covariates. Event-study models suggest that differential pre-trends in outcomes of interest are not driving the takeover effects. I also estimate a re-weighted event-study using cohort composition to adjust for heterogeneous effects (Abraham and Sun 2018). In addition, I test the robustness of the results to differential effects of oil price trends, measurement error in output, correlation with the Niger Delta conflict, and to the distinction between public and private local firms. I also calculate standard errors using randomization inference and the wild cluster bootstrap. Lastly, I run diagnostic tests from Goodman-Bacon (2019) and de Chaisemartin and D’Haultfoeuille (2019) on the differences-in-differences specifications, which indicate that the TWFE estimates appear reliable even in the presence of substantial treatment effect heterogeneity.

³ However, I do find some evidence of an aggregate shift toward self-employment, driven primarily by middle-aged women.

The results have implications for several literatures that span political science and economics. While the results do not undercut the substantial literature on the productivity advantages of multinationals and the spillover benefits of foreign direct investment in developing countries (reviewed in Harrison and Rodriguez-Clare 2010 and Alfaro and Chauvin 2020), they do introduce another perspective to a seemingly settled question. The vast majority of this literature studies manufacturing or service firms in environments that, while corrupt, are relatively stable. I show that resource sectors in conflict-affected countries have dynamics that may be very different; in these cases local advantage can massively outweigh the productivity gains from foreign investment. The policy implication is that indigenization in troubled extractive sectors can be justified on efficiency grounds.

The results also relate to extensive work on firms and politics. It is well known that in corrupt environments, political connections are valuable to firms (Fisman 2001, Faccio 2006, Li et al. 2008, Khwaja and Mian 2005, Akcigit et al. 2018). However, this literature typically emphasizes the negative equilibrium effects of political favoritism: inefficient firms are protected from competitive pressures. I show that in a context in which local government is closely linked to organized crime, political connections – and the corruption they engender – are actually welfare-improving; they substitute inefficient black markets for legitimate production and incentivize greater output on the part of firms. The results demonstrate a new mechanism by which political connections matter for firm outcomes – they protect against criminal activity by lowering the costs of bargaining with organized crime.

Lastly, the results join a growing empirical literature that uses microdata to unpack the local resource curse. This literature has looked at the negative spillover effects of natural resource booms on violent conflict (Berman et al. 2017, de la Sierra 2019, Dube and Vargas 2013, Fetzer and Kyburz 2018, Nwokolo 2018), social unrest (Sexton 2019, Christensen 2019), politics (Kyburz 2018, Fetzer and Kyburz 2018), and the environment (Aragon and Rud 2011, Sexton 2019).⁴ This work is one of the few to analyze the black markets and organized crime which are central to resource curse dynamics (Couttenier et al. 2017, Buonanno et al. 2015). Unlike previous work, I also analyze firms as strategic participants in the resource curse nexus, whose optimal strategies may vary by ownership type. I am the first to show that local ownership can mitigate some of the worst symptoms of the resource curse.

2 The Nigerian oil sector

Nigeria is the world's 11th largest oil-producer, and the largest in Africa. Rich deposits of Bonny Light are located onshore and in the shallow water of the Niger Delta, a geopolitical region in the

⁴ For reviews of this literature, see Cust and Poelhekke (2015) and Aragon et al. (2015)

far-southern tip of Nigeria that forms where the mouth of the Niger River meets the Gulf of Guinea. The Niger Delta comprises both coastal and inland portions of nine Southern states,⁵ home to 22% of Nigeria's population (NBS 2017), and populated by numerous ethnic minority groups, the largest of which are the Ijaw. Since oil discovery in 1956, the sector has historically been dominated by oil supermajors Shell, Chevron, ExxonMobil, Total, and Eni (Agip). In 2004, these multinational companies produced 93.5% of Nigeria's 2.49 million barrels per day. In that year, the sector was valued at 45.8 billion USD in 2019 dollars, accounting for 98% of Nigeria's export earnings.

Oil multinationals operate in a complex environment. All foreign companies operate profit-sharing agreements with the state-run oil company, the Nigerian National Petroleum Company (NNPC). Ownership contracts include joint ventures involving government and one or many multinationals, production sharing contracts, and fee-for-service contracts. Shares in new or expiring oil blocks are awarded by the state in a competitive bid process. This leads to variation in the share of profits claimed by the operator of a given oilfield. Figure A1 displays a histogram of operator shares for all producing oil blocks as of 2016, which range from 0 to full ownership, with a mean stake of 52%.

Nigeria's oil sector is also a byword for corruption. In 2012, one estimate claimed that the Nigerian government had lost nearly 400 billion dollars in oil income due to corruption since independence.⁶ In the most recent of a long history of corruption scandals, an Italian court is considering charges against Shell and Eni for their participation in bribing government officials to the tune of 1.1 billion USD to receive improved terms on an oil prospecting lease.⁷ Multinationals in Nigeria must contend with the added costs of corruption, which expose them to legal liabilities in their home countries.

Most importantly, companies must also interact with the indigenes of the Niger Delta communities in which they operate. This region is best described as the prototypical example of the local resource curse – a dizzying constellation of violent armed groups interact with oil companies, local and federal government, and each other in an organized chaos that blurs the line between civil war and organized crime (Obi and Rustad 2011, Watts 2007). Local politicians – particularly state governors – are notorious for corruption and the promotion of electoral violence and fraud (Watts 2007). Oil spills are common, affecting soil, fisheries, and drinking water, and increasing infant mortality (Bruederle and Hodler 2019). Despite its oil wealth and disproportionate federal budget allocations, state-level poverty rates in the region range from 39-64%.⁸

⁵ Abia, Bayelsa, Delta, Rivers, Akwa Ibom, Imo, Ondo, Edo, and Cross River states.

⁶ For a survey of current corruption issues in Nigeria, see <https://carnegieendowment.org/2018/07/17/new-taxonomy-for-corruption-in-nigeria-pub-76811>

⁷ <https://www.bbc.com/news/business-46336733>

⁸ Still, the Delta compares favorably to other parts of Nigeria in this respect, with only one state ranking in the top-20 in Nigeria.

Between 2000-2009, The Niger Delta Crisis saw the emergence of well-armed militant groups from longstanding criminal, political, and ethnic militia groups in the Delta (Watts 2007, Asuni 2009). Militants declared war on the federal government and oil companies, destroying oil facilities in an attempt to obtain concessions for themselves and the region. This decade of civil conflict ended in amnesty to nearly 25,000 combatants, as well as lucrative “pipeline surveillance contracts” awarded to major militant commanders – payments which amounted to private transfers to top militants (SDN 2019c). Rexer and Hvinden (2019) show that the amnesty process led to the emergence of a thriving black market in stolen oil, comprised both of mid-level ex-militants who did not secure contracts and more recent entrants. Figure A2 charts the evolution of the black market by plotting the monthly number of incidents of pipeline sabotage.⁹ Oil spills due to theft are declining in the months prior to the amnesty, but then rise steadily afterward. Oil spills due to operational failure, in contrast, are consistently declining over the whole period. The value of losses to Shell, Agip, and Chevron alone was nearly 10 billion dollars between 2009-2011 (NEITI 2011).

In this two-tiered market, smaller downstream entrepreneurs refine about 25% of the stolen crude locally for sale to the domestic market, while larger criminal syndicates typically export the remainder (SDN 2019a, SDN 2019b). The Niger Delta’s network of over 5,000 kilometers of unguarded onshore pipeline makes assets extremely vulnerable and theft a highly profitable activity for black market entrepreneurs. Protection rackets are common: oil companies must surreptitiously bargain with gangsters and local communities in order to prevent theft on their assets. Payments to local communities – which range from direct transfers and contracts for local chiefs and militant groups to community-wide development projects – are viewed as a cost of doing business. In this market, local politicians and security forces play an important role. Many militant groups have historically been supported by political patrons (Asuni 2009), while local security forces facilitate the smooth functioning of the black market through bribes for protection, in many cases even selling rights to operate lucrative upstream tapping operations (SDN 2019a). Connections to these entities would therefore prove useful for oil companies attempting to deter theft.

Against this backdrop, two related trends have emerged in the sector. In response to challenging onshore conditions, multinationals have opted to reallocate resources to the shallow and increasingly deepwater reserves of the Gulf of Guinea. Offshore assets are costly to reach for oil thieves and militants, though they entail much larger fixed and operational costs for firms. As Figure 1 (Panel A) demonstrates, between 2002 and 2015, the share of Nigerian oil produced from onshore fields fell by

⁹ I argue in Section 3 that this is a reasonable proxy for the number of theft incidents and therefore the size of the black market.

half, from 60% to just above 30%. This trend suggests that theft and crime impose significant constraints to onshore operations – firms will undertake costly and irreversible investments and abandon producing fields to avoid it.

FIGURE 1 HERE

At the same time, the sector is becoming increasingly Nigerian. According to Figure 1 (Panel B), the share of national oil output produced by independent private Nigerian oil companies has steadily risen over the past decade. In 2004, this fraction was 3.5%, while by 2015 it had risen to 9.9%. Over the same period, the number of independent local firms operating an oilfield rose from 8 to 30, while the number of fields operated by local companies rose from 9 to 71. In Figure A3, I plot local participation in fields (Panel A) and output (Panel B) by asset type over time. The data show that this growth in local participation is concentrated primarily in onshore assets. The multinational divestment from onshore and move to offshore has created space for local firms to enter the onshore market.

At the same time, this growth has been in part aided by the 2010 Nigerian Local Content Act, a law that enshrined preference for local firms in bidding on new oil blocks. The dotted vertical line in Figure A3 demonstrates that the timing of the law correlates with growth in local onshore participation. Variation in local participation is driven both by multinational divestments and the preferential awarding of new and expiring oil blocks. These are often what are classified as “marginal fields” by the Nigerian government – a category of small or underexploited fields reserved almost exclusively for local companies.

3 Data and summary statistics

3.1 Data description

Oil production and infrastructure data: Information on 279 active Nigerian oilfields forms the core of the data. These field-level data come from annual public reports from the NNPC, augmented with confidential data from the Department of Petroleum Resources (DPR)¹⁰ for years in which NNPC data is unavailable. Between these two sources, I observe the intensive and extensive margin of oil production for each oilfield from 1998-2016.¹¹ Because of uneven coverage, some fields are missing in certain years after the field first appears in the data. I assign output in these field-years to missing, while coding output as zero only when it is explicitly indicated as such in a DPR or NNPC

¹⁰ The DPR is Nigeria’s primary petroleum sector regulatory body.

¹¹ Unfortunately, disaggregated data are unavailable for 2009.

source. A “shut-in” field is defined as a field that is nonproducing in a given time period.

The DPR-NNPC dataset also includes information on the firm operating each field in each year. I code local participation as a dummy that equals one if a local firm is listed as the field operator. There are a few drawbacks to this data: first, there is no detailed existing panel of field ownership – ownership stakes are only observed in 2016 from DPR annual reports. Using the operatorship measure overlooks cases in which local firms are non-operating shareholders, which may also be important. This represents a strict treatment criteria that is likely to bias our results toward zero. Secondly, the DPR-NNPC data contain 249 field-years in which a field appears under multiple operators. I assign these fields to the treatment group if any of the operators are local. To allay concerns about double-counting, I also check that results are robust to excluding these observations.¹²

From DPR I also get field-level time-invariant covariates: the number of wells (field size), date of completion of the first well (field age), and the depth of the deepest well. Finally, I use infrastructure maps to obtain centroid locations for the fields in the DPR-NNPC data, which are then used to link fields to information on oil theft, militancy, piracy, and various control variables. The fields are mapped in Figure 2, with the color of the point indicating the year in which the observation was treated. Over the sample period, there are 71 ever-treated fields and 208 never-treated.

Oil spill and theft data: Data on oil theft comes from the Nigerian Oil Spill Detection and Response Agency (NOSDRA), a division of the Federal Ministry of the Environment. NOSDRA data is taken from the Oil Spill Monitor (OSM), a comprehensive database of all 11,587 reported oil spills from 2006-2017. For each oil spill, NOSDRA investigates and files a Joint Investigative Report (JIV), verified by local communities, the oil company, and the DPR. For each spill, I observe the location and cause of the spill, as well as a text description. For those without coordinates, I georeference based on site description in the JIV, resulting in 11,145 spills with coordinate information.

68.45 % of all oil spills are classified as being caused by “sabotage.” I take this to be my sample of oil theft incidents, since sabotage is a reliable indicator of illegal oil tapping.¹³ For each field, I define theft as the sum of all sabotage incidents that occur annually within 15 km of the centroid of the field. To measure the technical efficiency of oil production, I use all field-level spills that are not due to sabotage. In the OSM, the majority (65.3%) of these non-sabotage incidents are caused by “equipment failure” and “corrosion.” They are thus a reasonable measure for losses incurred by oil companies during the normal course of business that can be controlled by the firm directly.

¹² Treatment is coded as a staggered adoption, so for years in which production is missing, operatorship is assumed to be the same as in the previous year.

¹³ See Rexer and Hvinden (2019) for a discussion about measuring oil theft.

Conflict outcomes: I also estimate the extent to which local ownership affects militant activity. To do this, I use data from the Armed Conflict Location and Event Dataset (ACLED) from 1998-2016. To measure militant activity, I use all conflict events in which at least one participant is a rebel or political militia group. This captures attacks on the oil sector perpetrated by organized armed groups. For each field, I aggregate the sum of annual attacks and fatalities due to militant activity within 15 kilometers of the field centroid.

Boardmembers, managers, and shareholders data: For each of the 40 firms – foreign and domestic – that ever appear as operators in the NNPC-DPR data, I attempt to obtain data on the identities of boardmembers, managers, and shareholders. I first use the Bureau van Dijk Orbis global company database, which contains information on name, position, and demographics of boardmembers, managers, and shareholders for reporting companies. I find 29 of the 42 oil companies in the Orbis data, obtaining personnel information for 451 individuals. Since the Orbis data is incomplete both in its coverage of firms and reporting for a given firm, I augment this data from two sources. Firstly, I scrape company websites for all information on boardmembers and senior management. In this process, I find basic personnel data for 602 individuals across 39 firms, 10 of which are uncovered by Orbis. Lastly, I use the Oil and Gas Map of Nigeria, an “independent initiative to monitor the Oil and Gas industry of Nigeria,” for additional information on 376 shareholders across 73 Nigerian oil firms.¹⁴ In total, I obtain some personnel information on 1,037 unique individuals in all 40 firms.

I then scrape biographies on these individuals from Wikipedia, Google, and individual company websites; in total, I obtain biographical information for 431 individuals over 37 companies.¹⁵ I use this biographical information to code several field-level dummy variables. In particular, I identify fields in the data in which the operator employs or is owned by an individual that has ever served at any level of Nigerian government. I also refine this by considering connections to technocratic regulatory agencies (DPR and NNPC), elected politicians, politicians in the state in which the field is located, and members of the army and police. The data have several drawbacks: firstly, they are incomplete and the extent of incompleteness is unknown. For this reason, I use the relatively inclusive criteria of any connection to minimize the dependence on the number of individuals that were able to be identified in the scraping procedure. Most importantly, the data do not contain information on tenure or starting dates. It is therefore impossible to identify whether a company-specific connection is actually active at a given date. However, I still obtain field-time variation in these variables because of ownership changes at the field level. Thus, I estimate the effect of being operated by a firm that

¹⁴ However, many of these firms do not show up in the DPR-NNPC data because they have not yet started producing.

¹⁵ The three missing companies cover 166 field-year observations, or roughly 3% of the data.

contains any personnel ever satisfying some criterion.

Data on militant groups: Finally, I use data on militant camps, described in detail in Hvinden and Rexer (2019). These data – collected by the author from local NGOs and augmented by data from Blair and Imai (2013) – measure the location, commander, militant group affiliation, and amnesty status of 69 militant camps, as of roughly 2009. These camps are relevant to understanding oil theft activity, since much of the post-2009 spike in black market activity is concentrated in nearby areas (as shown in Hvinden and Rexer 2019), suggesting that they are strategic sites for oil theft activities. This is supported by the observation that ex-militants are important players in the post-conflict bunkering economy, with many transitioning from rebel activity to organized crime (SDN 2019c). These ex-militants typically operate in their previous geographical spheres of influence, either by directly participating in the bunkering economy or providing protection for those who do.

I use these data to measure several variables of interest. Firstly, if we accept that these camps represent epicenters of zones of militant influence, then fields very near to militant camps are likely to be low-cost targets for ex-militant-run (or sanctioned) oil theft syndicates. As such, I use distance between a field and its nearest camp to proxy for theft costs. Using the data on group affiliation of each camp, I am also able code the number of groups surrounding each oilfield within a certain radius – a measure of the competitiveness of the black market. Lastly, I take a measure of group military strength derived and validated in Hvinden and Rexer (2019) which identifies the strongest camps based on the number of local allies along the pipeline network.

Sample construction: The various data sources have different time series and degrees of completeness. To harmonize the results, I take as the sample 2006-2016, for which panel data on militant attacks, piracy, theft, and oil output is all available at the field-level. Within this period, oil production data is missing for some fields in each year because of incomplete coverage in the DPR-NNPC reports.¹⁶ Therefore, while the estimation sample for all non-production outcomes is 3,069 field-years, the sample for regressions in which production is the outcome falls to only 2,310 field-years.¹⁷

3.2 Summary statistics

Figure 2 maps the pipelines and oil fields of the Niger Delta in relation to Nigeria’s southern coastline. The points, representing the geographic center of each oilfield, are colored to indicate their treatment cohort. The 208 never-treated fields, in white, are always operated by a multinational firm.

¹⁶ I do not observe the cause of missingness. I therefore assume this data is missing at random. Table A1 shows that outcomes and covariates are very similar across these samples, supporting this assumption.

¹⁷ I choose not to restrict the sample for all estimation in order to make full use of available data for non-production outcomes.

Untreated fields are clustered in the tidal mangroves of Delta, Bayelsa, and Rivers states – the heart of Nigeria’s oil sector – as well as in the shallow waters off Akwa Ibom state. The red points indicate the 71 locally-owned fields and their takeover dates. Indigenized fields are clustered primarily in the inland Niger Delta – an underexplored reserve – with a cluster of recently-divested fields in coastal Rivers state and a handful of offshore assets. Localized fields are more likely to be in peripheral Niger Delta states like Ondo, Imo, and Edo, and are disproportionately located in central Delta state.

FIGURE 2 HERE

Summary statistics are presented in Table 1. In the top panel, I compare time-invariant field-level characteristics between the ever-treated and untreated fields. Treatment here is defined as ever having a local operator listed in the DPR-NNPC data. Given Figure 2, there are ample reasons to believe that localized and multinational fields are likely to be different. Despite these concerns, fields are not significantly different in their distance from the coast, the Niger River, the state capital, or from militant camps. They are of a similar age,¹⁸ on average initiated in 1972-73. They have similar maximum well depth, indicating that they do not belong to substantially different geological types.

However, treated fields do differ in a few important ways. Firstly, they have a greater latitude, indicating a more northerly position. This is driven by two dynamics – firstly, new blocks and marginal fields are more likely to be in the inland Niger Delta. Secondly, the trend of offshoring by multinationals, as discussed in Section 2, implies that treated fields are likely to be onshore. This is reflected in the geographic distribution of treated fields in Figure 2. 84% of ever-treated fields are onshore, while only 71% of multinational ones are, a difference that is significant at 5%. Treated fields are also slightly smaller, with 6.5 fewer wells per field though this is only significant at 10%. This fits with the prior that multinationals have not yet divested of their largest onshore holdings, and that locals are overrepresented in smaller marginal fields.

TABLE 1 HERE

I also compare differences in outcomes and other variables of interest for the analysis in the bottom panel of Table 1. Given the staggered event date, these comparisons use all the data and therefore mix before and after periods for the treated group. Treated fields experience less asset sabotage and theft, but more militant violence. There are no differences in rates of shut-in, but annual production

¹⁸ Defined as the date of completion of the first oil well.

is on average 1.1 million barrels lower, a 29% difference. This is likely driven by smaller field sizes. In order to determine which of these relationships are causal, and control for the level differences across covariates, we move to a staggered-adoption differences-in-differences approach in Section 4.

Treated and untreated fields also differ in the composition of boardmembers, managers, and shareholders of the operating company. Untreated fields are much more likely to be operated by firms led by people of the same ethnicity as the communities in which assets are located. At the same time, rates of political connection are similar, at around 30% of all fields, and 25% for elected politicians. However, multinationals are much more likely to lean on connections to technical agencies, such as the Department of Petroleum Resources, while locally-operated fields are connected to the security forces and politicians from the state in which the field is located.

Treatment and control fields also differ in their annual time series of black market activity. The growth of the black market in Figure A2 is highly heterogeneous. Figure A4 plots mean annual field-level sabotage incidents over time separately for ever-treated and un-treated fields. The two series start at similar levels in 2006 but diverge quickly. The plot suggests that the bulk of the aggregate spike in theft is driven by fields that had no local participation over the decade. In contrast, fields that experienced indigenization see only a mild rise in theft on average, followed by a leveling. Of course, the suggestive correlation of the descriptive data may or may not correspond to a causal effect of localization. For that, we turn to the differences-in-differences strategy in Section 4.

4 Empirical strategy

To test whether local firms affect outcomes at the field-level, I estimate the following differences-in-differences (DD) regression for field i at time t :

$$y_{it} = \alpha + \psi local_{it} + \delta_t + \zeta_i + X'_{it}\beta + \varepsilon_{it}$$

Where y_{it} is the outcome of interest, $local_{it}$ indicates that the field has a local operator, and ψ measures the average effect of localization. Fixed effects for year δ_t and field ζ_i complete the TWFE specification of the DD model, while X_{it} includes an additional vector of time-invariant covariates interacted with year dummies. Throughout, I use a parsimonious set of controls that includes the distance from the field to the state capital, the nearest river, and the coast. Standard errors are clustered at the field level. The key outcomes of interest output, shut-ins, and non-theft oil spills, as well as measures of criminality – oil theft, violence, and piracy.

Throughout, $local_{it}$ is measured as a dummy variable equaling one if the operator of an oilfield is an indigenous Nigerian firm. Variation in this variable comes from several sources. Firstly, cross-sectional variation in ownership exists in any given year; early-entrant Nigerian firms like Pan Ocean and Dubri have been operating since 1998 and account for much of the cross-sectional variation before 2010. Secondly, variation is driven by the entry of new Nigerian firms over time, a process which takes two forms: *i*) awarding of new blocks and *ii*) divestments of existing ones. Each of these presents a source of endogeneity: early Nigerian operators may have obtained smaller fields, new fields coming online may be initially more productive, and fields put up for sale could be languishing. The direction of the bias, much less its magnitude, is unclear.

In a TWFE specification, variation in $local_{it}$ is driven by changes in ownership within a field over time, holding common time-trends fixed. This means that fixed differences in the age, size, or productivity of fields allocated to different types of firms are controlled for. Only trends in output correlated with ownership changes should contaminate the results. Local takeovers might occur when oil prices are low, or following a deterioration of output and theft trends on a given asset. Localization could also be spatially and temporally correlated with specific policy changes – such as the amnesty – that influence theft in other ways. As a standard omnibus test for the presence of parallel pre-trends, I estimate the event-study specification

$$y_{it} = \alpha + \sum_{\tau=-T}^T \psi_{\tau} L_{it}^{\tau} + \delta_t + \zeta_i + X'_{it} \beta + \varepsilon_{it}$$

Where $L_{it}^{\tau} = 1(t - t_i = \tau) * local_i$, where $local_i$ indicates that i ever has a local operator, t_i is the year of treatment for unit i and τ is the year in event-time. The event-study specification also has the benefit of dealing with downweighted early-treated cohorts and bias introduced by time-varying treatment effects (Goodman-Bacon 2019). In addition to the event-study, I include robustness tests controlling for amnesty and other policy changes, controlling for differential responses to prices by treated status, addressing measurement error in output by focusing on single-operator fields, including wide array of interacted controls, and calculating standard errors with randomization inference. I investigate the role of heterogeneity and the implicit weighting of the TWFE specification using results from de Chaisemartin and D'Haultfoeuille (2019) and Goodman-Bacon (2019), address bias in the TWFE induced by using early-treated units as controls, and estimate event-study regressions that are robust to cohort-specific heterogeneous effects (Abraham and Sun 2018).

5 Main results

The main results of the TWFE models are in Table 2. In Panel A, I estimate the model for shut-in probability in (1)-(2), output in millions of barrels (3)-(4) and malfunction (5)-(6). For each outcome, I estimate the TWFE model with and without controls. Panel B contains results for crime and violence outcomes: oil theft incidents in (1)-(2), militant deaths in (3)-(4) and piracy attacks in (5)-(6).

5.1 Technical

Table 2 Panel A tests the production advantage of local firms. In column (1)-(2), I find that a local takeover reduces the shut-in probability of a field by 18.8-19.2 percentage points. Local firms therefore revive moribund fields when they assume operatorship. Output also rises by roughly 2.5 million barrels on average per field annually. This is a very large effect size, at roughly 80% of the control group mean. Furthermore, it is not driven exclusively by decreasing shut-ins. To see this, I estimate the output TWFE regression with controls on the sample of producing assets in Table A5, column (5). The coefficient falls to 1.74, or 56% of the control group mean. Therefore, the main effect on output operates not only on the extensive margin, but also requires increases on the intensive margin. This suggests that the benefits of localness outweigh the cost disadvantages. Lastly in columns (5)-(6), I estimate the effect of local ownership on equipment malfunctions that result in oil spillage. Local fields experience roughly 1-1.9 more spills annually, or roughly 13-25% of the sample mean. Across all specifications in Panel A, the inclusion of controls slightly weakens the effect, but not substantially.

I interpret the effects in (5)-(6) as evidence that local operators are less efficient in their operations, resulting in greater rates of equipment failure and oil spillage. These spills could be driven by lower-quality physical capital, human capital, or management practices and standards. I do not have the data to disaggregate these effects, but rather subsume all of these under efficiency differences.

The increase in output translates to substantially higher average revenue upon local takeover, as seen in Table A2. Here, I estimate the effect on revenue in columns (1)-(4) and log revenue in columns (5)-(8). I control for differential price effects, field-specific time trends, and controls interacted with year dummies. The increase in output corresponds to 167 million dollars in revenue per year in column (4), or 69% of the control group mean. Such huge effect sizes suggest a very important advantage for local firms operating in Nigeria.

TABLE 2 HERE

5.2 Criminality

What are the benefits of localness driving this gap in performance? In Table 2 Panel B, I explore differences in criminality and militancy, the key risks faced by oil companies in the Niger Delta. Fields taken over by local companies do better across a range of conflict indicators. They experience 2.8-3.3 fewer theft incidents annually, or 25.5-28.9% of the control group mean, significant at 5%. Locally-operated fields experience lower levels of violence: annual militant attacks perpetrated within an 15 km radius of the field centroid fall by 1.4-1.5 attacks, a result insensitive to controls and significant at 1% (columns 3-4). This effect is roughly 79.2-85.9% of the control group mean. Lastly, local fields reduce piracy on their assets by 0.072-0.089 annual attacks, although this effect is not significant (columns 5-6). The local advantage seems to be comprised primarily of the ability to manage the political risk of operating in the Niger Delta, as evinced by lower levels of theft and violence.

Are these differences large enough to be driving the output effect? In the model presented in Section 6, theft is related to output in two ways. On the intensive margin, there is the somewhat mechanical effect of theft reducing output as measured and received by the firm – which is what we observe – because oil is stolen or lost. On the extensive margin, however, there is an additional effect on shut-ins that arises from the increased fixed costs imposed by theft. Militant attacks can also affect the both margins, as demonstrated by the history of the Niger Delta conflict, in which output was strongly negatively correlated with periods of intense fighting (Rexer and Hvinden 2019).

One way of testing this is to split the sample by aggregate theft levels. If theft has not, historically, been an issue for a given set of oilfields, these fields should not experience any output impact of localization under the hypothesis that localization affects output primarily through theft. I classify the sample into the 50 fields that never experience theft in the sample period and the 229 that do. I find that the coefficient for the latter group who experience theft is 2.38, significant at 1%, while for the former it is 0.87 and insignificant. While we cannot conclusively rule out other mechanisms by which localization increases output, we can say with some certainty that reduced theft plays an important role.

5.3 Additional tests

5.3.1 Parallel trends

I test for divergent pre-trends using a standard event-study model, described in Section 4. In Figure 3 I estimate the model for output (Panel A) and malfunctions (Panel B). For each regression, I omit $\tau = -1$ as the pre-event reference year, and estimate the specification with fully interacted

field-level controls. Overall, the pre-trends for both outcomes appear relatively parallel across treated and control fields. None of the coefficients for the ψ_τ for $\tau < 0$ are ever significant, whereas nearly all of those for $\tau > 0$ are positive and significant. For malfunctions I also find that the coefficients for $\tau < 0$ are never significant and are typically very near zero. The post-local takeover coefficients are positive although imprecisely estimated. They increase more or less steadily over the years, while the output effects in Panel A display an initial boost followed by a leveling.

FIGURE 3 HERE

In Figure 4 I estimate and plot the same regressions for theft and militant attacks. For $\tau < 0$, the coefficients are insignificant for both outcomes. The post-event coefficients for oil theft incidents are negative, imprecise, significant in 8 out of the 17 post-event periods. They display an initial drop, followed by a long and sustained decline in oil theft for treated oilfields over time, relative to the difference at year zero. Militancy outcomes take substantially longer to improve, with effectively no impact until 5 years after a localization. 6 of the subsequent coefficients are negative and significant. Despite the noisiness of some of the estimates, these findings generally support that identification assumptions necessary for DD to deliver a causal effect are likely to be satisfied.

FIGURE 4 HERE

5.3.2 Asset type

As a falsification test that the local advantage is driven by reductions in theft which trade off against lower efficiency among local firms, I re-estimate the main regressions by asset type in Tables 3 and Tables 4. Offshore assets have higher technology requirements and equipment costs but are less susceptible to theft and crime due to distance from gangsters and the sophistication required to tap them. Onshore assets, with their unprotected pipelines, are highly susceptible to low-cost oil theft but also comparatively easy to operate for the firm.

We should therefore expect to see that reductions in theft are concentrated in onshore assets, which are operate in high-political risk environment and therefore benefit most from local takeover. If theft primarily drives the output advantage, then in turn we should expect that this advantage should be driven primarily by onshore assets. In addition, the differential technological requirements between asset types highlight differences in technical efficiency between firm types. In particular, if local firms are less efficient, their increased malfunctions should show up relatively more in complex

offshore assets rather than simple onshore ones where scope for efficiency gaps is smaller. This argument is suggestively supported by the revealed preference of different firms sorting into asset types, as shown in Figure A3. Over the study period we have seen local companies grow their onshore market share, while the offshore market remains firmly the purview of multinationals. This trend has occurred even as the offshore market has grown from 44.5% to 68.2% of national output (see Figure 1, Panel A). This revealed preference suggests comparative advantages in onshore production for local firms and offshore for multinationals.

TABLE 3 HERE

Table 3 supports these hypotheses. This table replicates Panel A of 2, but splits the sample into onshore (Panel A) and offshore (Panel B) fields. For shut-ins in columns (1)-(2), the effect of localness is indeed stronger for offshore assets (16.9 pp vs. 9.1 pp). However, the output effect for onshore assets is roughly 2.4 million barrels, while offshore it is only 0.8-1.7 million barrels and insignificant at conventional levels. While these differences themselves are not statistically significant, the local output effect seems to be driven primarily by higher risk onshore assets. In contrast, local operators cause substantially more spills in offshore sites. Local takeover of an onshore field increases malfunctions by only 0.7-1.4, significant only at 10%, while on offshore fields this number rises to 3.3-4.4, significant at 1%. The greater technological requirements of offshore extraction result in greater efficiency costs of local ownership, particularly when political risk is absent. At the same time, the political risks of onshore extraction give rise to a comparative advantage for local firms, highlighted by the concentration of their output gains in onshore fields.

Patterns of heterogeneity in crime effects across asset types should follow those of output if we believe that local firms increase output by mitigating onshore political risk. I find this to be the case in Table 4. For each outcome, I find that the effects of localization in offshore fields are very close to zero.¹⁹ This is also true to a lesser extent for pipeline sabotage – although still 36% of offshore fields experience sabotage at some point, compared with nearly 98% of onshore fields. Perhaps unsurprisingly, I find that the effect of indigenization on theft is entirely concentrated in offshore fields, where the coefficient ranges from 3.4-3.5, significant at 5%. The same is true of militancy and piracy, the latter of which however is still not statistically significant.

TABLE 4 HERE

¹⁹ In the case of militant attacks, it should be noted, this is by construction since very few militant attacks take place in offshore waters during the sample period.

Onshore fields drive the effects on both theft and output, suggesting they are linked by the onshore presence of criminal gangs that generate local comparative advantage. In contrast, increased malfunctions are concentrated offshore, highlighting the presence of technological barriers that generate multinational advantage. However, the aggregate output effects clearly show that the local political advantage dominates their efficiency disadvantage.

5.3.3 Additional robustness tests

In Appendix C.1, I test the robustness of the main TWFE results to potential sources of bias in the estimate of ψ beyond the generic parallel trends tests in Section 5.3.1. I first show that the effects in Table 2 are driven primarily by local private companies, rather than divestment to the state-owned oil company. I then rule out that differential responses to oil price fluctuations are driving the observed localization advantage. I also incorporate recent results from de Chaisemartin and D’Haultfoeuille (2019), Goodman-Bacon (2019), and Abraham and Sun 2018 on potential sources of bias in the TWFE estimator in staggered-adoption designs. Applying these methods, I test robustness of the main estimate to dynamic treatment effect heterogeneity, re-weighting, and cohort-specific heterogeneity. Despite the presence of dynamic and cohort-specific heterogeneous effects, I find that these issues are unlikely to affect the main TWFE and event-study results. Details are in Appendix C.1

6 Model

In this section, I develop a simple model to explain why local companies may do better at curbing crime and therefore increase output. In the model, firms set output quantities and then bargain with organized crime. The equilibrium of this game determines the level of theft and, the level of observed output, and the incentives to produce on the extensive margin. I identify several frictions in the bargaining process – bargaining costs, corruption costs, and partial ownership – that affect these equilibrium outcomes. The first predictions of the model are a set of comparative statics relating the levels of theft, shut-ins, and output to these bargaining frictions, as well as other parameters such as strength and costs of gangs, oil prices, and firm marginal costs. I argue that if local firms possess advantages across certain dimensions of the bargaining process, they may indeed produce lower theft, fewer shut-ins, and greater output as a consequence. The model concludes by looking at cross-partial derivatives of the key outcomes with respect to the bargaining frictions and other variables observable in the data, suggesting important patterns of heterogeneity in the main effects.

6.1 Set-up

The interaction is a simple two-stage game with players firms, indexed by $f \in F$, and gangs, indexed by $g \in G$. In the first stage of the game, firm chooses a level of output to produce Q using technology $c_f(Q)$. Assume any fixed costs are already sunk. The second stage of the game is a bargain between the firm and the gang over this output. The firm chooses a bribe b to the gang to dissuade them from theft. Firm strategies are a pair $(Q(f), b(f))$ mapping the type space F to \mathbb{R}_+^2 . The gang can either accept A or reject R the offer b . If the gang rejects, it steals a constant amount of output q , paying fixed cost $c - \epsilon_g$, where ϵ_g is random. Theft is inefficient, destroying output $\kappa > 0$. If the gang accepts, it receives b and all of Q goes to the firm. Firms are price takers at world oil price p .

Firms may differ in a number of important ways related to the cost of bargaining. If a bargain is consummated, firm f may pay a penalty Λ_f with probability λ_f if the behavior is discovered. For simplicity, normalize $\Lambda_f = 1$. This captures the fact that different firms may be subject to different legal and/or reputational costs of corrupt payments. Firms are also subject to differing bargaining costs. For a bribe b paid, a firm may pay a cost $1 + \tau_f$, an additional “tax” of doing business with gangsters. This captures payments made to intermediaries, frictions in the bargaining process, or principal-agent problems within the firm. For example, a firm that is well-connected to local political-criminal networks may costlessly interact with organized crime, thus nesting $\tau_f = 0$. Lastly, firms have limited liability, in that they only receive a share γ_f of Q , a common feature of joint-ventures in Nigeria, where the state has a roughly 40-60% stake in all oilfields.

6.2 Analysis

The subgame perfect Nash equilibrium is solved via backward induction. First consider the one-shot bargain.

Definition 1. Bargaining Range. *The bargaining range B will be the set of mutually acceptable bribes, defined as the interval $[\underline{b}_g, \bar{b}_f]$, where \underline{b} is the lowest bribe g is willing to accept and \bar{b} is the highest bribe f is willing to pay. The gangsters will accept whenever $b > pq - c + \epsilon_g$. The firm will offer $b > 0$ whenever*

$$\gamma_f p(Q - q - \kappa) - c_f(Q) < \gamma_f pQ - c_f(Q) - \lambda_f - b(1 + \tau_f)$$

This yields the reservation points

$$\underline{b}_g = pq - c + \epsilon_g \quad \bar{b}_f = \frac{\gamma_f p(q + \kappa) - \lambda_f}{1 + \tau_f}$$

Definition 2. Firm bribe offer. Firms offer a take-it-or-leave-it bribe to the gangsters. The optimal bribe makes the gangster indifferent, therefore, $b^* = \underline{b}_g$.

Assumption 1. Information structure. Assume that the firm only observes the cost type of the gangster in the bargaining phase. When choosing quantities, the firm takes ϵ_g as stochastic. Assume it takes a uniform distribution on the interval $[0, c]$.

Assumption 2. Cost of corruption. Assume that: i) the firm has positive willingness to pay, $\bar{b}_f > 0$ and ii) the firm can't afford to bribe at least the lowest cost gangsters, $\epsilon_f = c$, i.e. $\bar{b}_f < pq$.

This translates to the following condition on the cost of corruption, derived in Appendix B.1.

$$\lambda_f \in [\min\{0, \lambda_f > \gamma_f p \kappa - pq(1 + \tau_f - \gamma_f)\}, \gamma_f p(q + \kappa)]$$

A bargain occurs whenever $\underline{b}_g < \bar{b}_f$. Using the uniform distribution of ϵ_g , the probability of a bargain occurring is

$$Pr(B) = 1 - \frac{pq}{c} \left(1 - \frac{\gamma_f}{1 + \tau_f}\right) + \left(\frac{\gamma_f p \kappa - \lambda_f}{1 + \tau_f}\right) \frac{1}{c}$$

Note that for theft to occur with positive probability, we must have $Pr(\neg B) > 0$, which reduces exactly to $\lambda_f > \gamma_f p \kappa - pq(1 + \tau_f - \gamma_f)$.²⁰

Proposition 1. Comparative statics. The likelihood of theft is increasing in τ_f, λ_f, q , and decreasing in c, γ_f, κ . Theft is increasing in p whenever $\frac{\kappa}{q} > \frac{1 + \tau_f}{\gamma_f}$.

The proof is in Appendix B.2. Note the condition for the comparative static on prices. If losses are high relative to theft, then an increase in price affects the company's reservation price relatively more than the gangster's, increasing \bar{b}_f and expanding the bargaining range. If the opposite is true, then the bargaining range contracts because \underline{b}_g rises relatively more. Note that under perfect bargaining, where $\tau_f = 0$ and $\gamma_f = 1$, the inefficiency of theft implies that $\frac{\partial Pr(B)}{\partial p} > 0$ is always true. For a given increase in p , gangsters increase \underline{b}_g by q while \bar{b}_f rises by $q + \kappa$.

In the first stage, the firm maximizes expected profit

$$\max_Q Pr(B)[\gamma_f p Q - (1 + \tau_f)E[\underline{b}_g | \epsilon_g \in B] - \lambda_f] + Pr(\neg B)\gamma_f p(Q - q - \kappa) - c_f(Q)$$

Assuming the A2.2 is met, then

$$E[\underline{b}_g | \epsilon_g \in B] = \frac{\bar{b}_f - pq + c}{2}$$

Assumption 3. Firm technology. Firms have quadratic costs, $c_f(q) = \frac{c_f q^2}{2}$. Firms may differ on

²⁰ See B.1.

technical efficiency, indexed by parameter c_f . The optimal quantity is

$$Q^* = \frac{\gamma_f p}{c_f}$$

Since the costs of bargaining and theft are additive, they enter the firm's problem as a fixed cost. That being the case, we must analyze the decision to produce on the extensive margin, given the expected losses due to theft and bribes. Firm will produce on a field if:

$$\gamma_f p Q^* - c(Q^*) > Pr(B)[(1 + \tau_f)[E[b_g | \epsilon_g \in B] + \lambda_f] + Pr(\neg B)\gamma_f p(q + \kappa)]$$

Assumption 4. Lowest-type gangster. Assume that theft is profitable even for the weakest gangster, $\epsilon_g = 0$, so that $pq > c$.

Proposition 2. Comparative statics: shut-ins. A shut-in refers to a firm's corner solution, when production at the optimal quantity yields negative expected profits. Under A1-A4, shut-ins are increasing in λ, τ and decreasing in γ . Proof is in B.3.

Proposition 3. Comparative statics: output. Define observed output as $\tilde{Q} = Q^* - Pr(\neg B)(\kappa + q)$. Conditional on $Q^* > 0$, \tilde{Q} is increasing in all the same things as $Pr(B)$ as long as A2 holds and theft is inefficient (e.g. $\kappa > 0$). Proof follows directly from Proposition 1.

7 Mechanisms

The model suggests that local advantage can arise from lower bargaining costs τ , lower corruption costs λ , or greater output shares γ . In addition, if local firms produce greater spillover benefits than multinationals, it is conceivable that gangsters might face a greater c . This would be especially salient if local firms create local jobs, bidding up the opportunity cost of joining gangs for young men and therefore increasing the cost of labor to the gangs. In this section, I test each of these mechanisms.

7.1 Bargaining costs

Foreign firms may face greater bargaining costs. I study one observable feature that is likely to drive bargaining costs: political connections. Prominent Niger Delta politicians have been linked to organized criminal groups, while observers argue that complicity among the security forces stationed in the region also plays an important role in the illegal oil market (SDN 2019a, SDN 2019b, Asuni 2009). Such connections are likely to reduce the cost of interacting with gangs, as these agents can

act as go-betweens and guarantors of informal bribe contracts. Local firms may be more likely to cultivate these ties, as they are less constrained by multinational hiring practices or disclosure requirements enforced by international stock exchanges, and generally more connected with the local political economy. The descriptive evidence in Table 1 shows that political connections are widespread, affecting 28% of all field-years, a share that is relatively similar for treated and untreated. However, the proportions heavily favor local firms once we look at connections to army and local politicians, where the share is 10 and 13%, respectively, for treated fields but roughly 0% for untreated. It does appear that local firms, while not more likely to cultivate connections in general, are indeed more likely to cultivate the strategic connections likely to affect bargaining.

7.1.1 Political connections

To identify the political connections channel, I proceed as follows. First, I test whether political connections are indeed important determinants of theft. Secondly, I disaggregate political connections to identify whether effects are heterogeneous across types of connections – in particular concentrated among local politicians and security forces. Having established these facts, I re-estimate the TWFE specification for local advantage with the additional political connections terms included on the right-hand side of the model. Throughout, I focus on theft, since as per the model this is the outcome most directly related to bargaining costs.

Table 5 contains the results of the TWFE regression of theft on political connections.²¹ As with local ownership, this estimate is identified from field takeovers by politically connected companies. The political connections variables are defined as follows: “any politician” indicates that field i is operated by a company with a current or former member of any level of Nigerian government on board, management, or shareholder. “Elected politicians” refer only to elected officeholders, such as senators, governors, or representatives. “Technocrats” are those associated with ministerial posts or regulatory agencies, typically the NNPC, DPR, or Ministry of Petroleum Resources. “Same-state politician” indicates that the company operating field i in state s is connected to a politician specifically in state s . Lastly, “security forces” are those linked to the military or police forces.

TABLE 5 HERE

Each pair of columns in Table 5 indicates the impact of a specific type of connection, estimated

²¹ Note that the sample here is 2941 field-years as opposed to 3069 in the main estimation. This is because of the three firms with missing political connections data.

with and without control variables. Having any political connection reduces field-level theft by 2.2-2.7 incidents per year on average (columns (1)-(2)), although these effects are only significant without controls. The effects of technocratic politicians (columns (2)-(5)) are also negative, but small and insignificant. In contrast, elected politicians provide a larger advantage, resulting in between 2.4-2.9 fewer theft incidents, the latter of which is significant at the 5% level. However, by far the most pronounced effects are for connections to same-state politicians (columns (7)-(8)) and the security forces (columns (9)-(10)), at 3.2-4.2 and 4.9-6.2, respectively, all of which are significant. These effect sizes are summarized in Figure A12, which plots the coefficients and confidence intervals for each category of political connection. Addition of controls makes only slightly affects the estimates.

Clearly, then, political connections matter: politically connected firms are able to mitigate the activities of organized crime. Importantly, the effects are concentrated among elected politicians, local politicians, and security forces. The smaller effects for connections broadly, and technocratic ones in particular, suggests that political connections do not work via better relationships with regulatory agencies. Instead, connections to local and elected politicians give firms access to the local patronage networks that sustain black market activity. At the same time, access to security forces allows firms to leverage the selective enforcement of these agencies. Since the decision to enforce or collude by security forces is perhaps the primary determinant of the cost of theft, access gives connected firms the means to redirect theft away from their assets. Both of these should reduce costs of doing business with gangsters, ultimately resulting in greater bargaining space and lower theft.

To probe the validity of the TWFE approach for estimating the effect of political connections, I estimate event-study models for local politicians and security forces. The coefficients are plotted in Figure 6 for the event-study regressions of theft on local political connections (Panel A) and security force connections (Panel B). In both cases, they are insignificant for $\tau < 0$, and observe a steadily increasing negative impact after the event. In addition, both types of political connections result in an immediate drop in field-level theft, however, in the case of army connections it takes several years for the negative effects to become consistently significant.

FIGURE 6 HERE

We have established that local firms are more likely to have certain types of political connections, and that these specific connections result in lower theft. But can we then conclude that political

connections are the mechanism for local advantage? To test this, I estimate the following equation:

$$y_{it} = \alpha + \psi_1 local_{it} + \psi_2 pol_{it} + \delta_t + \xi_i + X'_{it}\beta + u_{it}$$

Which includes time-varying treatment indicators for both whether the field is controlled by a local company, and whether the field is controlled by a politically connected company, pol_{it} . Under the assumption that political connections are an important mechanism by which localness affects bargaining, we should have $\psi_1 > \psi_2$ – the magnitude of the local advantage estimate falls once political connections are accounted for. The results are presented in Table A9.

TABLE A9 HERE

Column (1) re-prints the main effect for reference. The local advantage with respect to theft remains large and significant even after controlling for connections to any politician (2) and technocrats (3). In fact, for each of these specifications the estimated coefficient on $local_{it}$ actually slightly increases when the particular form of political connections is controlled for. In contrast, the coefficient falls to 2.3 and 2.4, once local politician and security forces connections are accounted for, respectively. Neither of these estimate is statistically significant. Therefore, we can conclude that these specific political connections, which are relatively more common among local firms, are indeed a source of the local advantage in dealing with organized crime. Local firms are more likely to be connected to local political-criminal networks, and thus face lower frictions to interacting with gangs.

7.2 Corruption cost

Multinational firms may face higher expected costs of λ of engaging in corrupt behavior. In general, these costs are driven by home anti-corruption statutes that prohibit multinationals from improper payments to foreign officials, such as the Foreign Corrupt Practices Act (FCPA) in the United States. Given the relatively broad definitions of foreign officials contained in these laws, and the need to employ local agents – some of whom may be government officials – to conduct side-payments, the prospect of legal liabilities could plausibly deter multinationals from bargaining with gangsters. If this does matter, we should observe that even within multinationals, exposure to these laws should explain variation in levels of theft. Restricting the sample to multinationals also allows me to remain agnostic about the content, quality, and enforcement of Nigeria’s own anti-corruption laws.²²

²² This is preferable to assessing the effectiveness of these laws, which legal analysis suggest are basically ineffective (Aigbovo and Atsegbua 2013).

Every multinational firm in Nigeria’s oil sector currently falls under some form of foreign anti-bribery statute. In order to test this hypothesis in a TWFE model, I employ the staggered nature of law passages. The US FCPA was passed in 1977, but the UK Bribery Act, which covers Shell, was only passed in 2010. The Italian statute governing Agip was passed in 2012, the Swiss statute governing Addax (until its sale to SINOPEC in 2009) was passed in 2000, while the French law governing Total was not passed until 2017. Thus, there is considerable variation in the timing of laws governing each oilfield over the sample period, allowing for a DD approach.

TABLE 6 HERE

The results of this estimation for each of the six major outcomes are contained in Table 6. The sample is all field-years with a multinational operator. In general, foreign corruption laws have limited effect on the actual production decisions of the firm (Panel A) – the signs of the coefficients are not consistent and none of the estimates are significant. However, in Panel B columns (1)-(2), we can see that increased corruption costs do impact the ability of multinational firms to mitigate theft on their assets. The passage of a home-country corruption law is associated with 2.7-6.7 increase in theft, or 24.4-58.6% of the multinational sample mean, significant at the 1% level.

I test for divergent pre-trends in Figure A13, again employing an event-study specification and plotting the coefficients for $\tau \in [-5, 6]$. In Panel A, I plot coefficients from the unweighted event-study specification, while in Panel B I re-weight these coefficients by cohort shares following Abraham and Sun (2018) to account for non-convex weights and heterogeneity across cohorts. In Panel A, the pre-law coefficients are generally not significant. However, there is some evidence of an upward trend in these coefficients that could be accounting for the increased level of sabotage. In Panel B, however, the re-weighting procedure produces pre-event coefficients that are much closer to zero and exhibit no discernible trend. The post-event coefficients are all positive and significant with the exception fo year six, by which time the effect appears to fall to zero.

7.3 Alternative explanations

Several alternative explanations outside of the bargaining model may account for the results. In this section, I consider several salient issues.

Local employment spillovers: Part of the rationale behind indigenization is that local firms may increase the positive spillover effects of oil production to local communities. If this is the case, then it’s possible that the effects we see are driven by higher opportunity costs for attracting labor into the

criminal sector. If spillovers improve employment opportunities for young men, then the gangster's cost c may rise as labor costs rise. To test this hypothesis, I use data from three rounds of Nigeria's General Household Survey, a panel survey covering 32,537 Nigerians in 500 villages from 2010-2016. I link each village to its nearest oilfield, and estimate a TWFE model the effect of localization of nearby fields on employment and consumption outcomes in Appendix C.2. I find no evidence of employment spillovers from local ownership.

Differences in discount rates: The local advantage in production may be driven not by organized crime and differential bargaining frictions, but rather by discount rates. Local companies may simply have different optimal extraction profiles given underlying preferences, a plausible mechanism if local companies have shorter time horizons or differential accountability to investors. This alternative mechanism implies that extraction profiles should decline more rapidly under local ownership, with more oil extracted earlier in the life cycle of a field. Figure A16 plots extraction profiles for each type of ownership using a local polynomial fit, as well as the full sample for comparison in Panel A. Interesting, local private fields do display substantially different extraction profiles from both multinational and state-owned fields. However, these extraction profiles are significantly flatter over the field's lifetime, and actually increasing over time. This suggests, if anything, greater patience on the part of local private firms, implying that shorter time horizons are not driving the results.

Selection into field takeover: Fields that are divested to local investors are not randomly chosen. While Figures 3 and 4 suggest that selection in to takeover does not lead to a violation of parallel trends, it is possible that unobserved field-level characteristics might drive the result without appearing as pre-trends. For example, if multinationals sell fields with declining reserves, then this might incentivize faster extraction by local owners even as the path of output is flat before the sale.

8 Heterogeneity

The model predicts substantial heterogeneity; scrutinizing the expressions for the main comparative statics in Section B.2 suggests interesting cross-partials with respect to the costs of gangs c , their bunkering strength q and the price of oil p .

8.1 Costs

As crime costs increase, the bargaining range expands. This has the effect of reducing the sensitivity of theft to bargaining costs, since it is only where the bargaining range is somewhat restricted that frictions bite. These cross-partials imply signs on the interaction term of a heterogeneous effects

regression. In particular, the negative effect of localization on theft should fall in magnitude as costs rise. To proxy for trade costs in the black market, I consider three geographic variables that should affect bunkering costs – *i*) distance to the Niger river, which reduces costs of transporting stolen oil, *ii*) distance to the coast, the location of export markets, and *iii*) distance from the field to the nearest militant camp. The latter is perhaps the most important cost-shifter – from the perspective of the gang, closer fields are lower-cost targets.

I test the prediction that the benefits of localization for oil theft should be decreasing in costs by estimating the main regression with an added interaction term $c_i \times local_{it}$, where c_i is one of the bunkering cost variables. I consider specifications with each individual variable, as well as the full set of interactions. The results are in Table 7. Columns (1)-(3) gives the results for distance to Niger River, coast, and militant camp, respectively. As predicted by theory, interaction terms on both (1) and (3) are positive significant, though (2) is negative but not significant. When we consider all three cost-shifters together in columns (4)-(7), then only the interaction with distance to militant camps becomes positive and significant, increasing substantially in magnitude relative to column (3). The coefficients on the interactions between local ownership and distance to coast and river are insignificant.

TABLE 7 HERE

These differences are visualized in Figure 7, which plots the coefficients from the regressions in columns (6) and (7), without and with controls. The most noticeable effects are the interactions with distance to militant camps. Quantitatively (from column (3)), the impact of localization for fields located directly next to a camp is -5.1, which increases toward zero by 0.099 with every additional kilometer away from a militant camp; this is sufficient to eliminate the entire effect after only 51 kilometers. Clearly, the benefits of localization are concentrated in those fields very close to militant camps. The model explains that these are the lowest cost fields for gangs to tap, where bargaining costs would be more likely to act as binding constraints on dealmaking.

FIGURE 7 HERE

Note that this result does not rule out explanations outside of the model. In particular, fields near militant camps might be dominated by larger, more centralized syndicates with whom a negotiation can actually occur, increasing the importance of bargaining costs. However, the fact that we see small or null effects for rivers and distance to the coast casts doubt on this explanation. The bunkering

market exhibits a two-tiered structure, populated both by large, export-driven criminal syndicates and local entrepreneurs serving the domestic market (SDN 2019a). Since trade costs don't matter, this suggests that the smaller syndicates who serve the local market are driving the localization effect.²³ This observation would also be consistent with the results of Section 8.3, which find that the benefit of localization is largest among lower-strength gangs. This would imply that the interactive effect of militant camp distance is driven by a cost mechanism and not by exposure to larger syndicates.

8.2 Oil prices

Recall that the sign of the effect of oil prices on theft is ambiguous: for an increase in prices to increase theft, we must have $\frac{\kappa}{q} < \frac{(1+\tau)}{\gamma} - 1$. This expression depends on the bargaining friction τ , although not on fixed corruption cost λ . Assume this condition is met for some combination of model parameters. Then, all else equal, as τ falls this condition is more difficult to meet, so that response of prices may become negative for low frictions. The intuition is that price increases raise the reservation bribes of both actors. But as frictions fall, the willingness to pay of the firm is affected relatively more, to the point where this effect eventually dominates so that price increases widen the bargaining range. To test these implications, I estimate an interaction specification of the TWFE model

$$y_{it} = \alpha + \theta_0 p_t + \theta_1 local_{it} + \theta_2 local_{it} p_t + \delta_t + \zeta_i + X'_{it} \beta + v_{it}$$

Where p_t is the demeaned world price of crude oil, relative to the long-run mean.²⁴ The empirical implication is that while the price effect alone may be negative or positive, its interaction will clearly be negative—that is, higher prices have increasingly negative effects on theft as bargaining costs fall. So the estimate of θ_2 should be negative. Another implication of the model is that the sign of the price effect itself can tell us something about the size of $\frac{\kappa}{q}$, the ratio of spillage losses to illicit gains. In order to identify θ_0 in the context of a fixed-effects model, I exclude δ_t in some specifications. θ_0 then measures the responsiveness of theft to oil prices on multinational-operated fields.²⁵

The results are given in Table 8, which displays the estimated coefficients $\hat{\theta}$. Specifications (1)-(4) omit the time fixed effect to identify θ_0 ; columns (1) and (2) estimate the model without any fixed effects, while columns (3) and (4) include ζ_i . In columns (2) and (4) controls are additionally interacted with p_t to control for potential omitted variables correlated with localization that might

²³ Or that export and transport costs are low enough everywhere to drive substantial variation in the overall cost structure.

²⁴ The mean is calculated over the period 1998-2017.

²⁵ This is collinear with the year fixed effects and thus unidentified in the full TWFE model.

respond similarly to oil price trends. Columns (5)-(6) estimate the full TWFE specification.

TABLE 8 HERE

As predicted, the sign of the interaction θ_2 is robustly negative and significant in all specifications. The estimates imply that among locally-operated fields, the average responsiveness of theft to price is roughly 0.091 to 0.141 incidents lower than among multinationals. The coefficients on p_t in columns (1)-(4) show that, in fact, the level price effect is positive – on multinational assets, theft increases in prices. This is positive in all specifications and significant at 5% in two. Putting these estimates together, the effect of prices on theft is reduced from a positive, significant effect to essentially zero and insignificant among local firms with lower bargaining costs. Furthermore, the positive coefficient θ_0 also implies that pure losses from theft κ are actually low relative to the quantity stolen q , which leads the reservation price of firms to be less sensitive to oil prices. Finally, the estimates of θ_1 are generally negative and significant. Since p_t is demeaned, this implies that at long-run average prices, localization effects remain negative. As prices rise, so too do the benefits of localization, since for low frictions higher prices make reaching a deal more valuable to the firm.

8.3 Capacity for violence and competition

Additional evidence suggests that the benefits of localization are concentrated on assets where nearby gangs have lower capacity for violence. Hvinden and Rexer (2019) show that in the Niger Delta conflict, militant groups with more allies connected locally along the pipeline network have greater capacity for violence and receive more generous amnesty deals as a result. Using distance-weighted allied connections along a pipeline as a proxy for violence capacity among the nearest gang, I find that treatment effects for both oil theft and militancy are largest among areas with weaker groups. In Table 9, I interact the main localization regression with this measure of the destructive capacity of the nearest militant group, measured as of 2009.

TABLE 9 HERE

The results indicate that the benefits of localization are largest – between 3.4-4.1 fewer thefts per annum (Panel A) – when the nearest militant group has no local allies. The coefficient then attenuates toward zero with each additional ally, indicating a militarily stronger group, though this interaction is not significant. These results are unaffected by the addition of controls or using the onshore-only

sample (column 4). Similar results obtain for the analysis of militant activity in Panel B. Localization causes 1.7-2.7 fewer militant-related deaths annually when nearby militants are weak, an effect that reverses among stronger militant groups. Across both outcomes the heterogeneity results are also strongest for fields nearer to militant camps, where exposure to gangs is stronger. The effects disappear when we consider fields further than 20 kilometers from a militant camp in (6).

As the model makes clear, only groups that are on the margin can be bribed – for certain values of the model parameters, it may be the case that, on average, some subpopulation of gangs are either never or always worth bribing. Groups that can credibly threaten violent retaliation may fall into the latter category. These groups may be powerful enough to always obtain a deal if the threat of violence increases willingness to pay enough, so that among this population, changing bargaining frictions on the margin has no effect.²⁶ This interpretation is consistent with the results of Table 7, which implies that export costs are not an important dimension of heterogeneity. Those affected by localization are likely to be smaller groups who both do not export and have lower capacity for violence.²⁷

It is important to note that this effect is not mechanically driven by densely allied groups being surrounded by a larger number of groups (i.e., more nearby groups makes it harder to negotiate/coordinate/enforce bargains). When I include controls for the interaction between number of nearby groups within d kilometers and localization, the positive and significant interaction coefficient on group strength does not change.²⁸ In these regressions, I also find suggestive evidence that fields in more saturated black markets actually respond more to localization, once group strength is conditioned on.

Figure A17 plots coefficients from regressions where the interactions between localization and both group strength and group number are included, for oil theft and militancy outcomes. I find that the coefficient on the group number interaction is either zero or negative and sometimes significant, suggesting that localization benefits are weakly increasing in the number of local gangs. These effects are smallest when we measure gangs as the number of individual camps, and monotonically larger for commanders and groups. This is because aggregation captures coordination effects between camps under the same commander and commanders under the same umbrella group. Varying the size at which the market is defined,²⁹ I find that the results are stronger for larger markets.

Within the framework of the model, this result is consistent with greater competition leading to

²⁶ In other words, A2.ii is violated so that even the lowest-cost gangsters are worth bribing given the threat of violence.

²⁷ Export syndicates and capacity for violence are likely to be positively correlated, according to local qualitative research (SDN)

²⁸ Results available upon request.

²⁹ Defined as the radius around a given oilfield

higher q , which in turn increases the sensitivity of theft and militancy to bargaining costs as shown in Section B.2. If we think of q as arising out of a Cournot or tragedy of the commons-type interaction in which gangs compete over oil and individual quantities stolen are strategic substitutes, then it follows that monopolists should steal less, while competitive markets steal more. Analysis of the cross-sectional correlation between various measures of competition and oil theft provide some suggestive support for this in Figure A18. Across three different geographic market definitions (10, 50, 100 km) and each level of aggregation in defining market participants, I find that more competition is associated with greater quantities of theft. All results are conditional on the standard set of controls and year fixed effects. The relationship is subject to diminishing returns, which is consistent with an N -player symmetric Cournot equilibrium, where output scales at rate $\frac{N}{N+1}$. Of course, these results are purely correlational and should be taken only as suggestive support for the interpretation that greater competition among gangs increases q , thereby increasing the returns to localization.

9 Conclusion

Multinationals have substantial advantages over local firms in many markets. They use better technology, hire better workers, have greater access to foreign capital markets, and employ better management practices. Yet our understanding of the multinational advantage comes almost exclusively from manufacturing and service firms in relatively politically stable contexts. In this paper, I show that in the troubled natural resource sectors of countries suffering from pervasive violence, criminality, and corruption, the multinational advantage can become a substantial liability.

I demonstrate this in the context of Nigeria's oil sector – the paradigmatic example of the resource curse. In a context in which militant groups and organized crime are ever-present threats to firm operations and corruption buys protection for assets, local companies may possess distinct advantages. Using data on Nigeria's active oilfields from 2006-2016, I find that fields operated by multinationals are substantially less productive than those operated by local firms. For the average oilfield, a local takeover increased revenue by more than 160 million dollars per year, nearly a doubling. Local firms accomplish this feat in part by reviving moribund fields: the likelihood that a field is "shut-in" falls dramatically upon local takeover. Importantly, local fields operate at a technical disadvantage – they experience greater numbers of operational failure and equipment malfunctions.

The key to the local advantage is in dealing with the multi-billion dollar black market for stolen oil. I find that local takeovers reduce incidents of theft and militant violence by 30% and 85%, respectively. I further find that these gains are concentrated in the onshore fields most susceptible to

crime and violence, whereas the losses from equipment failure are concentrated in technically sophisticated offshore fields. This further underscores that while multinationals have a technology advantage, the black market generates a much larger local advantage. Private Nigerian firms drive the improvements in output, crime, and violence, while the state-owned company drives the increase in malfunctions. In fact, for local private firms there appears to be no efficiency cost to indigenization.

A model of the bargaining interaction between firms and organized crime shows that if local firms have lower bargaining cost and corruption penalties, then they may indeed perform better. To test this, I draw on the literature on firm-politician connections in developing countries. Using scraped data on the biographies of boardmembers, shareholders, and managers, I measure firm connections to elected politicians, local politicians, technocrats, and the security forces. Using these connections as a proxy for bargaining costs, I show that connections to local politicians and security forces – two groups intimately linked to the black market – reduces theft substantially and explains much of the negative effect of local ownership on black market activity. I also show that multinationals exposed to higher penalties from foreign corruption laws in their home country experience a spike in theft after those laws are passed. The results suggest that superior political connections and lower costs to corrupt behavior contribute to local advantage.

The findings imply that, at least in the Nigerian case, indigenization can reduce conflict in natural resource sectors at minimal or no cost to efficiency. In particular, when political and social conflict in the sector is extreme, the gains in output may be large enough to swamp the loss of multinational productivity and therefore justify indigenization on efficiency grounds. More broadly, the results point to a type of “greasing the wheels” corruption – in which local firms have a comparative advantage – that may indeed be welfare-improving given a particular set of institutional and economic constraints. However, given that I find no evidence that indigenization of oil assets improved local employment prospects in oil-producing communities, we should temper optimism that local ownership will fundamentally alter the enclave nature of oil extraction.

It may be tempting to dismiss the Nigerian experience as an extreme case of the resource curse. But in extractive sectors across the globe – from Congolese minerals to Ghanaian gold mining – firms face a political economy characterized by black markets, organized crime, violent armed groups, and corrupt politicians. It is unwise to apply conventional economic wisdom from manufacturing and service firms in middle-income countries to the specific dynamic of the natural resource sectors in poor ones.

References

- Abraham, Sarah and Liyang Sun (2018), "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Working Paper*.
- Aigbovo, Osaretin and Lawrence Atsegbua (2013), "Nigerian anti-corruption statutes: an impact assessment." *Journal of Money Laundering Control*, 16, 62–78.
- Aitken, Brian and Ann Harrison (1999), "Do domestic firms benefit from direct foreign investment? evidence from venezuela." *American Economic Review*, 89, 605–618.
- Akcigit, Ufuk, Salome Baslandze, and Francesca Lotti (2018), "Connecting to power: Political connections, innovation, and firm dynamics." *NBER Working Paper*, 25136.
- Alfaro, Laura and Jasmina Chauvin (2020), "Foreign direct investment, finance, and economic development." In *Encyclopedia of International Economics and Global Trade* (Francisco L Rivera-Batiz, Mariana Spatareanu, and Can Erbil, eds.), Singapore: World Scientific.
- Alfaro, Laura and Maggie X. Chen (2018), "Selection and market reallocation: Productivity gains from multinational production." *American Economic Journal: Economic Policy*, 10, 1–38.
- Aragon, Fernando M., Punam Chuhan-Pole, and Bryan Christopher Land (2015), "The local economic impacts of resource abundance: What have we learned?" *World Bank Policy Research Working Paper*.
- Aragon, Fernando M. and Juan Pablo Rud (2011), "Polluting industries and agricultural productivity: Evidence from mining in ghana." *Journal of International Money and Finance*, 126, 1980–2011.
- Aragon, Fernando M. and Juan Pablo Rud (2013), "Natural resources and local communities: Evidence from a peruvian gold mine." *Journal of International Money and Finance*, 5, 1–25.
- Arnold, Jens Matthias and Beata S. Javorcik (2009), "Gifted kids or pushy parents? foreign direct investment and plant productivity in indonesia." *Journal of International Economics*, 79, 42–53.
- Asuni, Judith Burdin (2009), "Understanding the armed groups of the niger delta." *Council on Foreign Relations Working Paper*.
- Balsvik, Ragnhild (2011), "Is labor mobility a channel for spillovers from multinationals? evidence from norwegian manufacturing." *The Review of Economics and Statistics*, 93, 285–297.
- Berman, Nicolas, Mathieu Couttenier, Dominic Rohner, and Mathias Thoenig (2017), "This mine is mine! how minerals fuel conflicts in africa." *American Economic Review*, 107, 1564–1610.

- Blair, Graeme, Darin Christensen, and Valerie Wirtschafter (2019), "How does violence shape investment? evidence from mining." *Working Paper*.
- Blair, Graeme and Kosuke Imai (2013), "Civilians and the strategic use of information during conflict in resource-rich territory." *International Growth Centre Working Paper*.
- Bloom, Nicholas and John Van Reenen (2010), "Why do management practices differ across firms and countries?" *Journal of Economic Perspectives*, 24, 203–224.
- Bloom, Nicholas, Raffaella Sadun, and John Van Reenen (2012), "Americans do it better: Us multinationals and the productivity miracle." *American Economic Review*, 102, 167–201.
- Bruederle, Anna and Roland Hodler (2019), "Who is screened out? application costs and the targeting of disability programs." *PNAS*, 116, 5467–5471.
- Buonanno, Paolo, Ruben Durante, Giovanni Prarolo, and Paolo Vanin (2015), "Poor institutions, rich mines: Resource curse in the origins of the sicilian mafia." *Economic Journal*, 125, 175–202.
- Burger, Martijn, Elena Ianchochina, and Bob Rijkers (2015), "Risky business: Political instability and sectoral greenfield foreign direct investment in the arab world." *World Bank Economic Review*, 30, 306–331.
- Christensen, Darin (2019), "Concession stands: How mining investments incite protest in africa." *International Organization*, 73, 65–101.
- Couttenier, Mathieu, Pauline Grosjean, and Marc Sangnier (2017), "The wild west is wild: The homicide resource curse." *Journal of the European Economic Association*, 15, 558–585.
- Criscuolo, Chiara and Ralf Martin (2009), "Multinationals and u.s. productivity leadership: Evidence from great britain." *The Review of Economics and Statistics*, 91, 263–281.
- Cust, James and Steven Poelhekke (2015), "The local economic impacts of natural resource extraction." *Annual Review of Resource Economics*, 7, 251–268.
- de Chaisemartin, Clement and Xavier D'Haultfoeuille (2019), "Two-way fixed effects estimators with heterogeneous treatment effects." *NBER Working Paper*, 25904.
- de la Sierra, Raul Sanchez (2019), "On the origin of states: Stationary bandits and taxation in eastern congo." *Journal of Political Economy*, Forthcoming.

- Deshpande, Manasi and Yue Li (2019), "Who is screened out? application costs and the targeting of disability programs." *American Economic Journal: Policy*, Forthcoming.
- Dube, Oeindrila and Juan F. Vargas (2013), "Commodity price shocks and civil conflict: Evidence from colombia." *Review of Economic Studies*, 80, 1384–1421.
- Faccio, Mara (2006), "Politically connected firms." *American Economic Review*, 96, 369–386.
- Fetzer, Thiemo and Stephan Kyburz (2018), "Cohesive institutions and political violence." *Centre for Competitive Advantage in the Global Economy*, Working Paper.
- Fisman, Raymond (2001), "Estimating the value of political connections." *American Economic Review*, 91, 1095–1102.
- Goodman-Bacon, Andrew (2019), "Difference-in-differences with variation in treatment timing." *Working Paper*.
- Gormley, Todd A. and David A. Matsa (2011), "Growing out of trouble? corporate responses to liability risk." *Review of Financial Studies*, 24, 2781–2821.
- Guadalupe, Maria, Olga Kuzmina, and Catherine Thomas (2012), "Innovation and foreign ownership." *American Economic Review*, 102, 3594–3627.
- Guidolin, Massimo and Eliana La Ferrara (2007), "Diamonds are forever, wars are not: Is conflict bad for private firms?" *American Economic Review*, 97, 1978–1993.
- Harrison, Ann and Andres Rodriguez-Clare (2010), "Trade, foreign investment, and industrial policy for developing countries." In *Handbook of Development Economics, Volume 5* (Dani Rodrick and Mark R. Rosenzweig, eds.), chapter 63, 4040–4214, Elsevier B.V.
- Javorcik, Beata S. and Shang-Jin Wei (2009), "Corruption and cross-border investment in emerging markets: Firm-level evidence." *Journal of International Money and Finance*, 28, 605–624.
- Khwaja, Asim Ijaz and Atif Mian (2005), "Do lenders favor politically connected firms? rent provision in an emerging financial market." *The Quarterly Journal of Economics*, 120, 1371–1411.
- Kyburz, Stephan (2018), "The local political resource curse." *Working Paper*.
- Li, Hongbin, Lingsheng Meng, Qian Wang, and Li-An Zhou (2008), "Political connections, financing, and firm performance: Evidence from chinese private firms." *Journal of Development Economics*, 87, 283–299.

- Lippert, Alexander (2014), "Spill-overs of a resource boom: Evidence from zambian copper mines." *Oxford Centre for the Analysis of Resource Rich Economies Working Papers*, 131.
- Loayza, Norman and Jamele Rigolini (2016), "The local impact of mining on poverty and inequality: Evidence from the commodity boom in peru." *World Development*, 84, 219–234.
- NBS (2017), "Demographic statistics bulletin." Technical report, Nigeria National Bureau of Statistics, Abuja, Nigeria.
- NEITI (2011), "2009-2011 oil and gas industry audit report." Technical report, Nigeria Extractive Industries Transparency Initiative, Abuja, Nigeria.
- Nwokolo, Arinze (2018), "Oil price shocks and civil conflict: Evidence from nigeria." *HiCN Working Paper*, 274.
- Obi, Cyril and Siri Aas Rustad (2011), *Oil and Insurgency in the Niger Delta*. Zed Books.
- Poole, Jennifer P. (2013), "Knowledge transfers from multinational to domestic firms: Evidence from worker mobility." *The Review of Economics and Statistics*, 95, 393–406.
- SDN (2019a), "Communities not criminals: Illegal oil refining in the niger delta." Technical report, Stakeholder Democracy Network, Port Harcourt, Nigeria.
- SDN (2019b), "Economic dynamics of the artisanal oil industry in the niger delta over five years." Technical report, Stakeholder Democracy Network, Port Harcourt, Nigeria.
- SDN (2019c), "Pipeline surveillance contracts in the niger delta." Technical report, Stakeholder Democracy Network, Port Harcourt, Nigeria.
- Sexton, Renard (2019), "Unpacking the local resource curse: How externalities and governance shape social conflict." *Journal of Conflict Resolution*, Forthcoming.
- Teece, David (1977), "Technology transfer by multinational firms: The resource cost of transferring technological know-how." *Economic Journal*, 87, 242–261.
- Watts, Michael (2007), "Petro-insurgency or criminal syndicate: Conflict and violence in the niger delta." *Review of African Political Economy*, 34, 637–660.

Table 1: Summary statistics

	Untreated	Treated	Full Sample
Covariates			
Field latitude	4.97 (0.60)	5.29 (0.64)	5.05 (0.62)
Distance to coast (km)	33.51 (29.48)	28.87 (27.06)	32.34 (28.91)
Distance to Niger River (km)	77.69 (71.59)	64.68 (55.26)	74.40 (67.97)
Distance to state capital (km)	86.30 (52.21)	82.91 (54.72)	85.45 (52.77)
Distance to militant camp (km)	30.00 (24.15)	32.80 (20.46)	30.71 (23.27)
Number of wells	20.76 (31.57)	14.44 (21.32)	19.16 (29.41)
Year of first well	1972.54 (11.14)	1973.47 (12.22)	1972.77 (11.41)
Onshore field	0.71 (0.45)	0.84 (0.37)	0.74 (0.44)
Max well depth (m)	2690.70 (808.05)	2834.25 (889.94)	2726.99 (830.17)
Outcomes			
Sabotage events	11.79 (25.35)	6.39 (13.21)	10.41 (23.00)
Militant deaths	1.19 (6.81)	1.85 (20.08)	1.36 (11.71)
Piracy attacks	0.10 (0.64)	0.14 (0.81)	0.11 (0.69)
Shut-in field	0.14 (0.35)	0.14 (0.35)	0.14 (0.35)
Annual oil production (million barrels)	3.72 (7.58)	2.63 (4.81)	3.46 (7.04)
Operational failure oil spills	8.06 (12.17)	4.27 (7.94)	7.10 (11.37)
Same-ethnicity company	0.60 (0.49)	0.31 (0.46)	0.51 (0.50)
Any politician	0.27 (0.45)	0.31 (0.46)	0.28 (0.45)
Technocrat	0.26 (0.44)	0.16 (0.37)	0.24 (0.43)
Elected politician	0.26 (0.44)	0.25 (0.44)	0.26 (0.44)
Security forces	0.00 (0.00)	0.10 (0.30)	0.03 (0.16)
Same-state politician	0.00 (0.05)	0.13 (0.33)	0.04 (0.18)
Number of clusters	208	71	279

Table displays means of variables with standard deviations in parentheses. Sample is a panel of 279 oilfields. Panel A gives summary statistics of field-level covariates while Panel B gives time-varying outcomes. Sample sizes indicate the number of unique oilfields in each group. Treated refers to all oilfields that have any local operator from

Table 2: The effect of divestment on output and criminality

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Output and efficiency</i>						
Outcome	Shut-in		Output		Malfunctions	
Local operator	-0.192*** (0.062)	-0.188*** (0.060)	2.538*** (0.671)	2.422*** (0.562)	1.904*** (0.728)	0.978 (0.696)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls \times Year FE	No	Yes	No	Yes	No	Yes
Observations	2310	2310	2310	2310	3069	3069
R^2	0.645	0.661	0.841	0.855	0.573	0.631
<i>Panel B: Crime and violence</i>						
Outcome	Theft		Militancy		Piracy	
Local operator	-2.839** (1.167)	-3.275** (1.378)	-1.400** (0.565)	-1.517** (0.587)	-0.072 (0.085)	-0.089 (0.084)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls \times Year FE	No	Yes	No	Yes	No	Yes
Observations	3069	3069	3069	3069	3069	3069
R^2	0.719	0.754	0.196	0.288	0.244	0.316

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 279 oilfields from 2006-2016. Shut-in is defined as a field registering zero output in a given year. Output is measured in millions of barrels of oil per year. Malfunctions are the total number of non-sabotage spills within 15 km of the field. Theft is the total number of sabotage spills within 15 km of the field. Militancy is the total number of militant deaths within 15 km of the field. Piracy is the number of pirate attacks within 15 km of the field. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: The effect of divestment on output by asset type

	(1)	(2)	(3)	(4)	(5)	(6)
Outcome	Shut-in		Output		Malfunctions	
<i>Panel A: Onshore fields</i>						
Local operator	-0.209*** (0.069)	-0.169*** (0.064)	2.417*** (0.687)	2.431*** (0.638)	1.395* (0.785)	0.682 (0.705)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls \times Year FE	No	Yes	No	Yes	No	Yes
Observations	1680	1680	1680	1680	2277	2277
R^2	0.647	0.674	0.651	0.666	0.647	0.714
<i>Panel B: Offshore fields</i>						
Local operator	-0.159** (0.079)	-0.091 (0.088)	1.708 (1.349)	0.861 (2.103)	4.375*** (1.135)	3.319** (1.339)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls \times Year FE	No	Yes	No	Yes	No	Yes
Observations	630	630	630	630	792	792
R^2	0.610	0.642	0.861	0.896	0.511	0.552

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 279 oilfields from 2006-2016. Sample is divided into onshore fields (Panel A) and offshore fields (Panel B). Shut-in is defined as a field registering zero output in a given year. Output is measured in millions of barrels of oil per year. Malfunctions are the total number of non-sabotage spills within 15 km of the field. Theft is the total number of sabotage spills within 15 km of the field. Militancy is the total number of militant deaths within 15 km of the field. Piracy is the number of pirate attacks within 15 km of the field. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: The effect of divestment on criminality by asset type

	(1)	(2)	(3)	(4)	(5)	(6)
Outcome	Theft		Militancy		Piracy	
<i>Panel A: Onshore fields</i>						
Local operator	-3.492** (1.471)	-3.378** (1.601)	-1.945*** (0.742)	-2.241*** (0.791)	-0.035 (0.091)	-0.094 (0.081)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	No	Yes	No	Yes
Observations	2277	2277	2277	2277	2277	2277
R^2	0.714	0.749	0.194	0.317	0.275	0.379
<i>Panel B: Offshore fields</i>						
Local operator	0.020 (0.032)	0.014 (0.040)	0.000 (.)	0.000 (.)	-0.070 (0.269)	-0.034 (0.274)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	No	Yes	No	Yes
Observations	792	792	792	792	792	792
R^2	0.165	0.332	.	.	0.201	0.269

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 279 oilfields from 2006-2016. Sample is divided into onshore fields (Panel A) and offshore fields (Panel B). Shut-in is defined as a field registering zero output in a given year. Output is measured in millions of barrels of oil per year. Malfunctions are the total number of non-sabotage spills within 15 km of the field. Theft is the total number of sabotage spills within 15 km of the field. Militancy is the total number of militant deaths within 15 km of the field. Piracy is the number of pirate attacks within 15 km of the field. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: The effect of political connections on criminality

Outcome	Theft									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Any politician	-2.719** (1.217)	-2.227* (1.303)								
Technocrat			-1.946 (1.337)	-1.667 (1.360)						
Elected politician					-2.949*** (1.065)	-2.405** (1.111)				
Same-state politician							-3.160** (1.330)	-4.185*** (1.554)		
Security forces									-4.949*** (1.719)	-6.191*** (2.261)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Observations	2941	2941	2941	2941	2941	2941	2941	2941	2941	2941
R ²	0.718	0.755	0.718	0.755	0.718	0.755	0.718	0.755	0.718	0.755

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 279 oilfields from 2006-2016 for which political connections data is available. Outcome variable is oil theft, the total number of sabotage spills within 15 km of the field. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. Political connections variables are dummy variables indicating that the operator of a given field-year has a particular type of politician as a board member, shareholder, or manager. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6: The effect of corruption costs on output and criminality

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Output and efficiency</i>						
Outcome	Shut-in		Output		Malfunctions	
Home-country corruption law	-0.007 (0.032)	-0.013 (0.037)	0.918*** (0.339)	0.087 (0.396)	1.583* (0.853)	-0.607 (0.798)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	No	Yes	No	Yes
Observations	2090	2090	2090	2090	2765	2765
R^2	0.669	0.680	0.838	0.853	0.570	0.632
<i>Panel B: Crime and violence</i>						
Outcome	Theft		Militancy		Piracy	
Home-country corruption law	6.652*** (1.050)	2.771*** (0.910)	-0.402 (0.494)	-0.358 (0.521)	-0.272*** (0.079)	-0.165** (0.082)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	No	Yes	No	Yes
Observations	2765	2765	2765	2765	2765	2765
R^2	0.726	0.757	0.203	0.302	0.270	0.342

Standard errors, in brackets, are clustered at the field level. Sample is the panel of oilfields from 2006-2016 operated by multinationals. Shut-in is defined as a field registering zero output in a given year. Output is measured in millions of barrels of oil per year. Malfunctions are the total number of non-sabotage spills within 15 km of the field. Theft is the total number of sabotage spills within 15 km of the field. Militancy is the total number of militant deaths within 15 km of the field. Piracy is the number of pirate attacks within 15 km of the field. Home country corruption law indicates that a field is operated by a company under the jurisdiction of a foreign anti-corruption statute. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 7: The effect of divestment on oil theft by cost shifters

Outcome Sample	Theft						
	Full					Onshore	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Local operator	-5.684*** (1.997)	-0.172 (1.945)	-5.111* (2.678)	-2.719 (3.148)	-2.418 (3.088)	-3.524 (4.349)	-2.917 (4.189)
Local operator × Distance to Niger River (km)	0.037** (0.015)			-0.016 (0.026)	-0.020 (0.025)	-0.015 (0.047)	-0.017 (0.045)
Local operator × Distance to coast (km)		-0.092** (0.043)		-0.130** (0.065)	-0.131** (0.062)	-0.123* (0.072)	-0.129* (0.067)
Local operator × Distance to militant camp (km)			0.099** (0.047)	0.183** (0.082)	0.188** (0.082)	0.217** (0.098)	0.210** (0.100)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	No	No	No	Yes	No	Yes
Observations	3069	3069	3069	3069	3069	2277	2277
R ²	0.734	0.729	0.740	0.770	0.775	0.760	0.765

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 279 oilfields from 2006-2016 in columns (1)-(5) and the sample of only offshore fields in (6)-(7). Outcome variable is oil theft, the total number of sabotage spills within 15 km of the field. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 8: The effect of divestment on oil theft by prices

Outcome	Theft					
	(1)	(2)	(3)	(4)	(5)	(6)
Local operator	-7.336*** (1.373)	-7.107*** (1.600)	1.284 (1.015)	1.249 (1.071)	-4.035*** (1.298)	-4.604*** (1.579)
Crude oil price (USD/barrel)	0.118*** (0.021)	0.137 (0.195)	0.123*** (0.021)	0.196 (0.196)		
Local operator × Crude oil price (USD/barrel)	-0.133*** (0.029)	-0.141*** (0.035)	-0.091*** (0.027)	-0.095*** (0.029)	-0.118*** (0.029)	-0.132*** (0.035)
Field FE	No	No	Yes	Yes	Yes	Yes
Year FE	No	No	No	No	Yes	Yes
Controls × Year FE	No	No	No	No	No	Yes
Controls × Oil Price	No	Yes	No	Yes	No	No
Observations	3069	3069	3069	3069	3069	3069
R ²	0.017	0.189	0.680	0.687	0.720	0.755

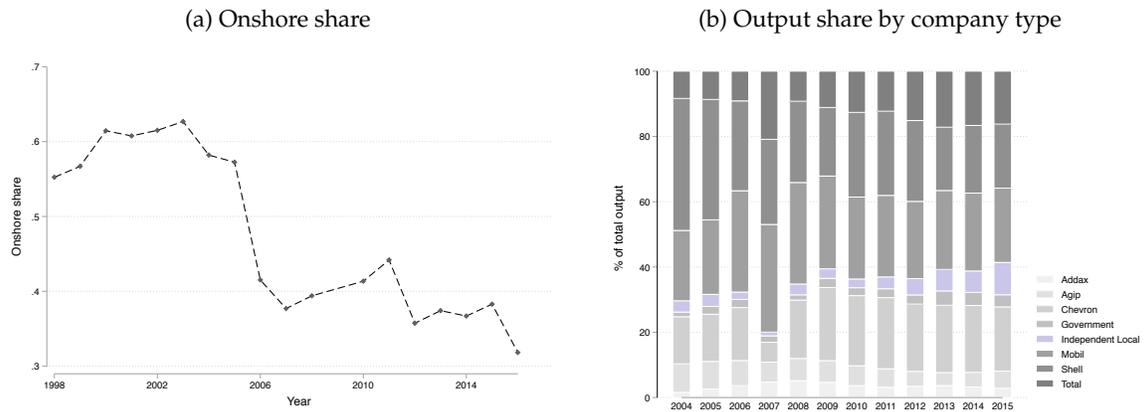
Standard errors, in brackets, are clustered at the field level. Sample is the panel of 279 oilfields from 2006-2016 in columns (1)-(5) and the sample of only offshore fields in (6)-(7). Outcome variable is oil theft, the total number of sabotage spills within 15 km of the field. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. Oil prices are the de-measured annual average world price, in dollars per barrel. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 9: The effect of divestment on criminality by militant strength

Sample	All			Onshore	$d < 20$	$d > 20$
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Theft</i>						
Local operator	-3.354** (1.659)	-3.398* (1.807)	-3.698** (1.643)	-4.128** (1.876)	-4.599 (4.334)	-1.161 (1.429)
Local operator × Distance-weighted camps along pipeline	0.424 (0.532)					
Local operator × Distance-weighted allies along pipeline		0.712 (0.723)	0.620 (0.684)	1.072 (0.816)	1.274 (1.310)	-0.826 (0.683)
Local operator × Distance-weighted non-allies along pipeline		0.363 (0.966)				
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	No	Yes	Yes	Yes	Yes
Observations	3069	3069	3069	2277	1309	1760
R^2	0.720	0.723	0.756	0.750	0.772	0.798
<i>Panel B: Militancy</i>						
Local operator	-1.694** (0.760)	-1.015 (0.810)	-2.055*** (0.693)	-2.935*** (0.914)	-3.414* (1.862)	-0.581 (0.490)
Local operator × Distance-weighted camps along pipeline	0.370 (0.318)					
Local operator × Distance-weighted allies along pipeline		0.863* (0.486)	1.292** (0.543)	1.620** (0.681)	1.780** (0.854)	0.730 (0.533)
Local operator × Distance-weighted non-allies along pipeline		-0.692 (0.617)				
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	No	Yes	Yes	Yes	Yes
Observations	3069	3069	3069	2277	1309	1760
R^2	0.238	0.247	0.316	0.344	0.424	0.236

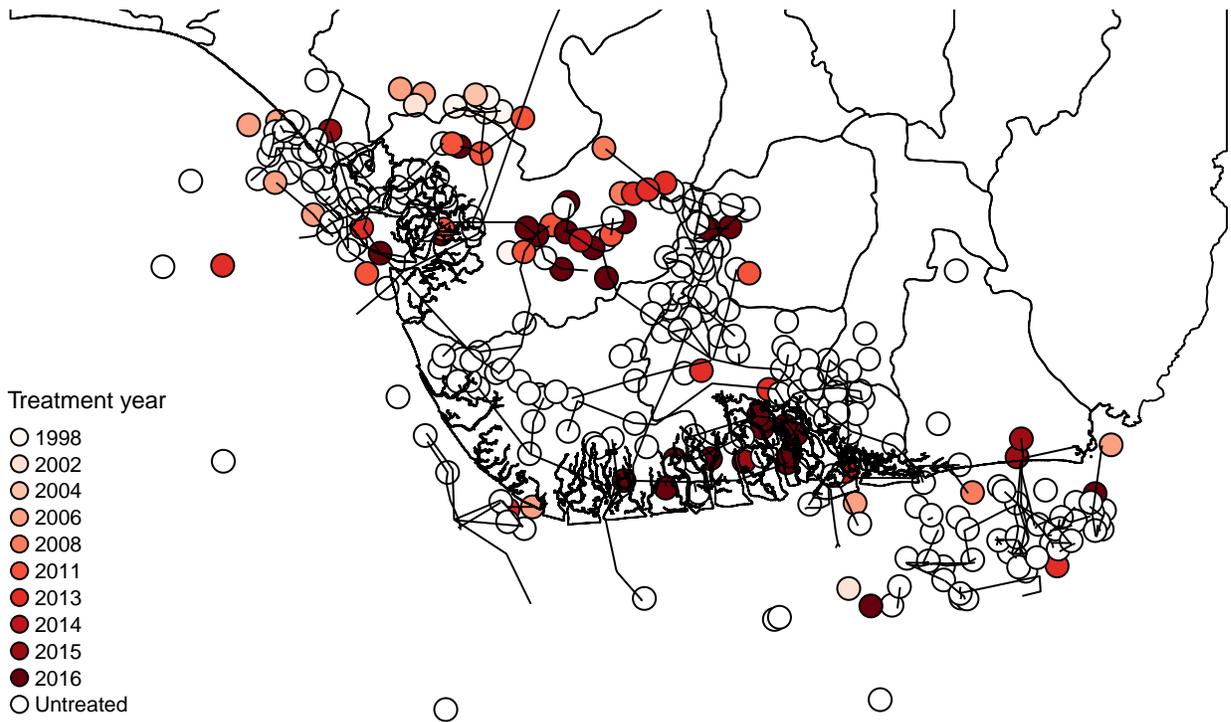
Standard errors, in brackets, are clustered at the field level. Sample is the panel of 279 oilfields from 2006-2016 in columns (1)-(3), the sample of only offshore fields in (4) and the sample of fields within (outside) of 20 kilometers of the nearest militant camp in (5)-(6). Outcome variable is indicated in panel headers. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. Distance-weighted connections along the pipeline is the number of allied/non-allied pipeline connections of the nearest militant camp, weighted by $1/\sqrt{distance}$. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 1: Indigenization and offshoring



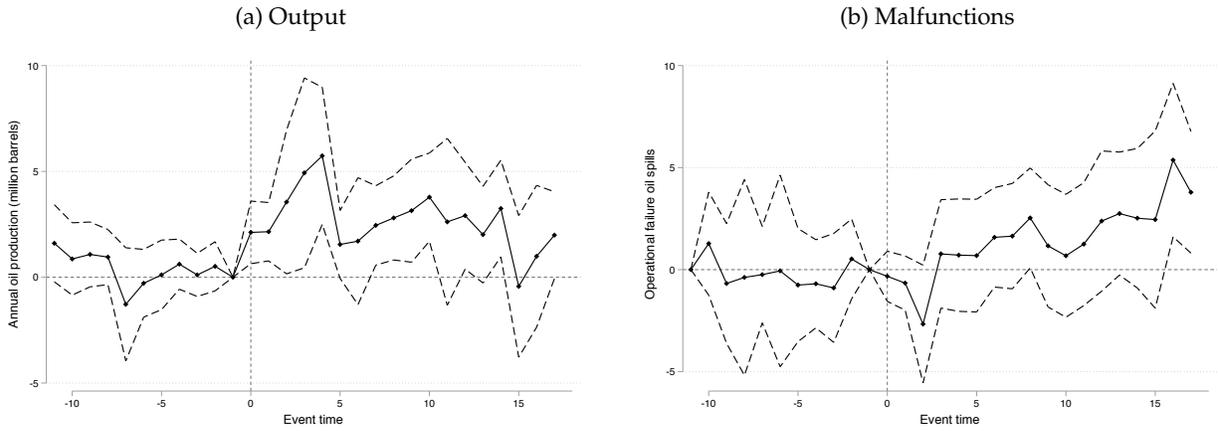
Note: Figure shows the share of total oil output produced onshore (Panel A), and by different companies (Panel B), over time. Onshore share is all output produced in onshore assets. Company categories are Addax, Agip, Chevron, Mobil, Shell, Total, the state-owned oil company, and independent private Nigerian companies, the latter of which is indicated in purple.

Figure 2: Map of treatment status



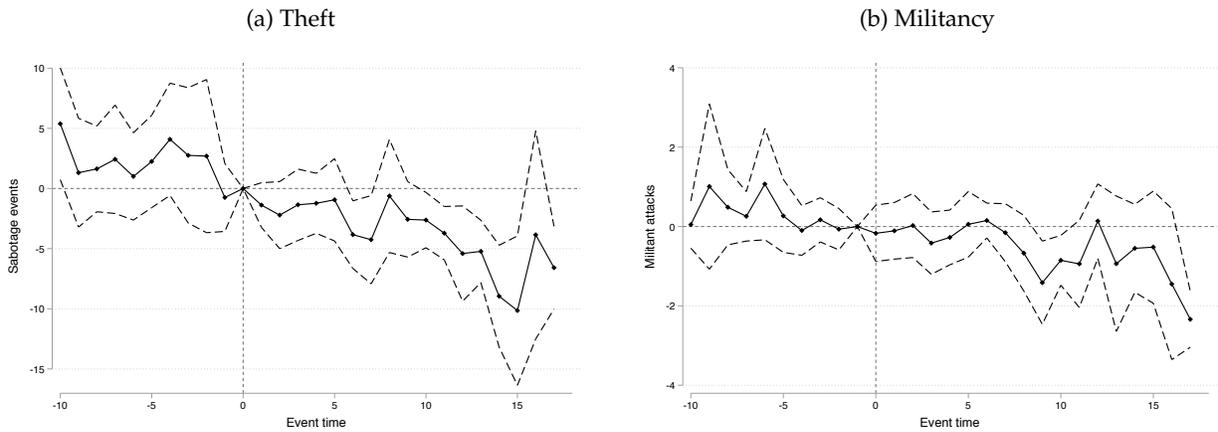
Note: Figure maps the centroids of 279 active Nigerian oilfields. Marker color indicates the year of local takeover of the field. White markers are never-treated fields. Basemap is Nigerian states of the Niger Delta region, while lines indicate oil pipelines.

Figure 3: Event study plots, technical



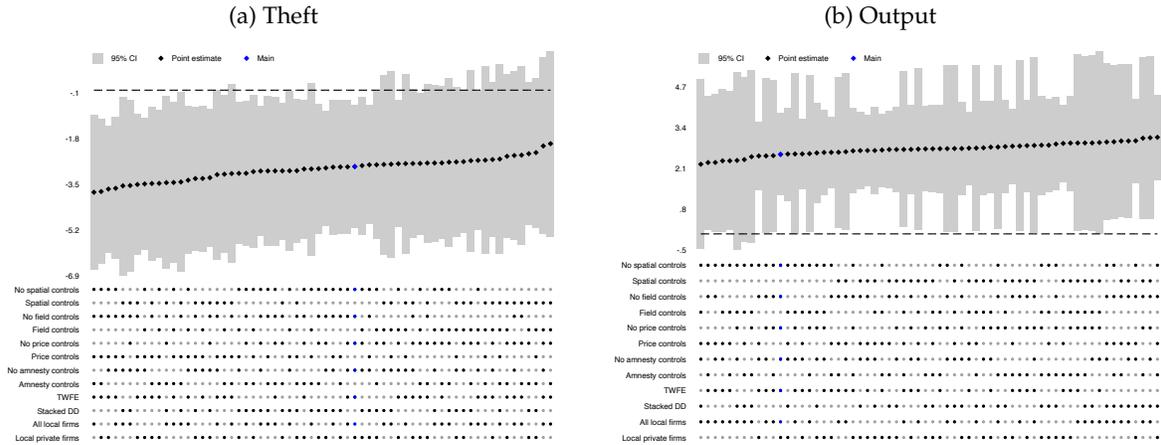
Note: Figure displays coefficients of event-study regressions of output (Panel A) and malfunctions (Panel B) on pre-and-post treatment indicators, conditional on unit and year fixed effect and controls interacted with year dummies. Malfunctions are the total number of non-sabotage spills within 15 km of the field. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. Sample is all nonmissing observations for the outcome in question.

Figure 4: Event study plots, criminality



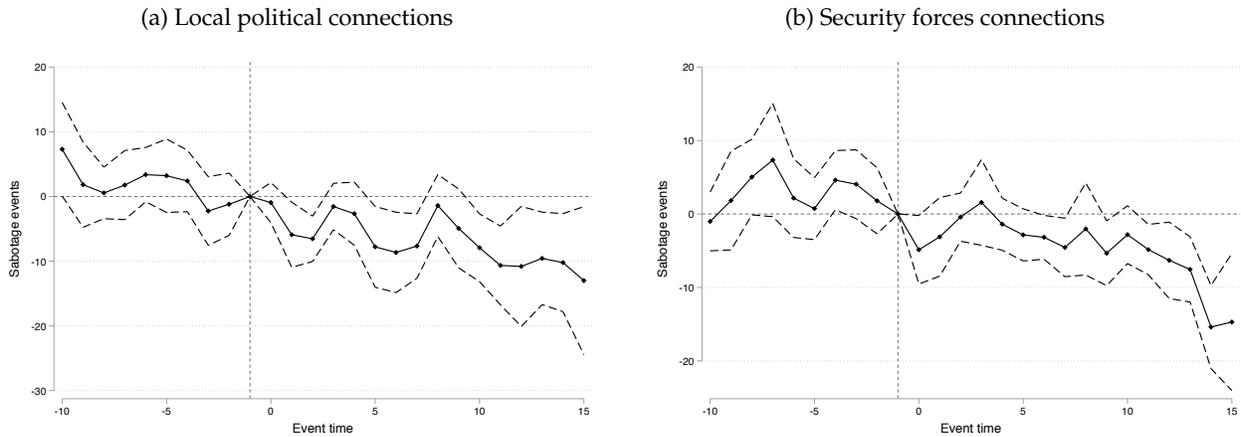
Note: Figure displays coefficients of event-study regressions of theft (Panel A) and militancy (Panel B) on pre-and-post treatment indicators, conditional on unit and year fixed effect and controls interacted with year dummies. Theft is the total number of sabotage spills within 15 km of the field. Militancy is the total number of militant attacks within 15 km of the field. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. Sample is all nonmissing observations for the outcome in question.

Figure 5: Robustness plots



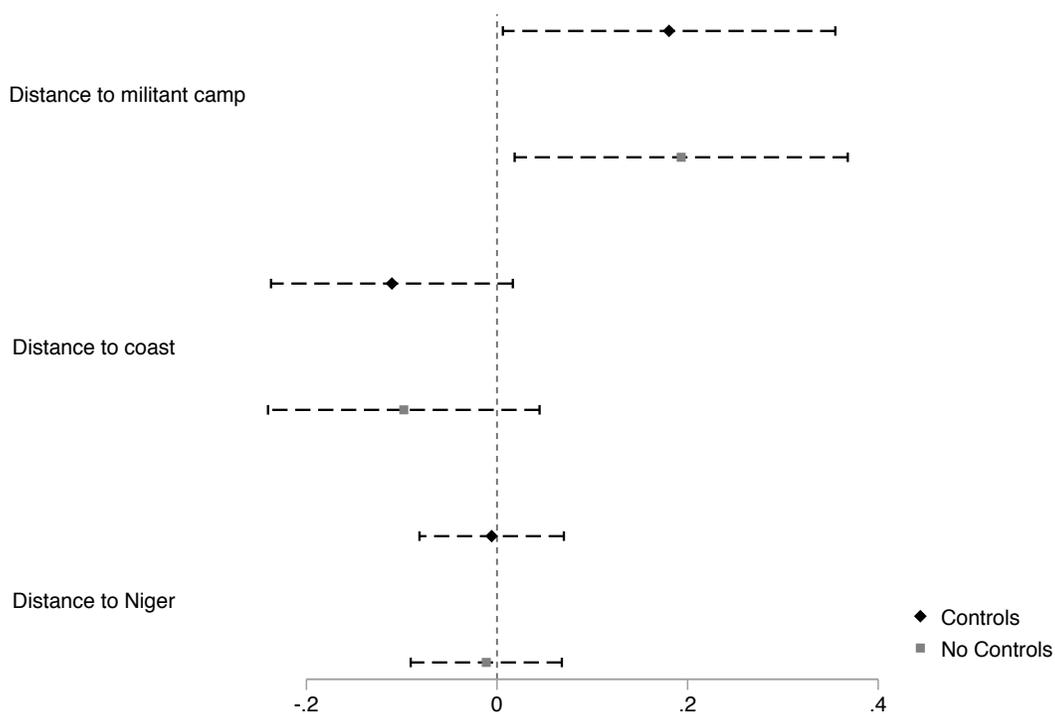
Note: Figure displays estimated coefficients on $local_{it}$ for robustness tests across specifications, for oil theft (Panel A) and oil production (Panel B) outcomes. Specification is indicated by points in the bottom of the figure. “Main” specification is that of Table 2, column (1). Spatial controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. Field controls are the number of wells, year of first well, mean well depth, and an onshore indicator. Price controls indicate inclusion of price interaction terms, as in Table A4. Amnesty controls indicate inclusion of controls for the 2009 Niger Delta amnesty. TWFE is the standard two-way fixed effects estimator, while Stacked DD is the specification in Table A7. Sample is all nonmissing observations for the outcome in question.

Figure 6: Event study plots, political connections



Note: Figure displays coefficients of event-study regressions of oil theft on pre-and-post treatment indicators for local political connections (Panel A) and security forces connections (Panel B). All estimates are conditional on unit and year fixed effect and controls interacted with year dummies. Theft is the total number of sabotage spills within 15 km of the field. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. Sample is all nonmissing observations for the outcome in question.

Figure 7: Localness and cost shifters



Note: Figure shows interaction coefficients from a regression of oil theft on localization interacted with cost shifters. Coefficients are Table 7, column (7). All distances are measured in kilometers. Sample is the panel of onshore oilfields from 2006-2016. Theft is the total number of sabotage spills within 15 km of the field. “Controls” estimates include year dummies interacted with controls, including latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km.

A Appendix tables

Table A1: Summary statistics by estimation sample

	Output (1)	Full (2)
Sabotage events	11.23 (25.01)	10.50 (22.81)
Operational failure oil spills	8.11 (12.49)	7.55 (11.73)
Piracy attacks	0.17 (0.92)	0.16 (0.85)
Militant deaths	1.36 (6.18)	1.65 (7.86)
Local operator	0.11 (0.32)	0.11 (0.32)
Field latitude	5.06 (0.63)	5.05 (0.62)
Distance to coast (km)	33.51 (29.51)	32.78 (29.15)
Distance to Niger River (km)	71.02 (68.97)	73.24 (67.35)
Distance to state capital (km)	88.25 (52.24)	85.54 (52.22)
Number of observations	2310	3069

Table displays means of variables with standard deviations in parentheses. “Output” sample in column (1) is the set of fields between 2006-2016 for which we have production information. Full sample is the full set of 279 fields between 2006-2016.

Table A2: The effect of divestment on revenue

Outcome	Revenue (millions of USD)				log(Revenue)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Local operator	226.421*** (57.064)	188.919*** (53.195)	204.116*** (53.957)	167.494*** (43.570)	0.953*** (0.233)	0.959*** (0.243)	0.948*** (0.246)	0.748*** (0.250)
Treated × Oil price (USD/barrel)			0.928 (0.752)				-0.001 (0.003)	
Field FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	Yes	No	No	Yes	Yes	No
Year × Field FE	No	No	No	Yes	No	No	No	Yes
Observations	2310	2310	2310	2310	1797	1797	1797	1797
R ²	0.821	0.843	0.843	0.873	0.766	0.780	0.780	0.851

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 279 oilfields from 2006-2016 for which output information is available. Outcome variable ins indicated in the table header. Revenue is measured as annual field output multiplied by annual average world oil prices. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A3: The effect of divestment on output and criminality, public and private

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Output and efficiency</i>						
Outcome	Shut-in		Output		Malfunctions	
Private local operator	-0.388*** (0.072)		2.462** (1.022)		-0.787 (0.933)	
Government operated		-0.007 (0.074)		2.119*** (0.669)		3.598*** (0.801)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2310	2310	2310	2310	3069	3069
R ²	0.665	0.656	0.854	0.854	0.631	0.632
<i>Panel B: Crime and violence</i>						
Outcome	Theft		Militancy		Piracy	
Private local operator	-3.399** (1.653)		-1.535** (0.683)		-0.192* (0.112)	
Government operated		-2.603 (1.883)		-1.266 (0.826)		0.086 (0.102)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3069	3069	3069	3069	3069	3069
R ²	0.754	0.754	0.288	0.287	0.316	0.315

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 279 oil-fields from 2006-2016. Output is measured in millions of barrels of oil per year. Private local operator is an indicator that the operator is a private Nigerian firm in a given field-year. Government operated is an indicator that the operator is the NPDC or NNPC in a given field-year. Malfunctions are the total number of non-sabotage spills within 15 km of the field. Theft is the total number of sabotage spills within 15 km of the field. Militancy is the total number of militant deaths within 15 km of the field. Piracy is the number of pirate attacks within 15 km of the field. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A4: The effect of divestment on output and criminality, robustness to prices

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Output and efficiency</i>						
Outcome	Shut-in		Output		Malfunctions	
Local operator	-0.185*** (0.062)	-0.183*** (0.061)	2.488*** (0.555)	2.216*** (0.488)	1.372* (0.746)	0.453 (0.717)
Treated × Oil price (USD/barrel)	0.000 (0.001)	0.000 (0.001)	-0.003 (0.012)	-0.013 (0.012)	-0.045*** (0.016)	-0.044*** (0.016)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	No	Yes	No	Yes
Observations	2310	2310	2310	2310	3069	3069
R ²	0.645	0.661	0.841	0.855	0.574	0.632
<i>Panel B: Crime and violence</i>						
Outcome	Theft		Militancy		Piracy	
Local operator	-3.674*** (1.205)	-4.193*** (1.524)	-0.734 (0.512)	-0.735 (0.559)	-0.040 (0.076)	-0.066 (0.078)
Treated × Oil price (USD/barrel)	-0.070* (0.036)	-0.078** (0.038)	0.056*** (0.014)	0.066*** (0.015)	0.003 (0.002)	0.002 (0.002)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	No	Yes	No	Yes
Observations	3069	3069	3069	3069	3069	3069
R ²	0.719	0.755	0.198	0.291	0.244	0.316

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 279 oilfields from 2006-2016. Output is measured in millions of barrels of oil per year. Oil prices are measured as the annual average world crude oil price in dollars per barrel. Malfunctions are the total number of non-sabotage spills within 15 km of the field. Theft is the total number of sabotage spills within 15 km of the field. Militancy is the total number of militant deaths within 15 km of the field. Piracy is the number of pirate attacks within 15 km of the field. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A5: The effect of divestment on output and revenue, measurement error

Sample Outcome	Single-operator				No shut-in	
	Q	$\log(Q)$	R	$\log(R)$	Q	R
	(1)	(2)	(3)	(4)	(5)	(6)
Local operator	3.336** (1.501)	0.948*** (0.363)	209.567 (155.515)	0.948*** (0.363)	1.741** (0.853)	142.446*** (42.558)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls \times Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2142	1653	2142	1653	1797	1797
R^2	0.845	0.779	0.832	0.784	0.880	0.864

Standard errors, in brackets, are clustered at the field level. Sample in columns (1)-(4) is the panel of single-operator field-years from 2006-2016. Sample in columns (5)-(6) is the panel of field-years with positive production from 2006-2016. Output is measured in millions of barrels of oil per year. Revenue is measured as annual field output multiplied by annual average world oil prices. Oil prices are measured as the annual average world crude oil price in dollars per barrel. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km.. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A6: Goodman-Bacon (2019) 2×2 DD weights

DD Comparison	Weight	DD Estimate
Earlier T vs. Later C	0.069	-4.875
Later T vs. Earlier C	0.022	-2.823
T vs. Never treated	0.832	-2.539
T vs. Already treated	0.076	0.376
Estimate	-2.486	

Table gives weights and estimates for all 2×2 DD comparisons, as derived by Goodman-Bacon (2019). Outcome variable is oil theft, the total number of sabotage spills within 15 km of the field.

Table A7: Stacked Differences-in-Differences estimates

Sample	Full			Ever-treated		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Theft</i>						
Local operator	-3.994*** (1.253)	-3.207** (1.252)	-2.552* (1.440)	-6.158*** (1.587)	-4.489** (1.744)	-4.363* (2.472)
Observations	26938	26938	26938	4058	4058	4057
R ²	0.699	0.722	0.722	0.561	0.599	0.602
<i>Panel B: Militancy</i>						
Local operator	-1.056** (0.436)	-1.285*** (0.476)	-1.651*** (0.588)	1.142** (0.526)	0.802 (0.507)	0.738 (0.654)
Observations	26938	26938	26938	4058	4058	4057
R ²	0.175	0.202	0.202	0.362	0.393	0.396
<i>Panel C: Piracy</i>						
Local operator	-0.057 (0.073)	-0.050 (0.075)	-0.068 (0.092)	0.217** (0.092)	0.142 (0.092)	0.152 (0.137)
Observations	26938	26938	26938	4058	4058	4057
R ²	0.231	0.260	0.260	0.259	0.323	0.327
<i>Panel D: Output</i>						
Local operator	3.576*** (1.285)	3.583*** (1.290)	3.102** (1.232)	3.748*** (1.337)	4.366*** (1.446)	4.982** (1.993)
Observations	20408	20408	20402	2718	2718	2708
R ²	0.865	0.865	0.867	0.547	0.555	0.632
<i>Panel E: Malfunctions</i>						
Local operator	1.061 (1.400)	1.393 (1.387)	1.180 (1.564)	0.291 (1.418)	0.814 (1.500)	0.682 (1.944)
Observations	26938	26938	26938	4058	4058	4057
R ²	0.546	0.577	0.578	0.636	0.657	0.667
<i>Panel F: Shut-ins</i>						
Local operator	-0.168** (0.069)	-0.179*** (0.068)	-0.162** (0.070)	-0.324*** (0.076)	-0.342*** (0.087)	-0.396*** (0.109)
Observations	20408	20408	20402	2718	2718	2708
R ²	0.667	0.678	0.680	0.461	0.539	0.551
Field FE	Yes	Yes	No	Yes	Yes	No
Event-time FE	Yes	Yes	No	Yes	Yes	No
Event-cohort FE	Yes	Yes	No	Yes	Yes	No
Calendar Year FE	No	Yes	No	No	Yes	No
Field-by-cohort FE	No	No	Yes	No	No	Yes
Event-time-by-cohort FE	No	No	Yes	No	No	Yes

Standard errors, in brackets, are clustered at the field-by-event-cohort level. Table presents results for the stacked-DD specification described in Section C.1. Columns (1)-(3) use the full sample while columns (4)-(6) restricts controls to be only ever-treated fields. Outcome variable is indicated in panel header. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A8: Political connections at the firm-level

	MNC	Local
Any politician	2	26
Any technocrat	2	12
Any elected politician	2	20
Any Niger Deltan politician	1	17
Any security forces member	0	8
Number of firms	5	31

Table displays counts of politically connected firms by type of connection, as well as the total number of firms, for multinational (MNC) and local firms. Sample is 36 firms for which political connections data on boardmembers, managers, and shareholders was available.

Table A9: The effect of political connections conditional on local ownership

Outcome	Theft					
	(1)	(2)	(3)	(4)	(5)	(6)
Local operator	-3.259** (1.414)	-3.870** (1.571)	-3.432** (1.324)	-2.981 (1.862)	-2.118 (1.757)	-2.378 (1.577)
Any politician		0.732 (1.297)				
Technocrat			0.343 (1.092)			
Elected politician				-0.390 (1.358)		
Same-state politician					-2.399 (1.921)	
Security forces						-4.013 (2.556)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls \times Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2941	2941	2941	2941	2941	2941
R^2	0.755	0.755	0.755	0.755	0.755	0.755

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 279 oilfields from 2006-2016 for which political connections data is available. Outcome variable is oil theft, the total number of sabotage spills within 15 km of the field. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. Political connections variables are dummy variables indicating that the operator of a given field-year has a particular type of politician as a board member, shareholder, or manager. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A10: Local ownership and local employment

Sample	All				Onshore	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Employed</i>						
Local operator	-0.016 (0.018)	-0.021 (0.017)	0.002 (0.022)	0.007 (0.020)	-0.018 (0.018)	-0.001 (0.025)
Observations	16518	16518	16518	16518	16017	16017
R^2	0.033	0.034	0.042	0.038	0.033	0.043
<i>Panel B: Employed outside home</i>						
Local operator	-0.020 (0.021)	-0.018 (0.020)	-0.009 (0.022)	-0.007 (0.017)	-0.028 (0.020)	-0.017 (0.018)
Observations	9040	9040	9040	9040	8749	8749
R^2	0.108	0.110	0.179	0.133	0.097	0.160
<i>Panel C: Self-employed</i>						
Local operator	0.080*** (0.024)	0.082*** (0.025)	0.068** (0.030)	0.093*** (0.030)	0.084*** (0.024)	0.074** (0.034)
Observations	9040	9040	9040	9040	8749	8749
R^2	0.073	0.074	0.138	0.107	0.072	0.137
<i>Panel D: Employed in household agriculture</i>						
Local operator	-0.032 (0.040)	-0.036 (0.043)	-0.002 (0.045)	-0.034 (0.031)	-0.032 (0.040)	-0.008 (0.045)
Observations	9040	9040	9040	9040	8749	8749
R^2	0.151	0.152	0.274	0.181	0.153	0.282
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	No	Yes	No	No	No	No
Year \times State FE	No	No	No	Yes	No	No
Controls \times Year FE	No	No	Yes	No	No	Yes

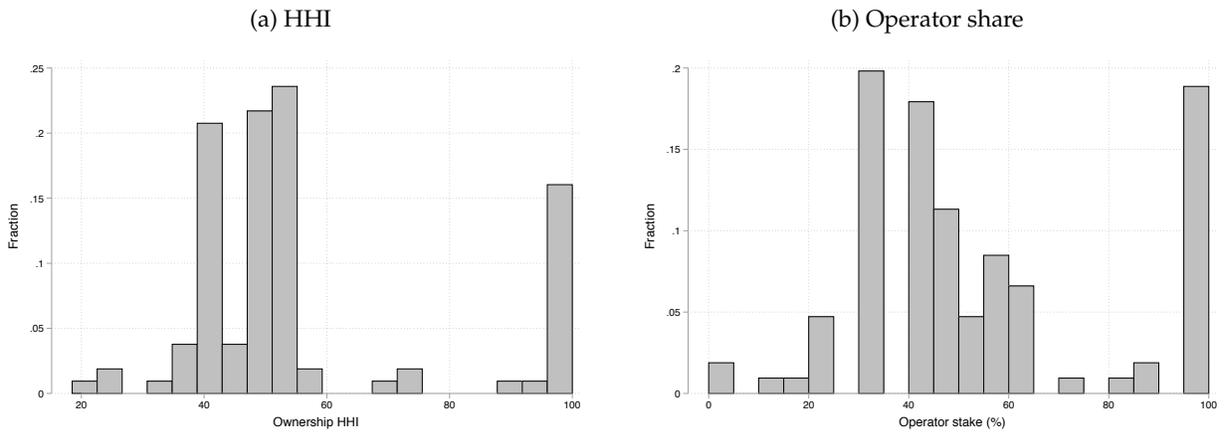
Standard errors clustered at the field level in brackets. Outcome variable is given in the panel header. Sample is all individuals in the three waves of the GHS between the ages of 15-60 living in clusters within 50 km of an oilfield. All regressions use household-level sampling weights. GHS controls are cluster distance to road, population center, market, border, and administrative center, a rural dummy, slope, elevation, and mean annual temperature and precipitation. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A11: Local ownership and local consumption

Outcome Sample	log(consumption)					
	All				Onshore	
	(1)	(2)	(3)	(4)	(5)	(6)
Local operator	0.089 (0.107)	0.091 (0.105)	0.079 (0.073)	-0.003 (0.056)	0.087 (0.105)	0.044 (0.063)
Observations	5019	5019	5019	5019	4877	4877
R^2	0.246	0.247	0.298	0.271	0.254	0.310
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	No	Yes	No	No	No	No
Year \times State FE	No	No	No	Yes	No	No
Controls \times Year FE	No	No	Yes	No	No	Yes

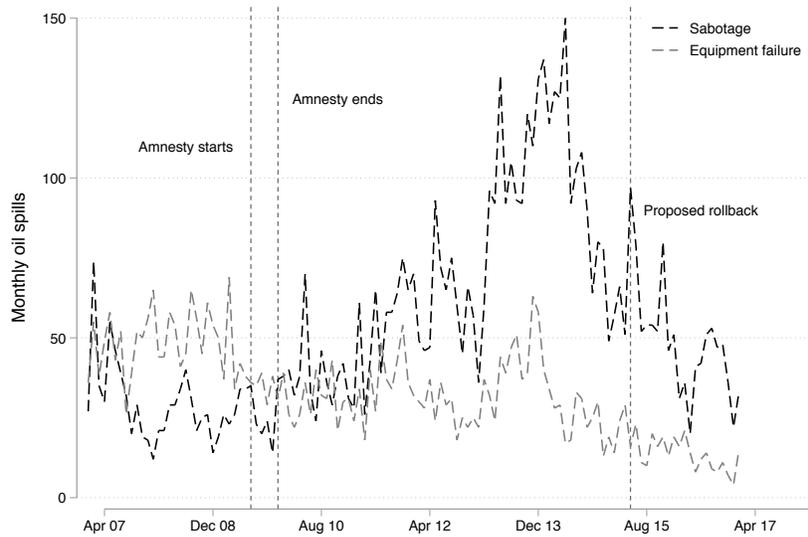
Standard errors clustered at the field level in brackets. Outcome variable is the log of per capita household consumption. Sample is all households in the three waves of the GHS living in clusters within 50 km of an oilfield. All regressions use household-level sampling weights. GHS controls are cluster distance to road, population center, market, border, and administrative center, a rural dummy, slope, elevation, and mean annual temperature and precipitation. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure A1: Ownership mixes in 2016



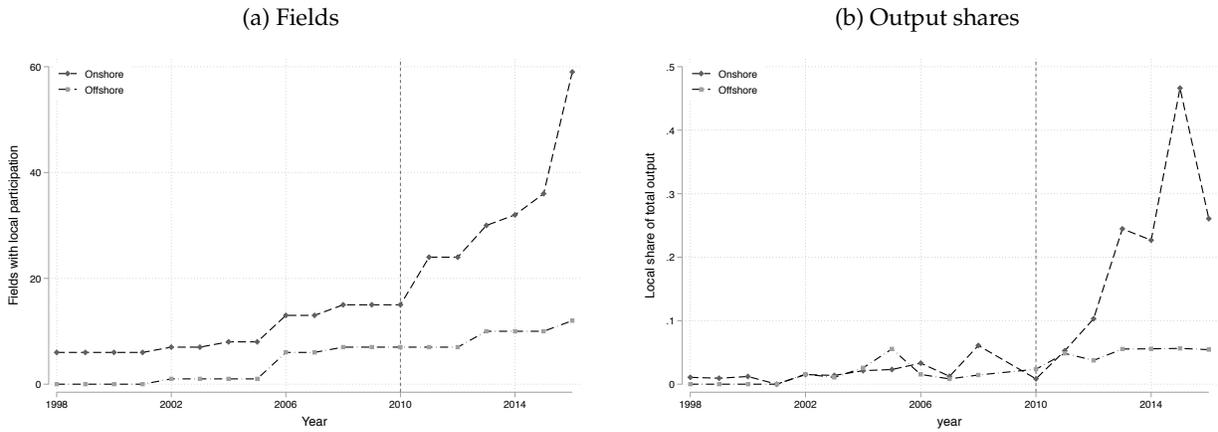
Note: Figure shows histograms of ownership concentration, measured as the Herfindahl index (Panel A), and the stake owned by the operating company (Panel B). Sample is a cross-section of 106 active oil blocks (licenses) in 2016.

Figure A2: Pipeline sabotage and operational malfunction over time



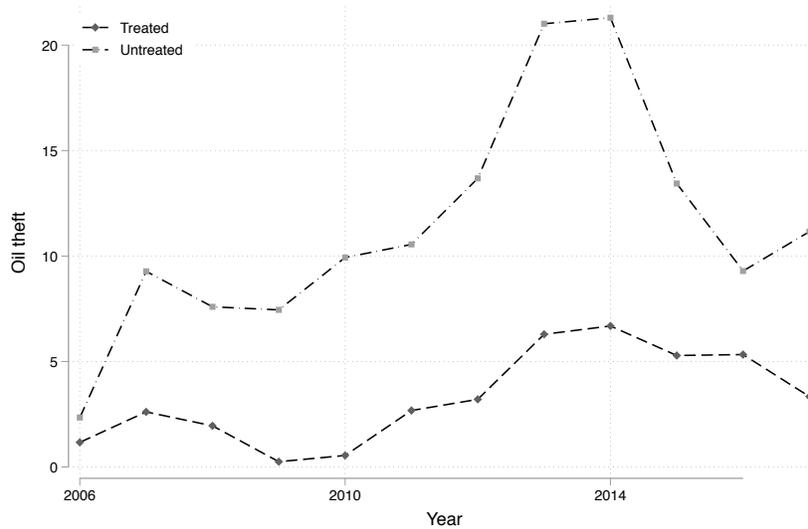
Note: Figure shows monthly totals of oil spills due to sabotage and non-sabotage (equipment failure) over time. Data come from 11,587 oil spills recorded by the NOSDRA OSM from 2006-2017. Vertical lines indicate the beginning of the federal amnesty program for ex-combatants, the end of the initial amnesty period, as well as the proposed rollback of amnesty benefits.

Figure A3: Indigenization



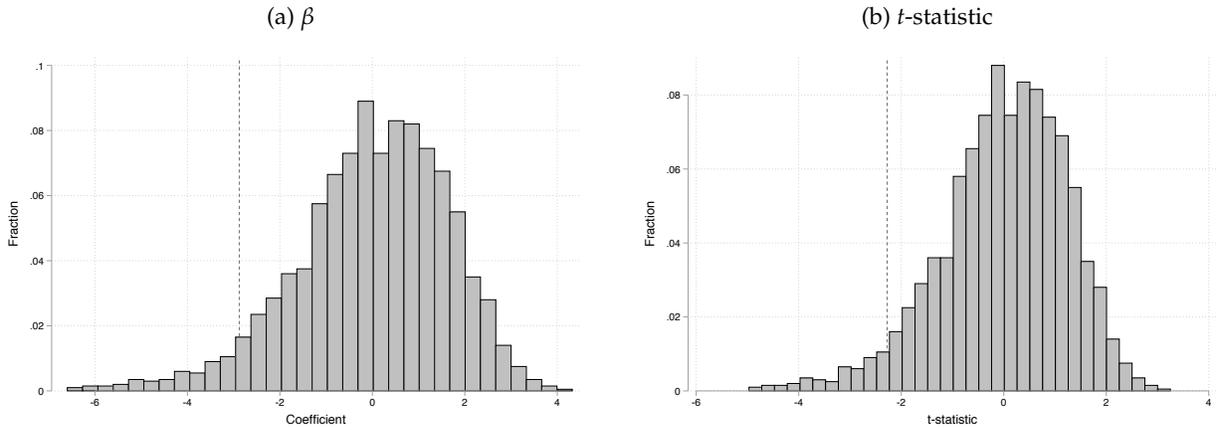
Note: Figure shows number of fields (Panel A) and output share (Panel B) of local Nigerian operators over time by type of asset (onshore vs. offshore). Vertical line indicates the 2010 passage of the Nigerian Local Content Act. Sample is an unbalanced panel of 279 oilfields from 1998-2016. Oil production data and output shares are missing for 2009.

Figure A4: Theft over time



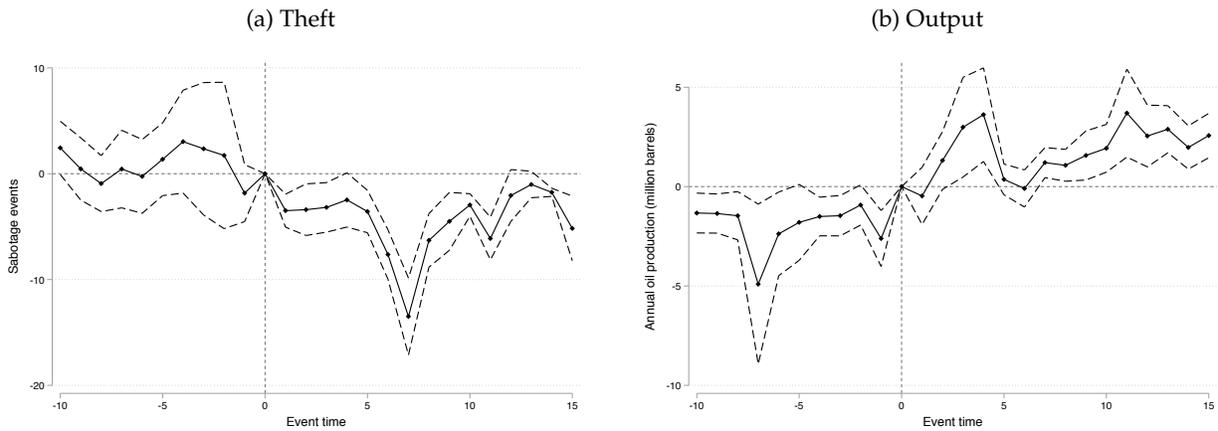
Note: Figure shows mean annual field-level sabotage incidents over time for a sample of 71 ever-treated and 208 never-treated oilfields.

Figure A5: Randomization inference



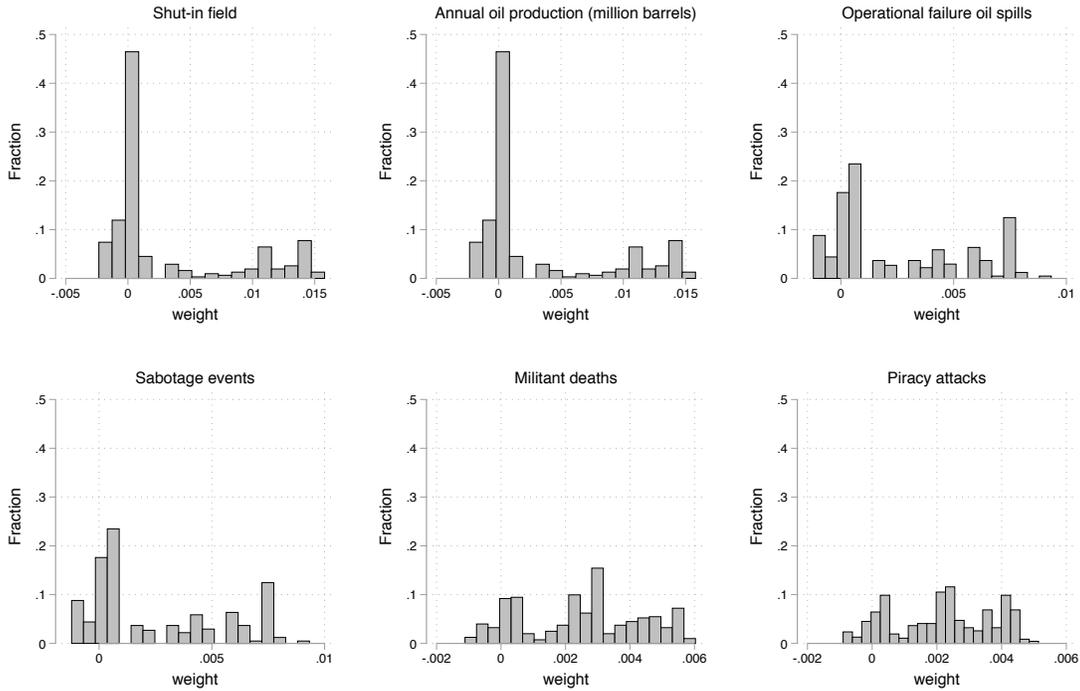
Note: Figure shows histograms of coefficient estimates (Panel A) and t -statistics (Panel B) for 2000 draws of a randomization inference routine. Outcome variable is theft, the total number of sabotage spills within 15 km of the field. Vertical line indicates the estimate for the observed data.

Figure A6: Re-weighted event-study



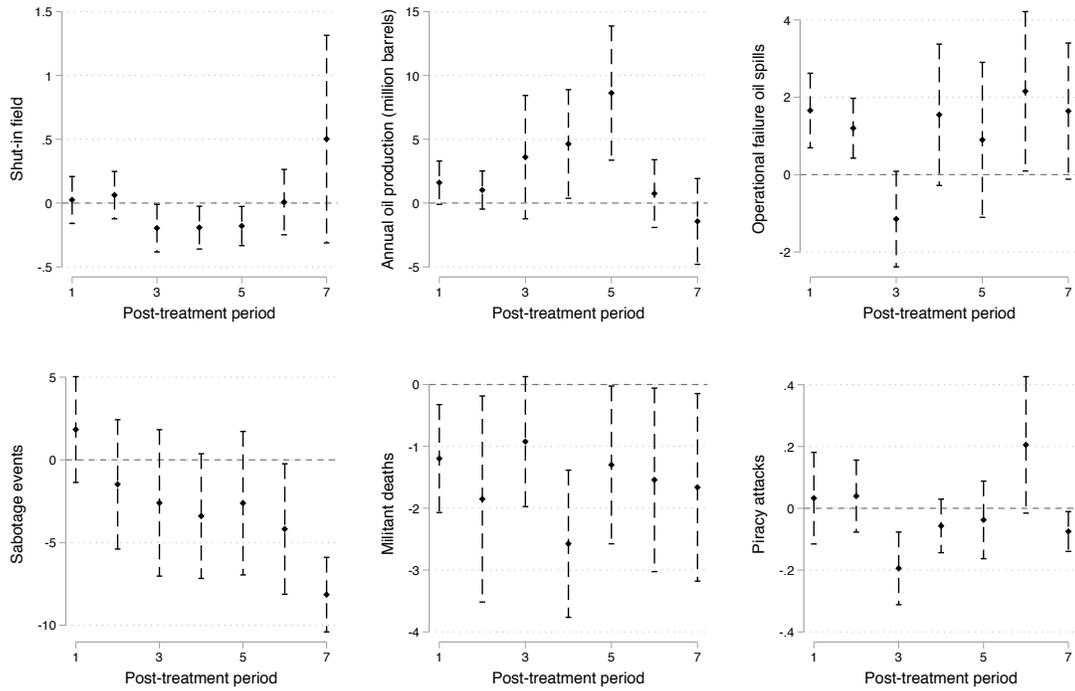
Note: Figure shows re-weighted coefficients from the cohort-specific event-study regression described in Abraham and Sun (2018). To obtain these point estimates, I estimate the fully-saturated event-study regression of the outcome of interest on pre-and-post-treatment dummies are interacted with dummies indicating the cohort of treatment, as well as unit and time fixed effects. I then weight these cohort-specific event-study estimates by the cohort share among the treated group in a given event-period. Standard errors are clustered at the field level and calculated using the delta method for a linear transformation of cohort-specific per-period effects.

Figure A7: Histogram of weights from de Chaisemartin and D’Haultfoeuille (2019)



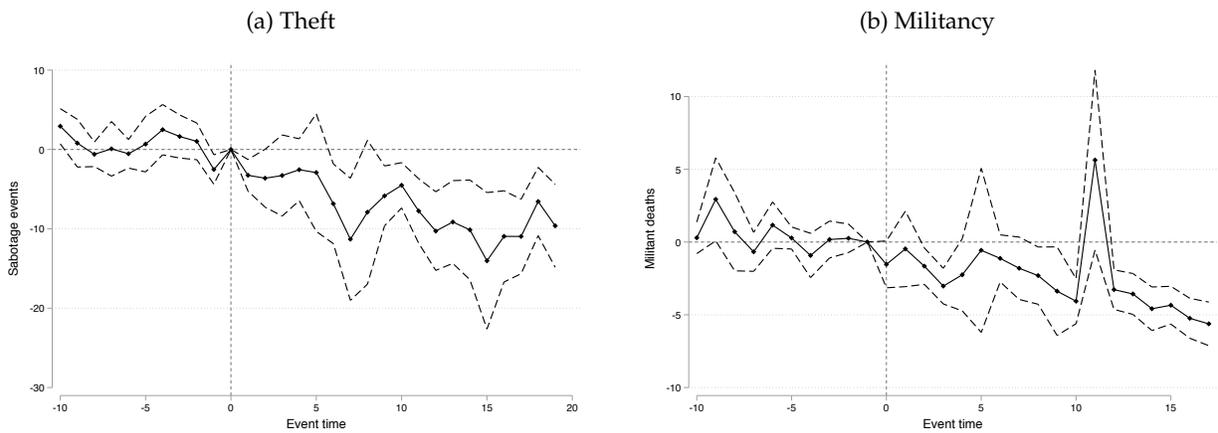
Note: Figure shows implied weights for unit-and-time-specific average treatment effects, as derived in the de Chaisemartin and D’Haultfoeuille (2019) decomposition results. I display histograms of the weights for each of the six key outcomes analyzed in Table 2. Sample is the panel of 279 oilfields from 2006-2016. Shut-in is defined as a field registering zero output in a given year. Output is measured in millions of barrels of oil per year. Malfunctions are the total number of non-sabotage spills within 15 km of the field. Theft is the total number of sabotage spills within 15 km of the field. Militancy is the total number of militant deaths within 15 km of the field. Piracy is the number of pirate attacks within 15 km of the field.

Figure A8: Dynamic effects from de Chaisemartin and D’Haultfoeuille (2019)



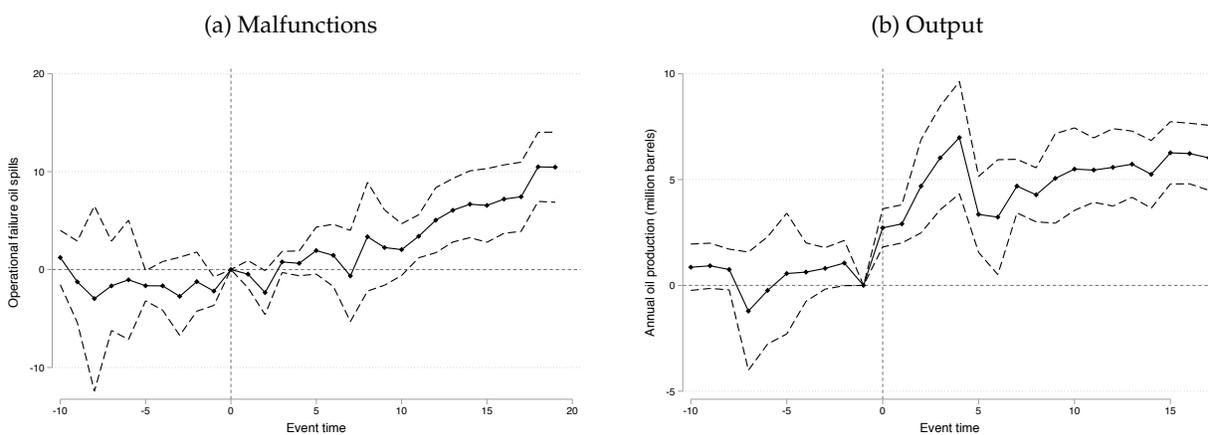
Note: Figure shows dynamic effects among switchers using the Wald-type estimator for staggered adoption designs described in de Chaisemartin and D’Haultfoeuille (2019). Standard errors are clustered at the field level and computed using a bootstrapping procedure with 1000 replications.

Figure A9: Stacked-DD event-study, crime



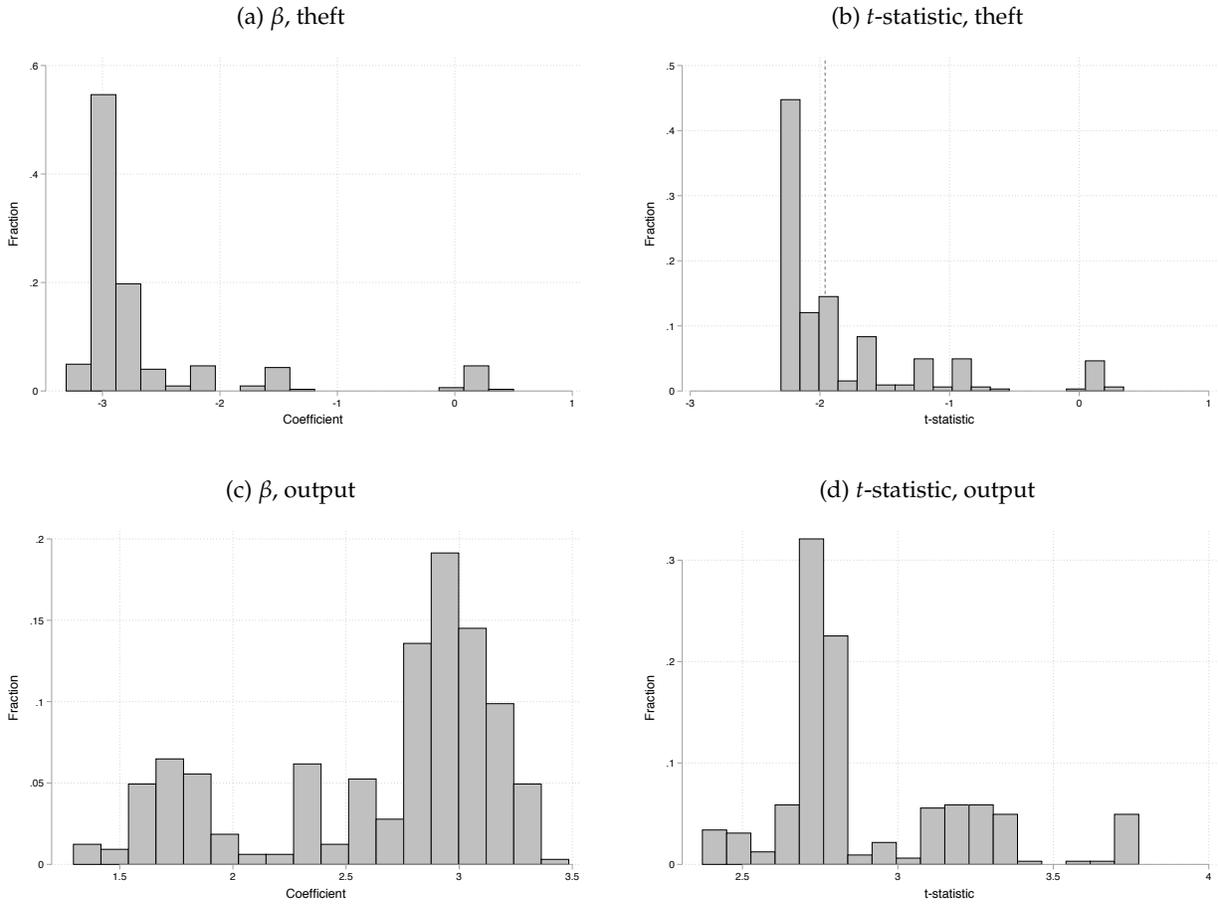
Note: Figure shows coefficients from event-study regressions of the stacked-DD specification described in Section C.1 for crime and violence outcomes. Standard errors are clustered at the field-by-event-cohort level. Theft is the total number of sabotage spills within 15 km of the field. Militancy is the total number of militant deaths within 15 km of the field.

Figure A10: Stacked-DD event-study, output



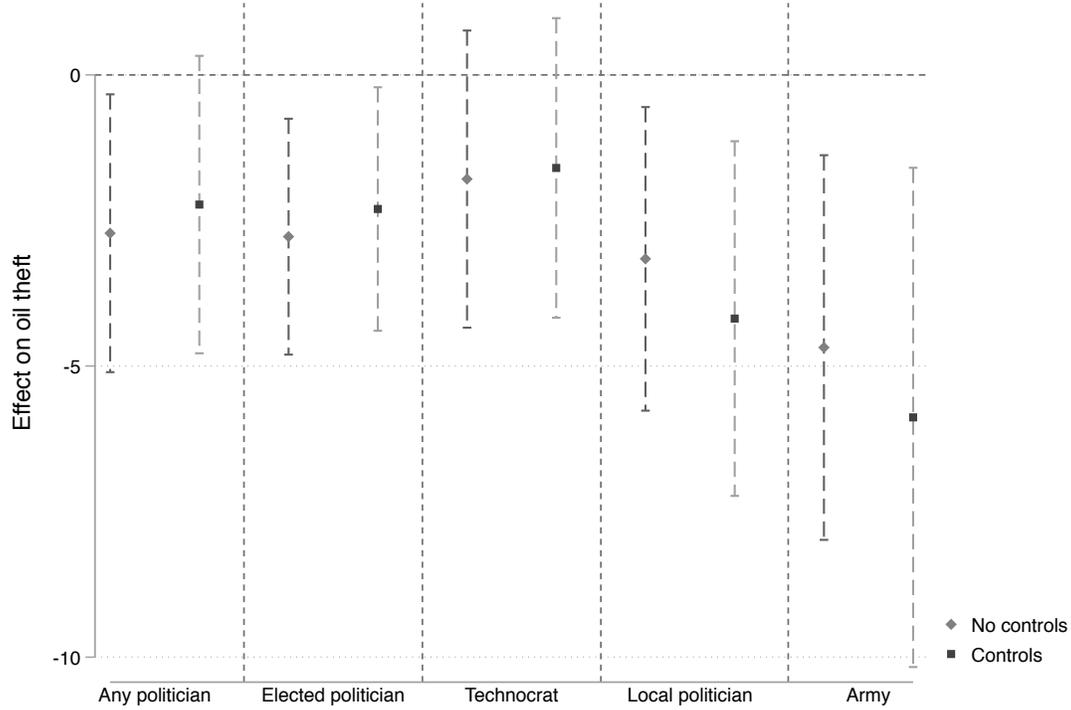
Note: Figure shows coefficients from event-study regressions of the stacked-DD specification described in Section C.1 for oil production outcomes. Standard errors are clustered at the field-by-event-cohort level. Output is measured in millions of barrels of oil per year. Malfunctions are the total number of non-sabotage spills within 15 km of the field.

Figure A11: Stacked-DD histogram over event-windows



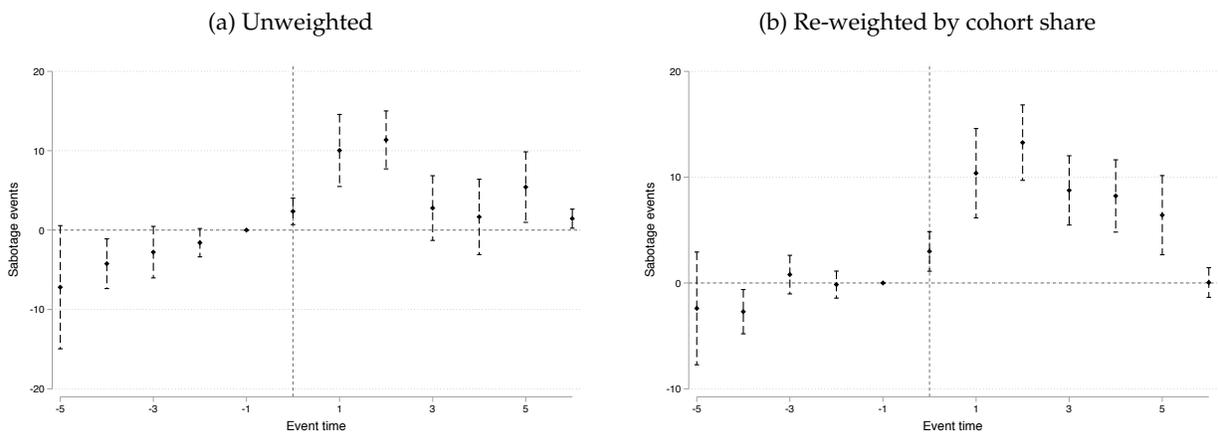
Note: Figure shows histograms of coefficients and t -statistics from the stacked-DD specification described in Section C.1 for oil production and theft outcomes. For each outcome, I estimate treatment effects for all possible combinations of event windows up to 18 years before and 18 years after the event and then plot these estimates. Standard errors are clustered at the field-by-event-cohort level. Output is measured in millions of barrels of oil per year. Theft is the total number of sabotage spills within 15 km of the field.

Figure A12: Political connections and theft



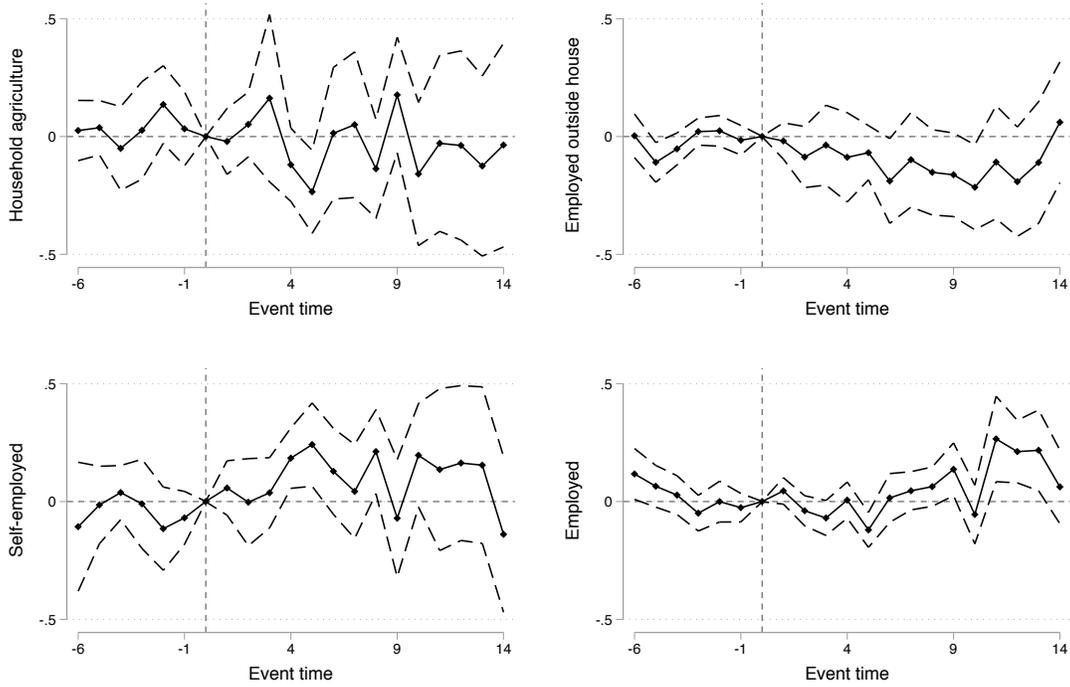
Note: Figure plots the estimates from Table 5. Sample is the panel of 279 oilfields from 2006-2016 for which political connections data is available. Outcome variable is oil theft, the total number of sabotage spills within 15 km of the field. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. Political connections variables are dummy variables indicating that the operator of a given field-year has a particular type of politician as a board member, shareholder, or manager.

Figure A13: Oil theft and anti-corruption laws



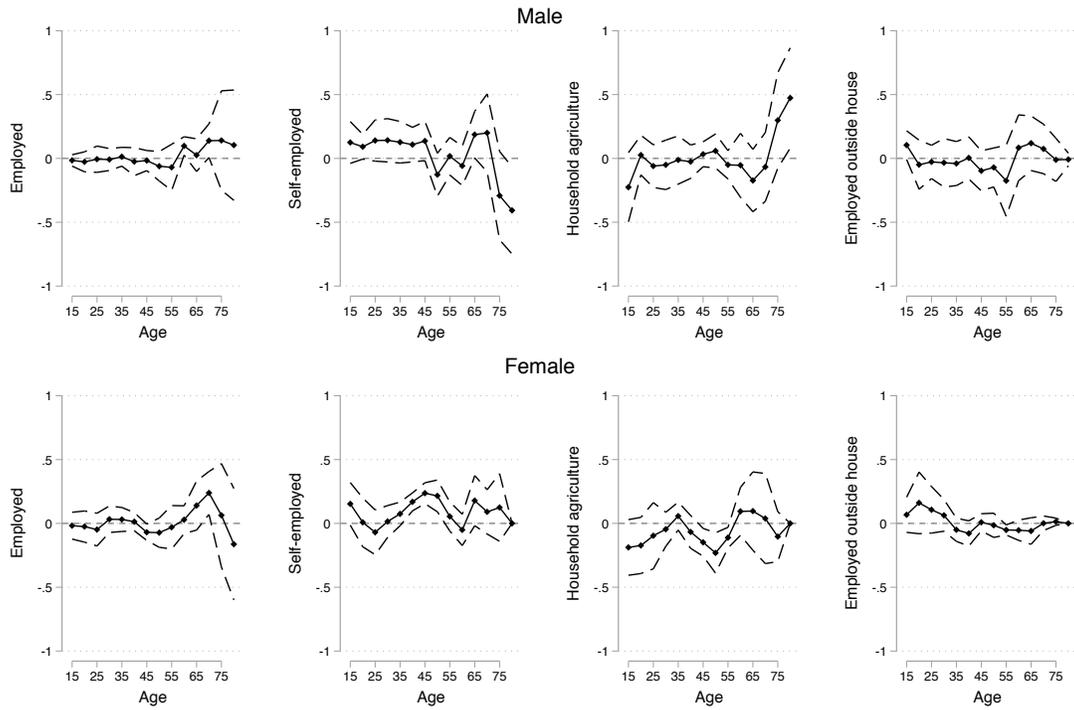
Note: Figure shows coefficients from an event-study regression of oil theft on indicators for before and after the passage of a home-country anti-corruption law, as well as time and unit fixed effects. Sample is the panel of oilfields from 2006-2016 where the operator is a multinational corporation. Outcome variable is oil theft, the total number of sabotage spills within 15 km of the field. Panel A shows unweighted event-study estimates while Panel B shows estimates re-weighted by cohort shares following Abraham and Sun (2018).

Figure A14: Local ownership and local employment, parallel trends



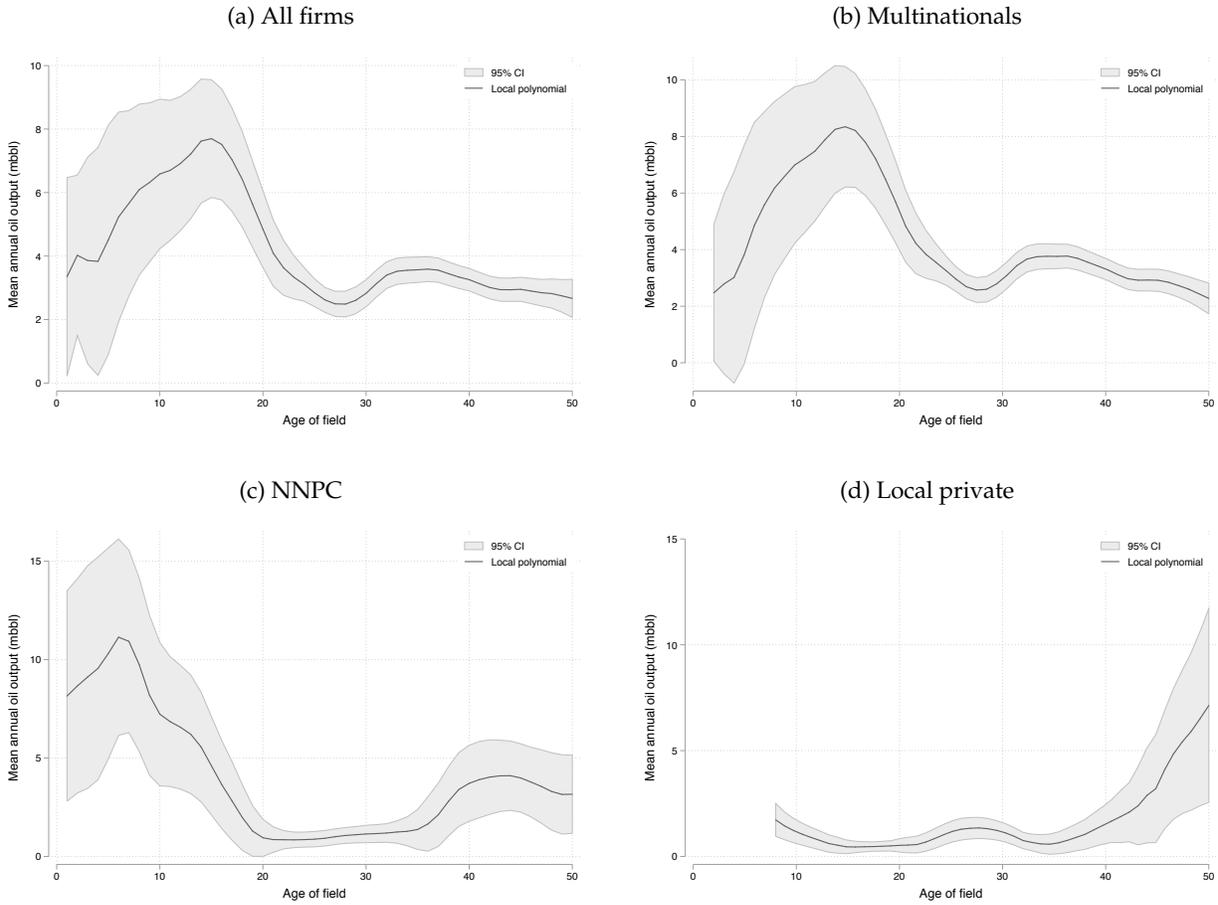
Note: Figure displays coefficients of event-study regressions of employment outcomes on pre-and-post treatment indicators for localization, conditional on unit and year fixed effect and controls interacted with year dummies. Employment outcomes are given in each subfigure. Sample is all individuals in the three waves of the GHS between the ages of 15-60 living in clusters within 50 km of an oilfield. All regressions use household-level sampling weights. GHS controls are cluster distance to road, population center, market, border, and administrative center, a rural dummy, slope, elevation, and mean annual temperature and precipitation

Figure A15: Local ownership and local employment by age and gender



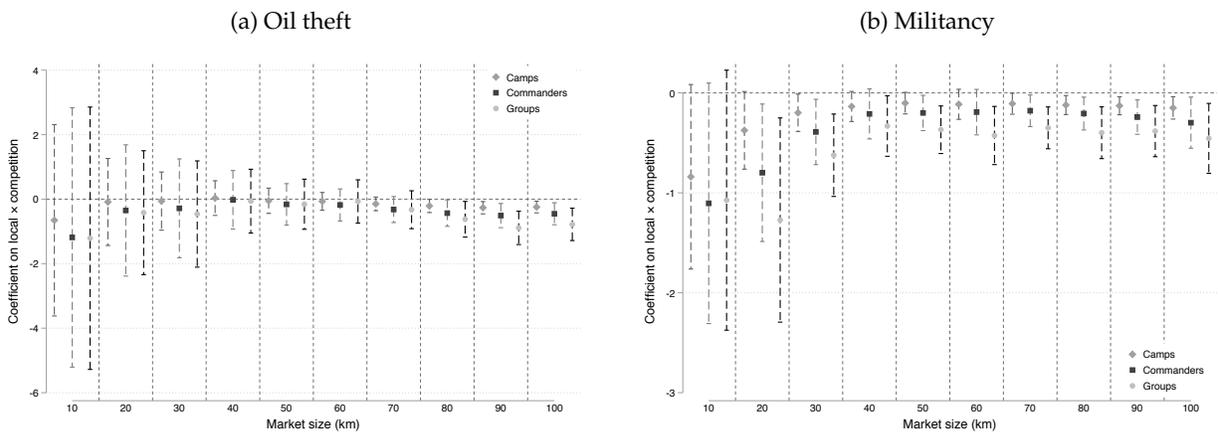
Note: Figure shows coefficients from differences-in-differences regressions of employment outcomes on local ownership of the nearest oilfield. Sample is all individuals in the three waves of the GHS above the age of 10 living in clusters within 50 km of an oilfield. All regressions use household-level sampling weights. Each point-estimate corresponds to a DD estimate for a particular gender-age subsample, as indicated in the plot. X-axis numbers indicate the midpoint of a ten-year age grouping (i.e. 15 corresponds to the 10-20 age bin). Standard errors are clustered at the field level.

Figure A16: Extraction profiles



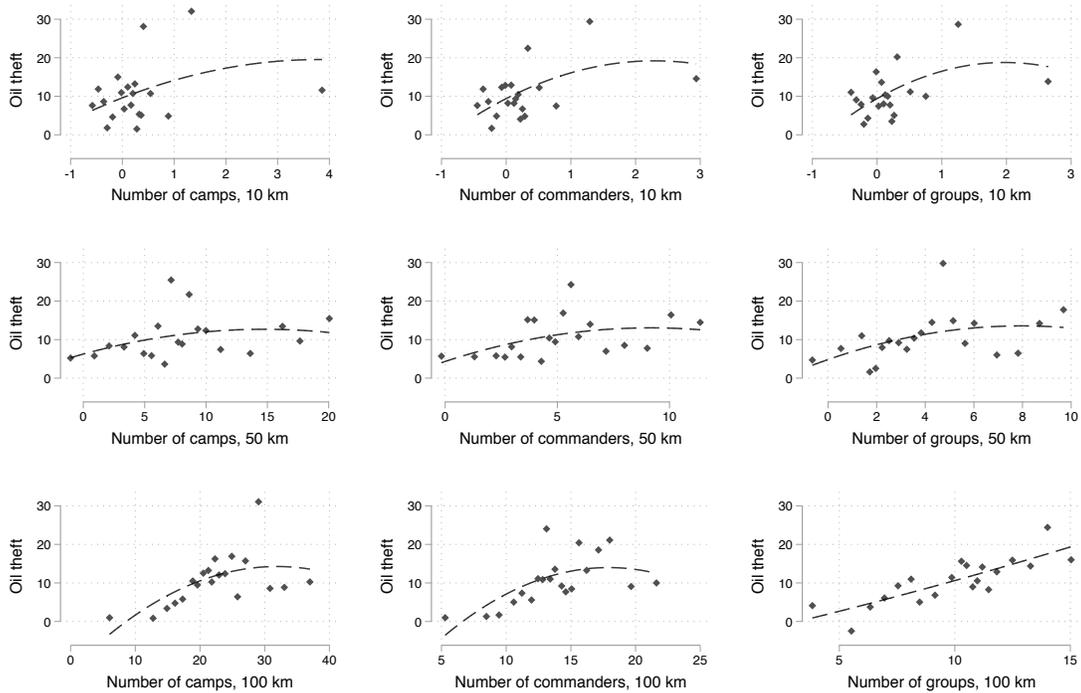
Note: Figure displays oil extraction profiles for the whole sample of fields (Panel A), multinational-operated fields (Panel B), NNPC fields (Panel C), and private locally-operated fields (Panel D). Age of field is defined as the current year minus the year when of completion of the first well.

Figure A17: Interaction effects on competition measures



Note: Figure plots interaction coefficients from a regression of oil theft (Panel A) or militancy (Panel B) on the localization treatment and its interactions with the number of camps, commanders, or groups within a given radius around the field. All regressions control for unit and time fixed effects, controls interacted with time dummies, and the interaction between localization and the number of distance-weighted pipeline allies and non-allies of the nearest militant group. Each estimate varies the radius around the camp and/or the level at which gangs are defined (at the camp, commander, or group level). Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km.

Figure A18: Quantity correlations



Note: Figure shows binned scatterplots relating oil theft to the number of camps, commanders, or groups within a given radius around the field. Line is a quadratic fit of the underlying data. Each scatterplot varies the radius around the camp and/or the level at which gangs are defined (at the camp, commander, or group level). All scatterplots are cross-sectional regressions that control for year fixed effects latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km.

B Theoretical Appendix

B.1 Cost of corruption

$\bar{b}_f > 0$ implies that

$$\frac{\gamma_f p(q + \kappa) - \lambda_f}{1 + \tau_f} > 0$$
$$\lambda_f < \gamma_f p(q + \kappa)$$

While at the same time $\bar{b}_f < pq$ gives

$$\frac{\gamma_f p(q + \kappa) - \lambda_f}{1 + \tau_f} < pq$$
$$\gamma_f p(q + \kappa) - \lambda_f < pq(1 + \tau_f)$$
$$\gamma_f p(q + \kappa) - pq(1 + \tau_f) < \lambda_f$$
$$\lambda_f > \gamma_f p\kappa - pq(1 + \tau_f - \gamma_f)$$

Yielding

$$\lambda_f \in [\min\{0, \lambda_f > \gamma_f p\kappa - pq(1 + \tau_f - \gamma_f)\}, \gamma_f p(q + \kappa)]$$

Note that for theft to occur with positive probability, we must have $Pr(\neg B) > 0$, which gives

$$pq \left(1 - \frac{\gamma_f}{1 + \tau_f}\right) > \left(\frac{\gamma_f p\kappa - \lambda_f}{1 + \tau_f}\right)$$
$$(1 + \tau_f)pq - pq\gamma_f > \gamma_f p\kappa - \lambda_f$$
$$\lambda_f > \gamma_f p\kappa - pq(1 + \tau_f - \gamma_f)$$

This is exactly the second part of A2, which assumes that at least the highest type gangsters are too expensive to bribe, implying that theft can occur in equilibrium.

B.2 Proof of Proposition 1

Proof:

$$\begin{aligned}
\frac{\partial Pr(B)}{\partial \lambda} &= -\frac{1}{(1+\tau_f)c} < 0 \\
\frac{\partial Pr(B)}{\partial \tau} &= -\frac{(\gamma_f p(q+\kappa) - \lambda_f)}{(1+\tau)^2 c} < 0 \\
\frac{\partial Pr(B)}{\partial q} &= \frac{p}{c} \left(\frac{\gamma_f}{1+\tau_f} - 1 \right) < 0 \\
\frac{\partial Pr(B)}{\partial c} &= -\frac{1}{c^2} \left[\left(\frac{\gamma_f p \kappa - \lambda_f}{1+\tau_f} \right) - pq \left(1 - \frac{\gamma_f}{1+\tau_f} \right) \right] > 0 \\
\frac{\partial Pr(B)}{\partial \gamma} &= \frac{p(q+\kappa)}{(1+\tau_f)c} > 0 \\
\frac{\partial Pr(B)}{\partial \kappa} &= \frac{\gamma_f p}{(1+\tau_f)c} > 0 \\
\frac{\partial Pr(B)}{\partial p} &= \frac{1}{c} \left(\frac{\gamma_f(q+\kappa)}{1+\tau_f} - q \right) \text{ is ambiguous}
\end{aligned}$$

The second is negative by A2.1, the third is negative since $\gamma_f < 1$ and $1 + \tau_f > 1$, the fourth is positive by A2.2. Lastly, $\frac{\partial Pr(B)}{\partial p} > 0$ whenever $\frac{\kappa}{q} > \frac{(1+\tau_f)}{\gamma_f} - 1$ and negative otherwise. If losses are high relative to theft, then an increase in price affects the company's reservation price relatively more than the gangster's, increasing \bar{b}_f and expanding the bargaining range. If the opposite is true, then the bargaining range contracts because \underline{b}_g rises relatively more. Note that under perfect bargaining, where $\tau_f = 0$ and $\gamma_f = 1$, the inefficiency of theft implies that $\frac{\partial Pr(B)}{\partial p} > 0$ is always true. For a given increase in p , gangsters increase \underline{b}_g by q while the company increases b $q + \kappa$.

B.3 Proof of Proposition 2

Proof: Denote the parameter vector by $\theta = (\gamma, \lambda, \tau, q, \kappa, c)$. Define R as the right-hand side of the shut-in equation.

$$R = Pr(B)[(1+\tau_f)b(\theta) + \lambda_f] + Pr(\neg B)\gamma_f p(q+\kappa)$$

Where, $b(\theta) = E[\underline{b}_g | \epsilon_g \in B] = \frac{\bar{b}_f - pq + c}{2} = \frac{Pr(B)c}{2}$. We want to know the sign of $\frac{\partial R}{\partial \theta_i}$ with respect to a parameter i , e.g., does it increase or decrease the likelihood of shut-in for a given level of variable profits.

With respect to λ :

$$\begin{aligned}
\frac{\partial R}{\partial \lambda} &= \frac{\partial Pr(B)}{\partial \lambda} [(1 + \tau_f)b(\theta) + \lambda] + Pr(B)[(1 + \tau_f)\frac{\partial b}{\partial \lambda} + 1] - \frac{\partial Pr(B)}{\partial \lambda} \gamma_f p(q + \kappa) \\
&= \frac{\partial Pr(B)}{\partial \lambda} [(1 + \tau_f)\frac{Pr(B)c}{2} + \lambda] + \frac{Pr(B)}{2} - \frac{\partial Pr(B)}{\partial \lambda} \gamma_f p(q + \kappa) \\
&= \frac{\partial Pr(B)}{\partial \lambda} [\lambda - \gamma_f p(q + \kappa)] > 0
\end{aligned}$$

Since $\lambda - \gamma_f p(q + \kappa) < 0$ by A2.1 and $\frac{\partial Pr(B)}{\partial \lambda} < 0$ by Proposition 1.

With respect to τ :

$$\begin{aligned}
\frac{\partial R}{\partial \tau} &= \frac{\partial Pr(B)}{\partial \tau} [(1 + \tau)b(\theta) + \lambda_f] + Pr(B) \left[\frac{\partial b}{\partial \tau} (1 + \tau) + b(\theta) \right] - \frac{\partial Pr(B)}{\partial \tau} \gamma_f p(q + \kappa) \\
&= \frac{\partial Pr(B)}{\partial \tau} (1 + \tau)b(\theta) + Pr(B) \left[\frac{\partial b}{\partial \tau} (1 + \tau) + b(\theta) \right] + \frac{\partial Pr(B)}{\partial \tau} (\lambda_f - \gamma_f p(q + \kappa)) \\
&= \frac{\partial Pr(B)}{\partial \tau} (1 + \tau)\frac{Pr(B)c}{2} + Pr(B) \left[\frac{\partial Pr(B)}{\partial \tau} \frac{c}{2} (1 + \tau) + \frac{Pr(B)c}{2} \right] + \frac{\partial Pr(B)}{\partial \tau} (\lambda_f - \gamma_f p(q + \kappa)) \\
&= \frac{\partial Pr(B)}{\partial \tau} (1 + \tau)Pr(B)c + Pr(B)^2 \frac{c}{2} + \frac{\partial Pr(B)}{\partial \tau} (\lambda_f - \gamma_f p(q + \kappa)) \\
&= \frac{\partial Pr(B)}{\partial \tau} [\gamma_f p(q + \kappa) - \lambda_f + (c - pq)(1 + \tau_f)] + Pr(B)^2 \frac{c}{2} + \frac{\partial Pr(B)}{\partial \tau} (\lambda_f - \gamma_f p(q + \kappa)) \\
&= \frac{\partial Pr(B)}{\partial \tau} (c - pq)(1 + \tau_f) + Pr(B)^2 \frac{c}{2} > 0
\end{aligned}$$

Which is positive by Proposition 1, A2, and A4.

With respect to γ :

$$\begin{aligned}
\frac{\partial R}{\partial \gamma} &= \frac{\partial Pr(B)}{\partial \gamma} [(1 + \tau_f)b(\theta) + \lambda_f] + Pr(B)(1 + \tau_f)\frac{\partial b}{\partial \gamma} - \left(\frac{\partial Pr(B)}{\partial \gamma} \gamma p(q + \kappa) + Pr(B)p(q + \kappa) \right) \\
&= \frac{\partial Pr(B)}{\partial \gamma} [(1 + \tau_f)b(\theta) + \lambda_f - \gamma p(q + \kappa)] + Pr(B) \left[(1 + \tau_f)\frac{\partial b}{\partial \gamma} - p(q + \kappa) \right] \\
&= \frac{\partial Pr(B)}{\partial \gamma} [\gamma_f p(q + \kappa) - \lambda_f + (c - pq)(1 + \tau_f) + \lambda_f - \gamma p(q + \kappa)] - Pr(B)\frac{p(q + \kappa)}{2} \\
&= \frac{\partial Pr(B)}{\partial \gamma} (c - pq)(1 + \tau_f) - Pr(B)\frac{p(q + \kappa)}{2} < 0
\end{aligned}$$

Which again is negative by Proposition 1, A2, and A4.

C Additional empirical results

C.1 TWFE robustness

Definition of treatment: Several additional tests lend credibility to a causal interpretation of the results. Until now, I have included all non-multinational firms in “local.” In Table A3 I disaggregate separate treatment indicators for fields operated the NPDC – the state oil company – and those operated by independent local firms. I find that the effect on shut-ins and output is primarily driven by private firms. In contrast, the efficiency costs of localness in terms of greater malfunctions essentially vanishes when we disaggregate the treatment with a negative and insignificant point estimate, while the effect size rises to 3.6 for state-run fields. At the same time, reductions in theft, militancy, and piracy are also large and significant for private firms but insignificant for the government. Private local firms appear to have no efficiency disadvantage, magnifying the output benefits of localness. In contrast, the efficiency costs of public production are quite large and the benefits smaller, resulting in a smaller output effect.

Oil prices: I also test robustness against oil prices in Table A4. I find no evidence that differential responses to oil price changes by localized fields are driving the results. Only for militant attacks do the main results of Table 2 lose their significance, although the point estimates remain negative.

Measurement error in output: I also consider robustness of output results to potentially non-random measurement error in output, including double-counting output for fields where multiple operators are observed in a given year. In Table A5, I restrict the sample to fields with one listed operator in columns (1)-(4) or to only producing fields in columns (5)-(6); I find the magnitudes of the main quantity and revenue effects unchanged. Lastly, to account for the fact that clustered standard errors that may be biased in cases where the number of treated clusters is small, I use randomization inference to calculate standard errors, the results of which are in Figure A5. The results clearly show that the estimated coefficient and t -statistic is in the far left tail of the distribution of estimates over 2000 random permutations of the treatment assignment, corresponding to a p -value of 0.045.

Difference-in-differences weights: Several related methodological papers show that the TWFE estimate can be decomposed into a weighted average of individual average treatment effects (ATEs) across units and time (de Chaisemartin and D’Haultfoeuille 2019, Goodman-Bacon 2019, among others). It can be shown that such weights may be negative because in staggered-event designs such as ours, already-treated units may later act as controls. The weighted TWFE estimate also tends to underweight units that are treated early or periods later in the panel. Under sufficient treatment effect heterogeneity, the TWFE estimate can differ markedly in size and sign from the individual

ATEs. de Chaisemartin and D’Haultfoeuille (2019) provide some guidance by deriving a general formula for the unit-time-specific weights of the treated observations, which allows diagnostic testing on the share of negative weights. I apply their analysis in A7, which displays histograms of estimated weights for each of the 6 outcomes in Table 2. In all cases, only a small share of the weights are negative, suggesting that it is unlikely that the TWFE estimate will be of a different sign than the individual ATEs. Furthermore, the authors suggest an alternative estimator that recovers the sample-weighted ATE at the period of switching and dynamically, under a refinement of the common trends assumption in staggered adoption designs. I estimate dynamic effects using their method for 10 post-treatment periods, bootstrapping standard errors, and display the results in Figure A8. In general, the results are similar to the standard TWFE event-study results, and the dynamic treatment effects are of the correct sign.

Goodman-Bacon (2019) decomposes the TWFE estimate into a weighted average of all two-by-two difference-in-difference comparisons. These weights depend on the size of the groups and the variance of the treatment in each 2×2 comparison. As such, the TWFE will tend to place lower weight on 2×2 estimates for units treated early or late in the panel, and will generally not correspond to the ATT, which is sample-share-weighted. The key insight is that these weights identify which comparisons are driving the TWFE results. Table A6 presents weights and average treatment effect estimates for each 2×2 DD comparison type. Because of the large sample of untreated clusters, the TWFE estimate heavily weights the “treated vs. never treated” 2×2 comparison, which accounts for 83% of the treatment effect. Still, every 2×2 group estimate is negative except for the “treated vs. already treated” comparison, which is near zero.

The second key insight from Goodman-Bacon (2019) and others is that early-treated groups act as controls in later periods when their treatment status does not change. If treatment effects vary over time, then these already-treated units may have differential post-treatment trends even as they are serving as controls for future switchers. This can introduce bias in the TWFE estimate by implicitly violating parallel trends for the 2×2 comparisons in which already-treated units act as controls.³⁰ One way to address this issue is to run event-study regressions as in Figures 3 and 4.³¹

An alternative estimation method is the stacked DD (see Gormley and Matsa 2011, Deshpande and Li 2019 for examples), as suggested by Goodman-Bacon (2019). In this method, treated units in each treatment-year cohort are paired with all not-yet-treated observations in the data as of year t . The cohorts are then “stacked” to obtain a dataset in which the control groups are always untreated,

³⁰ This is identical to the “negative weights” problem identified in de Chaisemartin and D’Haultfoeuille (2019).

³¹ Abraham and Sun (2018) show that event-studies are unbiased as long as there is no cohort-specific heterogeneity in the time-path of effects. Of course, I re-weight to correct for cohort heterogeneity in Figure A6

and the event-time takes the place of calendar year. This eliminates the negative weighting/ 2×2 bias problem by ensuring that already-treated observations are never used as controls. We then estimate the following equation, for unit i in cohort-stack c for event-time t

$$y_{ict} = \alpha + \beta local_{ict} + \delta_{ct} + \gamma_{ic} + \epsilon_{ict}$$

Standard errors are clustered at the stack-field level. The parameter β is a variance weighted average of cohort-specific causal effects, where each cohort-specific comparison is only between newly treated and not-yet-treated groups. An additional robustness test is to further restrict the sample only to ever-treated fields, eliminating any bias that may emerge from comparing ever-treated to never-treated fields. Then each c relies only on comparisons between an earlier-treated treatment group and later-treated controls. The results of this analysis are given in Table A7. I find that full-sample stacked-DD estimates (columns 1-3) are robustly negative and significant for theft, militancy, and shut-ins, and positive and significant for output. The magnitude of effects is in fact somewhat larger than the TWFE estimates in Table 2. The effect on malfunctions remains positive but not significant. The results indicate that using already-treated units as control is not a substantial source of bias in our main TWFE estimates, consistent with their low weights in Table A6.

In columns (4)-(6) of Table A7 I estimate the stacked DD on only the ever-treated sample. Results are of the correct sign, but now smaller and insignificant for militant attacks and malfunctions. In contrast, the results for theft, output, and shut-ins are all robustly significant. I also estimate event-studies in the stacked format, the results of which are displayed in Figure A10. The results look similar to the main and re-weighted event-study plots, although the estimates appear to be more precise. I also test robustness to estimating the stacked DD regression over all possible event-windows for output and theft, the two main outcomes. The resultant β coefficients and t -statistics are plotted in Figure A11. As desired, they are clustered around large negative and positive values, respectively.

Cohort-specific heterogeneity: Abraham and Sun (2018) show that the standard TWFE event-study specification produces estimates $\hat{\psi}_\tau$ that are a weighted average of cohort-specific estimates. These weights can be non-convex, which, in the presence of treatment effect heterogeneity, can render results difficult to interpret and undermine the validity of the test for pre-trends. They propose estimating cohort-specific event-study coefficients and then applying convex weights to these coefficients derived from the share of each cohort in the treated population for a given event-period τ . In other words, I estimate

$$y_{it} = \alpha + \sum_{\tau=-T}^T \sum_c \psi_\tau^c L_{it}^\tau 1(t_i = c) + \delta_t + \zeta_i + X'_{it} \beta + v_{it}$$

And then form the re-weighted event-study treatment effect $\tilde{\psi}_\tau = \sum_c \hat{\psi}_\tau^c \omega_\tau^c$, where $\omega_\tau^c = \frac{\sum_i L_{it}^c 1(t_i=c)}{\sum_i L_{it}^c}$. The results are in Figure A6, which re-weights the event-study for the two main outcomes – theft (Panel A) and oil production (Panel B). The parallel trends appear to hold.

C.2 Local employment spillovers

Part of the rationale behind indigenization is that local firms may increase the positive spillover effects of oil production to local communities. If this is the case, then it's possible that the effects we see are driven by higher opportunity costs for attracting labor into the criminal sector. In particular, if spillovers improve employment opportunities for young men, then the gangster's cost c may rise as labor costs rise. Theoretically, this could be responsible for reduced criminal activity and increased output, as $\frac{\partial Pr(B)}{\partial c} > 0$, since higher cost gangs are easier to buy off.

To test this hypothesis, I use data from three rounds of Nigeria's General Household Survey, a 3-wave panel survey covering 32,537 Nigerians in 500 villages from 2010-2016. I link each village to its nearest oilfield in order to identify villages treated by localization of nearby fields. I then drop all villages further than 50 km to their nearest oilfield. For individual-level regressions, the analysis sample is all individuals of working age, defined as 15-60. For individual (or household) i in village v near to field f at time t , I estimate the following

$$y_{ivft} = \alpha + \psi_{local} f_t + \delta_t + \zeta_f + X'_{ivft} \beta + \mu_{it}$$

For y_{ivft} I consider individual and household measures including employment, employment outside the home, self-employment, and employment in household agriculture, as well as the log of overall per capital household consumption. Household-level controls included in X are household distances to roads, population centers, markets, borders, and state capitals; village-level controls are slope, altitude, mean annual temperature, and annual rainfall. Each of these time-invariant conditions is interacted with year dummies. Standard errors are clustered at the field level

Results of this estimation are given in Table A10. Each Panel considers a different individual-level employment outcome. Columns (1)-(4) estimate using the entire sample of fields with various combinations of year, month, field, and state-by-year fixed effects, as well as the interacted controls. Columns (5) and (6) exclude all individuals residing in a village whose nearest oilfield was offshore, where spillovers are less likely to manifest.

TABLE A10 HERE

The results show no effect on the level of employment (Panel A). Across all specifications, the results are robustly zero. For the composition of employment, I do not find any statistically significant changes in employment outside the home (Panel B) or employment in household agriculture (Panel D), although the point estimate for both of these outcomes are consistently negative. However, there does appear to be an increase in self-employment (Panel C) by roughly 7-9 percentage points, significant at 1%. Since overall employment does not change, this effect seems to be offsetting small and statistically insignificant reductions in other categories. Lastly, I test the impact of localization on log household per capital consumption in Table A11. Again, there are no statistically significant effects, though the point-estimates are generally positive. Overall, there is no evidence that localization creates positive economic spillovers for nearby oil-producing villages.

I test for parallel pre-trends in Figure A14. All results suggest that pre-trends are essentially flat and insignificant for each outcome considered in Table A10. The pattern of dynamic effects does suggest some increase in self-employment, as well as decreases in employment outside the home and in household agriculture. Lastly, the aggregate employment effect does appear to have a small positive trend for years $\tau > 5$. However, as Goodman-Bacon (2018) and de Chaisemartin and D’Haultfoeuille (2019) show, late-adopters and later periods in the panel are down-weighted in the TWFE estimate, perhaps accounting for the zero aggregate effect in Table A10 Panel A, despite a small positive effect in some of the event-study coefficients.

FIGURE A14 HERE

Opportunity costs for young men – and not other demographic groups – are likely to determine wages offered by organized crime. If employment effects are heterogeneous across demographics, then it may be that the aggregate zero effects are masking effects on the demographic groups relevant for the gangsters’ cost structure. To test this hypothesis, I re-estimate the employment equation of each outcome by ten-year age bins and gender. The results are displayed in Figure A15, which plots coefficients by age bin and gender for each outcome. For men (top panel), the results indicate robust zeroes along each outcome and for each age group, with the exception of some noisy estimates for older age groups with small sample sizes. In contrast, the plot reveals that middle-aged women are driving the aggregate positive effect on self employment, which is offset by a reduction in agricultural employment for the same demographic group. For both men and women, the aggregate employment effects are zero at all ages. Therefore, while women observe some reallocations of labor across sectors as a result of localization, young men –our demographic of interest – do not experience

any changes. It is therefore unlikely that the effect of localization on theft and output is operating through opportunity cost mechanisms.