The local advantage: Corruption, organized crime, and indigenization in the Nigerian oil sector [*]

Jonah M. Rexer[†]

University of Pennsylvania

April 30, 2021

Abstract

Despite productivity advantages, multinational firms may operate less effectively than their local competitors in markets plagued by corruption and conflict. In natural resource sectors where firms face predation by armed groups, local firms may more easily engage in efficient corruption to buy law enforcement protection for their assets. I study a two-decade indigenization drive in Nigeria's turbulent oil sector, during which the share of local ownership grew substantially. Local takeover considerably increases oilfield output and reduces the share of nonproducing assets, despite evidence that local firms are of lower quality. Local firms increase output by reducing black-market activity: oil theft, maritime piracy, and militant violence all fall following local takeover. A simple bargaining model illustrates that political connections enable local firms to align law enforcement incentives, explaining their superior output performance. Data on anti-oil theft raids by government forces show that local firms receive preferential law enforcement protection. I find evidence that connections to high-level politicians and the security forces drive local firms' advantage in obtaining state protection and reducing criminal activity.

JEL Classifications: F2, L24, Q34, Q35
Keywords: foreign investment, hydrocarbons, political risk, organized crime, black markets, conflict, law enforcement, corruption.

[*]For their outstanding feedback and assistance, I thank Ann Harrison, Camilo Garcia-Jimeno, Guy Grossman, Santosh Anagol, Arthur van Benthem, Nina Harari, Michael Watts, Jacob Shapiro, Nicholas Ryan, Alexey Makarin, and Bruno Baranek. I also thank participants at the Pacific Development Conference at UC Berkeley, the Oxford University Center for the Study of African Economies Conference, the Northeast Workshop on Empirical Political Science, the Young Economists Symposium at UPenn, the African Development Bank, the NBER Development Economics meeting, and the Working Group on African Political Economy. I also thank representatives at the Nigerian Bureau of Statistics, the Department for Petroleum Resources, and the Federal Ministry of the Environment for facilitating access to data. Moses Ogunsanya provided outstanding research assistance. This work has been supported by the Kleinman Center for Energy Policy at the University of Pennsylvania and the Centre for Economic Policy Research. All errors are my own.

[†]Wharton School of Business. 3620 Locust Walk, Philadelphia, PA 19104. Email: jorexer@wharton.upenn.edu. https://www.jonahrexer.com/

1 Introduction

Global experience and a large body of evidence demonstrate that multinational companies (MNCs) are more productive on average than locally-owned firms in developing economies. They have better human capital, better technology, and better management practices (Aitken and Harrison 1999, Arnold and Javorcik 2009, Guadalupe, Kuzmina, and Thomas 2012, Criscuolo and Martin 2009).[1] Despite these advantages, foreign multinationals may be ill-equipped to deal with the corruption, conflict, and expropriation that often accompanies working in difficult markets (Burger, Ianchovichina, and Rijkers 2015 Blair, Christensen, and Wirtschafter 2019).[2] In these contexts, local companies may possess the political connections that allow them to protect against expropriation and violence, and navigate the political environment more broadly.[3]

The extractive industries in weakly-institutionalized, resource-rich states present local firms with a consequential but underappreciated advantage. In many countries, natural resource extraction is plagued by active black markets in oil and minerals run by organized crime and violent armed groups.[4] These rackets of theft and extortion impose substantial costs on firms operating in these sectors. Throughout history, such lucrative black markets have often been accompanied by law enforcement corruption as state security agents use their (partial) monopoly of violence to appropriate illicit rents.[5] Selective law enforcement creates the possibility that local firms in volatile natural resource sectors may leverage superior political connections to more easily obtain state security protection. Locals' propensity for "efficient corruption" – which substitutes away from inefficient black markets – suggests that multinational divestment to local firms may be an effective policy to boost resource output, reduce criminal activity, and alleviate the resource curse.

In this paper, I study the benefits of localness in the context of the Nigerian petroleum sector, an industry fraught with political violence and corruption. From 2000-2009, the oil-producing Niger Delta region witnessed an armed uprising in which militant groups attacked multinational oil infrastructure in order to wrest greater control over oil revenues. In the aftermath, a multi-billion dollar-a-year black market for crude oil stolen from onshore oil

[1]They also raise aggregate productivity by transferring technology (Teece 1977, Guadalupe, Kuzmina, and Thomas 2012) and skills (Bloom and Reenen 2010, Bloom, Sadun, and Reenen 2012), forcing inefficient firms to exit via competitive pressures (Alfaro and Chen 2018), and transmitting human capital to local firms through labor markets (Balsvik 2011, Poole 2013). See Alfaro and Chauvin (2020) for a review of the multinational spillovers literature.

[2]Although Guidolin and Ferrara (2007) shows that conflict was beneficial to multinational diamond firms during the Angolan civil war.

[3]Evidence suggests that corruption does encourage joint ventures as multinationals seek partners to navigate local politics (Javorcik and Wei 2009).

[4]In Colombia, the gold trade is marred by paramilitary taxation and theft, while in the DRC, armed militias control substantial mineral extraction (Sierra 2020). The black market for stolen crude oil is worth roughly 100 billion USD annually (Ralby 2018).

[5]For example, prohibition in the U.S. (García-Jimeno 2016) or the current drug trade in Mexico.

pipelines has emerged (Rexer and Hvinden 2020). In recent years, the sector has undergone an ownership transformation; the share of onshore output produced by local firms grew from 6.4% in 2008 to 35.8% in 2016.

Using detailed annual panel data on Nigeria's active oilfields, I leverage this wave of indigenization to study the effect of local ownership on oil output, crime, political violence, and law enforcement activity. Since locally-owned assets may differ from multinational ones across many underlying characteristics, I employ a difference-in-differences approach that exploits changes in field ownership driven by multinational divestment to local firms. I find that local takeover increases output by nearly 60%, while reducing the share of non-producing fields by 17 percentage points.

This "local advantage" occurs despite local firms exhibiting lower quality. Consistent with lower operating standards, local firms have outsize negative environmental impacts. Divested fields experience 22% more oil spills due to mechanical failure, reflecting lower quality safety standards, management, and maintenance of physical infrastructure. Localization also increases gas flaring – the environmentally damaging practice of burning natural gas byproduct on site – by more than 60%, equivalent to an additional 36,000 tonnes of CO_2 emissions per field annually. These effects cannot explained by increased production under local ownership.

Local firms increase output despite lower quality by mitigating criminal and violent activity: local takeover leads to a prolonged reduction in incidents of oil theft, oil-related militant violence, and maritime piracy. On average, a locally operated field experiences 3.4 fewer theft incidents per year, a 33% decline, and 0.7 fewer oil-related militant fatalities. As a falsification test of the causal pathway, I examine heterogeneous effects by asset type. I find that the local advantage in output, theft, and violence is concentrated onshore, consistent with greater political risk exposure. Instead, local companies' quality disadvantage is more pronounced offshore, where technology and capital requirements are higher.

Of course, divestments to local firms may be correlated with unobserved trends in oilfield quality. I bolster the claim of causal identification using detailed data on the universe of corporate transactions in Nigeria's oil and gas sector. I exploit the fact that in Nigeria, a weak legal framework creates substantial regulatory discretion over oil and gas transactions, leading many planned divestments to be stalled or terminated. I show that these planned divestments do not produce effects until a divestment is ultimately consummated. Since these fields are ostensibly subject to similar unobserved trends as those divestments unencumbered by capricious bureaucratic interference, this placebo test allays concerns that unobserved trends are driving the results.

The sizable rents from illegal oil theft in the Niger Delta create strong incentives for law enforcement corruption. Indeed, criminal armed groups in the sector operate primarily under the protection of the state security forces. As such, a key source of local advantage may be the ability of local firms to leverage political connections to increase the provision law enforcement on their assets. To illustrate this dynamic, I write down a simple model in which a firm bargains with the security forces to obtain law enforcement, where the government maintains the outside option of collecting criminal bribes and allowing theft. Corruption is "efficient" only when it incentivizes law enforcement provision, maximizing social surplus.

Despite the scope for efficient corruption, several frictions may prevent a deal: *i*) law enforcement is costly, *ii*) bribery laws impose costs on firms undertaking illicit payments, and *iii*) joint ownership implies that the losses from theft are only partially internalized by the operating firm, while the gains are fully internalized by organized crime. Nevertheless, I show that political connections are able to overcome these frictions by providing law enforcement with a claim to residual oil income, therefore aligning incentives between the firm and the state. A such, connected firms experience more enforcement and less crime. In addition, I derive numerous other comparative statics that demonstrate how equilibrium outcomes vary with features of the strategic environment. Importantly, I also show that when there are severe commitment problems, lump-sum bribes are rendered ineffective; only transferring claims to residual surplus can align law enforcement incentives to reduce theft. Political connections therefore become a binding constraint to law enforcement provision, without which firms are subject to state-sanctioned predation by organized crime. An obvious explanation for local firms' observed advantages, then, is their ability to leverage superior political connections to expand the space for efficient corruption.

The model predicts that when state security forces internalize firms' losses from theft, they are more willing to enforce the law, reducing theft and attendant output losses. The inclusion of powerful politicians and security agents with influence over law enforcement activity as shareholders and board members of oil companies may plausibly align incentives toward greater provision of security. I argue that local firms utilize superior political connections to increase enforcement, and provide several pieces of evidence that this mechanism is likely to explain local firms' ability to reduce crime and raise output.

First, using data on law enforcement activity collected from Nigerian news media, I show that local firms indeed receive greater law enforcement protection. In the TWFE model, local takeover increases state law enforcement actions against oil theft by 83%. This increase is robust to numerous specifications and falls on numerous activities across the black-market value chain, from illegal refining for domestic consumption to illegal export activity. Notably, however, placebo tests show that these enforcement effects are null for non-oil crime, ruling out unobserved differential trends in policing in divested areas. Event study regressions strongly suggest that preferential enforcement for local firms causes their lower levels of crime. Local takeover is immediately followed by a large and significant increase in enforcement, while crime reduction is initially small but grows over time as crime responds negatively to enforcement spike. After an initial period of elevation, enforcement effects ultimately taper toward zero as crime falls and black markets wane.

Next, I identify politically-connected firms using hand-collected data on the biographies of board members, shareholders, and managers. Using TWFE, I show that fields operated by politically connected firms experience lower oil theft. These associations are greatest for connections to the security forces, the group most intimately involved in black-market corruption and with the most direct control over enforcement activity; connections to technocrats in regulatory agencies and the state oil company have smaller effects. Local firms are also much more likely to have political connections in general and strategically important security connections in particular. Taken together, the evidence suggests that local firms leverage superior political connections security to align incentives between firms and the security forces and obtain law enforcement protection for their assets.

I also find suggestive evidence in favor of additional mechanisms contained in the model. Using variation across multinational companies, I find that exposure to a foreign corruption law *increases* field-level theft by 6.1 incidents and violence by 0.15-0.45 deaths annually, suggesting that part of the local advantage stems from the lower costs of corruption imposed by weak Nigerian anticorruption law. In addition, I use data on equity stakes in oil licenses to show that indigenization increases the equity share of the operating firm by 20%, driven both by lower state ownership requirements for local firms and consolidation of multinational stakes during joint venture divestments. With larger ownership stakes, local operators internalize a greater share of theft losses, increasing incentives to deter criminality. These additional results provide important insights into why multinational firms fail to engage in efficient corruption: first, they are subject to more onerous legal restrictions (in addition to reputational risks) governing the composition of their boards and shareholders, so they are unable to directly transfer rents to political and security figures. However, they are also unwilling to use joint ventures to subcontract this activity to local firms, because in order to incentivize a joint-venture partner, a mutinational would need to provide sufficiently large ownership stakes to the local operator. This may not be optimal, particularly when firms retain the outside option of going offshore.

The results are robust to controlling for interactions between year dummies and time-invariant field-level covariates. Event-study models indicate that differential pre-trends in outcomes of interest are not driving the takeover effects.[6] In addition, I test the robustness of the results to differential effects of oil price trends, measurement error in output, correlation with region-specific policies, and locality-by-year fixed effects. Lastly, I run diagnostic tests from Goodman-Bacon (2019), Chaisemartin and D'Haultfoeuille (2019), and Callaway and Sant'Anna (2019) on the two-way fixed effects specifications, which indicate that the estimates appear reliable even in the presence of treatment effect heterogeneity. Lastly, I consider

[6] I also estimate a re-weighted event-study using cohort composition to adjust for heterogeneous effects (Abraham and Sun 2018), finding similar dynamic treatment effects.

several alternative explanations for the results – local employment spillovers, differences in discount rates, targeted host-community investments, and local grievance toward multinationals. I find no support for any of these alternative mechanisms. Finally, I find substantial heterogeneity in the response of oil theft to localization with respect to oil prices and local military strength, consistent with additional predictions of the model.

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 add nuance 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 natural resource sectors in conflict-affected countries have very different dynamics; in these cases local advantage can massively outweigh the productivity gains from foreign investment. The policy implication is that indigenization in troubled extractive sectors may be justified on productivity grounds. However, these benefits must be balanced against the substantial welfare costs of increased environmental pollution. This paper illuminates new and important tradeoffs between local and multinational ownership.

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, H. Li et al. 2008, Khwaja and Mian 2005, Akcigit, Baslandze, and Lotti 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 law enforcement is closely linked to black-market activity, political connections – and the corruption they engender – substitute inefficient black markets for legitimate production and incentivize greater output on the part of firms. The results demonstrate a novel mechanism by which political connections matter for firm outcomes: they protect against criminal activity by incentivizing the provision of protection by corrupt local law enforcement agencies.

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, Sierra 2020, 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).[7] It has also studied the economic effects of local resource booms on income, employment, and prices (Aragon and Rud 2013, Lippert 2014, Loayza and Rigolini 2016). This work is one of the few to demonstrate the centrality of black markets and organized crime to resource curse dynamics (Couttenier, Grosjean, and Sangnier 2017, Buonanno et al. 2015). Unlike previous work, I also analyze firms as strategic participants in the resource

[7] For reviews of this literature, see Cust and Poelhekke (2015) and Aragon, Chuhan-Pole, and Land (2015)

curse nexus. I am the first to show that local ownership can mitigate some of the most violent pathologies 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 waters of the Niger Delta, a region in the 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 states,[8] home to 22% of Nigeria's population (NBS 2017), and populated by numerous ethnic minority groups. Since oil discovery in 1956, the sector has historically been dominated by oil supermajors Shell, Chevron, ExonnMobil, 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, or 98% of Nigeria's export earnings.

All multinationals operate profit-sharing agreements with the state-run oil company, the Nigerian National Petroleum Company (NNPC), structured as joint ventures often involving several multinationals, production sharing contracts, or 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 an average 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.[9] Multinationals in Nigeria must contend with the added costs of corruption, which expose them to legal liabilities in their home countries.

Oil companies' relationship with the Niger Delta communities in which they operate is fraught. The region is the prototypical example of the local resource curse – a constellation of armed groups interact with oil companies, local and federal government, and each other in a low-grade conflict that blurs the line between civil war and organized crime (Obi and Rustad 2011, Watts 2007). Local politicians 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

[8] Abia, Bayelsa, Delta, Rivers, Akwa Ibom, Imo, Ondo, Edo, and Cross River states.

[9] See https://carnegieendowment.org/2018/07/17/new-taxonomy-for-corruption-in-nigeria-pub-76811 for a survey of corruption issues in Nigeria. 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, https://www.bbc.com/news/business-46336733.

range from 39-64%.[10]

Between 2000-2009, The Niger Delta Crisis saw the emergence of well-armed militants from longstanding criminal gangs and ethnic political militias (Watts 2007, Asuni 2009). Militants declared war on the federal government and oil companies, destroying oil facilities and kidnapping staff in an attempt to obtain concessions for themselves and the region. In 2009, the Federal Government announced amnesty to nearly 25,000 combatants, as well as lucrative "pipeline surveillance contracts" awarded to several commanders – payments which amounted to private transfers to top militants (SDN 2019c). Rexer and Hvinden (2020) show that the amnesty process led to the emergence of a thriving black market in stolen oil, comprised both of ex-militants and more recent entrants. Figure A2[11] charts the evolution of the black market by plotting the monthly incidents of pipeline sabotage.[12] 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, decline over the whole period. In 2016, the black market totaled 4.2 billion dollars, or 15% of Nigeria's total production (NEITI 2016).

In this two-tiered market, smaller downstream entrepreneurs refine about 75% of the stolen crude locally for sale to the domestic market, while larger criminal syndicates typically export the remainder (SDN 2019a, SDN 2019b). The region's pipeline network, traversing thousands of kilometers of militant-controlled swampland, is extremely vulnerable to theft. Protection rackets naturally arise: oil companies must negotiate with gangsters and local communities in order to safeguard output. Payments to local communities – which range from direct transfers and contracts for local chiefs and militant groups to community-wide development projects – are a cost of doing business. Politicians and security forces play an important role in the black market. 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 lucrative illegal tapping operations (SDN 2019a).

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

FIGURE 1 HERE

[10] Still, the Delta compares favorably most of Nigeria in this respect, with only one state ranking in the top-20 in Nigeria.

[11] Reprinted from Rexer and Hvinden (2020).

[12] I argue in Section 3 that this is a good proxy for the number of theft incidents and the size of the black market.

At the same time, the sector is becoming increasingly Nigerian, in part because of multinational divestment from onshore. 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 9 to 31, while the number of fields operated by local companies rose from 9 to 70. 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, where the local share has grown from 2.4% to 34% over the same period. 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. This post-law growth in local participation is driven both by an increasing rate 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

Below I briefly describe the key sources of data I use to test the local advantage hypothesis. For greater detail on the sources of data, the cleaning process, and the construction of key variables, please refer to Appendix B.

3.1 Data description

3.1.1 Oil data

Information on 314 active Nigerian oilfields forms the core of the data. Field-level data on oil production come from the Annual Statistical Bulletin of the NNPC, augmented with data from the Department of Petroleum Resources (DPR), and covers the years 1998-2016. In each year I record total output, in millions of barrels, for the field, as well as the identity of the operating company. A field enters the dataset in the year it first appears in these administrative records, and remains thereafter. A "shut-in" field is defined as a field that is nonproducing in a given year. Time-invariant field-level covariates are the number of wells, date of completion of the first well, and the depth of the deepest well. I link fields to information on oil theft, violence, piracy, and geospatial control variables using centroid coordinates. The fields are mapped in Figure 2, with the color indicating the year in which the observation was treated. Over the sample period, there are 70 ever-treated fields and 244 never-treated.

FIGURE 2 HERE

With some exceptions, ownership of Nigerian oilfields is determined at the concession-level. Detailed data on concessions for the years 2013-2018 comes from the DPR and the Nigerian Extractive Industries Transparency Initiative (NEITI). These sources contain the concession size, location, operator, license type, and shareholder breakdown. Since this data is only available for a limited time period, I exclude it from the main analysis and use it only to test auxiliary model predictions.

Data on oil theft comes from the Nigerian Oil Spill Detection and Response Agency (NOSDRA), a division of the Federal Ministry of the Environment. I obtain 11,587 reported oil spills from 2006-2017. For each oil spill, I observe the location and cause of the spill, as well as a text description. 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] To measure the technical efficiency of oil production, I use all field-level spills that are not due to sabotage, but rather equipment malfunction or unknown causes.

Data on monthly gas flaring from 2012-20120 comes from the NOSDRA Gas Flare Tracker. The Tracker uses VIIRS Nightfire satellite data to identify flaring sites remotely and converts luminosity to measures of gas output using an algorithm from Hodgson (2018).[14] Gas flaring is measured at the field-month or field-year in thousand cubic feet (mscf). These location-specific volume estimates can then be converted to CO_2 emissions, since according to U.S. Energy Information administration, flared natural gas emits 54.75 kg of CO_2 per mscf.[15] In total, I obtain data on 180 flare sites corresponding to 136 fields that appear in the NNPC-DPR production data.

In general, I measure treated fields from administrative data by the identity of the operating company. However, I also use data from DrillingInfo (DI), a corporate database on the oil and gas sector, as an independent source of data on corporate transactions. From DI, I identify 171 Nigerian oil sector transactions and use these to generate a dataset of asset sales containing the asset (field or oil block) sold and the nationalities of buyers and sellers. DI data also allows me to observe local divestments that were initiated but either stalled or terminated for bureaucratic reasons. I match these transactions to the administrative production data at the field level.

[13] See Rexer and Hvinden (2020) for a discussion about measuring oil theft.

[14] For greater detail on the remote sensing methodology, consult the Gas Flare Tracker website https://nosdra.gasflaretracker.ng/data.html

[15] https://www.eia.gov/environment/emissions/co2_vol_mass.php

3.2 Political connections data

For each of the 40 firms – foreign and domestic – that ever appear as operators in the NNPC-DPR data, I obtain data on the identities of boardmembers, managers, and shareholders from several sources: Firstly the Bureau van Dijk Orbis global company database contains information on name, position, and demographics of boardmembers, managers, and shareholders for reporting companies. I augment this with information scraped from company websites on boardmembers and senior management. Lastly, I use the Oil and Gas Map of Nigeria, an "independent initiative to monitor the Oil and Gas industry of Nigeria," to identify additional shareholders. 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 400 individuals over 37 companies.[16] I use this biographical information to 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), cabinet-level politicians, and members of the army and police.

3.3 Conflict and militancy data

Data on conflict and violence comes from from the Armed Conflict Location and Event Dataset (ACLED) and covers 1998-2016. To measure violent oil-related activity, I use all conflict events in which the description mentions a set of key words about the oil industry. This captures attacks on the oil sector perpetrated by various armed groups. I then further distinguish between conflict events perpetrated by organized rebel or political militia groups, which I call "militant" attacks, and those perpetrated by unknown or unorganized groups, which I call "non-militant" attacks. I measure conflict intensity using total annual fatalities.

Finally, I use data on militant camps collected by the author from local NGOs and augmented by data from Blair and Imai (2013), which catalogue the location, commander, militant group affiliation, and amnesty status of 69 militant camps, as of 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 (Rexer and Hvinden 2020), 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). Ex-militants typically operate in their geographical spheres of influence by directly participating in the bunkering economy or providing protection for those who do. Using the data on group affiliation of each camp, I also code the number of groups surrounding each oilfield within a certain radius – a measure of the num-

[16] The three missing companies account for only 2.6% of total oil production.

ber of actors in the illegal market. Lastly, I take a measure of group military strength derived and validated in Rexer and Hvinden (2020) which identifies the strongest camps based on the number of local allies along the pipeline network.

The various data sources have different time series and degrees of completeness. To harmonize the results, I take as the main estimation sample an unbalanced field-level panel from 2006-2017.[17] Within this period, oil production data is missing for some fields in each year because of incomplete coverage in the DPR-NNPC reports. I do not observe the cause of missingness, and assume this data is missing at random. Table B1 shows that outcomes and covariates are very similar across these samples, supporting this assumption. Therefore, while the estimation sample for all non-production outcomes is 3,497 field-years, the sample for regressions in which production is the outcome falls to only 2,476 field-years. The estimation sample for the gas flares data is an annual panel of the 314 fields from 2012-2019.

3.4 Law enforcement data

Data on law enforcement comes from the text of news media reports on raids, seizures, arrests, and other oil theft-related activity from Nigerian and international publications. We assemble a comprehensive sample of possibly crime-related news articles, and then employ local research assistants to identify relevant articles and extract all law enforcement events, defined as a unique interaction between criminals and state security agents in a particular location. From this procedure, we obtain 5682 unique geocoded law enforcement events from 2000-2020, of which 3261 are related to the illegal oil sector. For each event, we observe all involved law enforcement agencies, the criminal activity[18], the items destroyed or seized[19] the number of arrests, and the number of fatalities. Importantly, our data on illegal activities allow us to disaggregate criminal behavior and enforcement actions along the illegal sector value chain. We collapse these to the field-year level by taking all enforcement actions within twenty kilometers of a given oilfield.

3.5 Summary statistics

Figure 2 maps the oil infrastructure 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 244 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 70 locally-owned fields and their takeover dates. Indigenized fields are clustered primarily in the inland Niger Delta, with a

[17] The panel is unbalanced because fields may first appear the DPR-NNPC data in different years. However, once they enter, they remain in the dataset.

[18] Some prominent examples: transportation of stolen oil, piracy, pipeline vandalism.

[19] Some examples are guns, illegal refineries, stolen oil, and boats.

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 A1. 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. 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,[20] on average initiated in 1974-75. 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, since new blocks and marginal fields are more likely to be in the inland Niger Delta and offshoring by multinationals implies that divested fields are likely to be onshore. Indeed, 82% of ever-treated fields are onshore, while only 69% of multinational ones are, a difference that is significant at 5%. Treated fields are also slightly smaller, with 5.65 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.

I also compare differences in outcomes and other variables of interest for the analysis in the bottom panel of Table A1. 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 is on average 780,000 barrels (23%) lower among the treated, likely driven by smaller field sizes. In order to determine which of these relationships are causal, and control for the differences across time-invariant covariates, I 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. Rates of political connection are 10 percentage points greater for locally-operated fields, at 43%. 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 local politicians.

Figure A4 plots mean annual field-level sabotage incidents over time separately for ever-treated and un-treated fields, revealing growth in the black market to be heterogeneous. The two series start at similar levels in 2006 but diverge quickly. The plot suggests that the bulk

[20] Defined as the date of completion of the first oil well.

of the aggregate post-amnesty 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 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:

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

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 ξ_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, but also test robustness to the inclusion of controls for field size (number of wells), age, onshore, and maximum well depth. Standard errors are clustered at the field level. The key outcomes of interest are output, shut-ins, and non-theft oil spills, as well as measures of criminality – oil theft, violence, and piracy.

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

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

Where $L^\tau_{it} = 1(t - t_i = \tau) * local_i$, 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 has the benefit of dealing with the down-weighting of early-treated cohorts and bias introduced by time-varying treatment effects (Goodman-Bacon 2019). In addition, I test robustness to controlling for amnesty and other policy changes, differential responses to prices by treated status, measurement error in output, fixed field-level covariates interacted with time dummies, and calculating standard errors with randomization inference. I investigate the role of heterogeneity and the implicit weighting of the TWFE specification using results from Chaisemartin and D'Haultfoeuille (2019) and Goodman-Bacon (2019), address bias in the TWFE induced by using early-treated units as controls by used a stacked-DD design,[21] and estimate event-study regressions that are robust to cohort-specific heterogeneous effects (Abraham and Sun 2018).

Empirical variation in local ownership comes from three sources. First, fields are divested from multinational to local operators in asset sales. Second, new blocks offered for bidding are awarded to local firms, a practice increasingly common after the 2010 local content law created a preference for indigenous bidders. Third, marginal fields – smaller discoveries within concessions owned by multinationals that have been undeveloped for more than 10 years – are eligible by law to be farmed out to local operators. Since 2002, 30 marginal fields have been awarded to indigenous firms, of which 11 appear as producing fields in the data. Under TWFE, identification comes exclusively from the ownership transitions in 1). In fact, 2) and 3) functionally serve as control groups in the TWFE model, since they do not change nationality.[22]

The critical question for identification is whether multinational-to-local divestments can be considered exogenous, conditional on TWFE. Of course, the choice of divested assets is not random, as Table A1 shows. Differential trends across observable field-level characteristics can be flexibly controlled for by interacting fixed field covariates with time dummies. But the main concern is that differential trends on unobservables will bias the results. Intuitively, such trends should tend to bias the results against local advantage, since incumbent multinationals possess inside information on trends in field quality and would likely divest fields that are trending poorly on unobservables.

Flat pre-trends merely provide suggestive evidence, and are neither necessary nor sufficient for unbiased treatment effects (Roth 2019). To bolster identification, I rely on the fact that the precise timing of divestment is highly idiosyncratic. In particular, Nigeria's oil sector has for many years operated without a unifying regulatory framework due to failure to pass the long-delayed Petroleum Industry Bill.[23] Asset sales lack clear rules and are frequently subject to a host of discretionary regulatory actions; in the most recent wave, many transactions were stalled or terminated by uncertainty, litigation, and the capriciousness of the Ministry of Petroleum Resources.[24] As such, the precise timing of local takeover is unlikely to be system-

[21]See Goodman-Bacon (2019), Gormley and Matsa (2011), and Deshpande and Y. Li (2019)

[22]Note, however, that stacked-DD design excludes these as controls, yielding similar results.

[23]For a recent timeline of the PIB, see this report by the think-tank Good Governance Africa https://gga.org/wp-content/uploads/2020/01/23.01-PIB-Final-Draft_OB-final-reviewd-7April.pdf

[24]See the article here for numerous examples https://www.energymixreport.com/controversies-in-oil-assets-divestments-hinder-local-participation/

atically correlated with unobserved field trends, since it is not directly manipulable by market participants. In Section 5.4.3, I provide evidence in support of this assumption in a placebo test that uses data on delays and terminations to show that fields targeted for divestment, but not ultimately divested, do not exhibit any "localization" effects.

5 Main results

The main results of the TWFE models are in Table 1. 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). All results indicate the control group mean for reference.

5.1 Technical performance

Table 1 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 16.3-1.67 percentage points. Local firms therefore revive moribund fields when they assume operatorship. Output also rises by roughly 1.7-1.9 million barrels on average per field annually. This is a very large effect size, at roughly 60% of the control group mean. The increase in output translates to substantially higher average revenue upon local takeover, as seen in Table A2. The increase in output corresponds to 143 million dollars in revenue per year in column (4), or 64% of the control group mean. These output gains are not driven exclusively by decreasing shut-ins. Table B6, column (5) shows that the output effects hold conditional on field production. Therefore, the main effect on output operates not only on the extensive margin, but also requires increases on the intensive margin.

TABLE 1 HERE

In Table 1 columns (5)-(6), I estimate the effect of local ownership on equipment malfunctions that result in oil spillage. Local fields experience 0.9-1.5 more spills annually, or 14-22% of the sample mean. Across all specifications in Panel A, the inclusion of controls only slightly weakens the effect. I interpret the effects in (5)-(6) as evidence that local firms are less due to lower operating standards, though the precise source of this efficiency gap is unclear.[25] I also find that the effect on malfunctions is not mechanically driven by the effect of greater oil production.[26]

[25]These operational spills could be driven by lower-quality physical capital, human capital, or management practices and standards; I subsume all of these under efficiency differences.

[26]In Figure A5 I adjust the estimates to account for the fact that increasing production naturally leads to more malfunctions by subtracting the output-malfunctions elasticity times the effect of ownership on output from the estimate of $\hat{\psi}$. The results remain positive and significant.

5.2 Criminality and conflict

What drives this gap in performance? In Table 1 Panel B, I find that local takeover reduces crime and violence. Localized fields experience 3.4-3.5 fewer theft incidents annually, or 33-34% of the control group mean, significant at 1%. Locally-operated fields experience lower levels of violence: oil-related fatalities fall by 0.7 deaths per year, a result insensitive to controls and significant at 1% (columns 3-4). Lastly, local fields reduce piracy on their assets by 0.11-0.14 annual attacks, although this effect is not significant (columns 5-6). The local advantage seems to be comprised primarily of the ability to mitigate losses due to crime and violence.

In Table A3, I re-estimate the main TWFE equation using different types of violence, showing that the reduction in violence is not uniform across categories. In particular, disaggregating oil-related violent fatalities into those attributable to militant and rebel groups (columns 5-6) vs. those by less organized armed actors (columns 7-8) reveals that the main effect on oil-related violence is driven by the latter. Essentially, localization reduces gang violence rather than organized militancy, consistent with the different bargaining dynamics that affect these two types of violence in the Niger Delta.[27] The disaggregated violence effect sizes are visualized in Figure A6.

5.3 Environmental outcomes

Divestment to local firms clearly increases output and reduces criminality. Still, divestment may not be unambiguously welfare-improving. Localization may entail social costs in the form of increased environmental damage if local firms have less stringent operating standards. The environmental effects of divestment implied by Table 1 are ambiguous. On one hand, local firms are less efficient and therefore spill more oil during the normal course of operations. At the same time, the reduction in oil theft corresponds directly to fewer oil spills, since pipeline sabotage invariably results in oil spillage. Since oil spills from operations and theft may differ in magnitude, the net environmental effect on oil pollution is difficult to assess without more detailed data on spillage quantities.

In addition to oil spills, gas flaring represents a major source of environmental pollution from oil production in the Niger Delta. Flaring occurs when natural gas created as a byproduct from oil production is not economically viable to capture and transport to market, and is therefore burned on site. Gas flaring pollutes air quality, vegetation, and waterways, worsens health outcomes,[28] and contributes to climate change with CO_2 emissions. The practice has been subject to regulation since 1969, but meagre fines of 10 Naira per mscf flared (roughly

[27] Rexer and Hvinden (2020) show that organized militant violence is primarily driven by bargaining interactions with the Federal Government rather than firms.

[28] See Ologunorisa (2009) for a review of studies on the negative impacts of Niger Delta flaring.

0.03 USD in 2016) failed to deter flaring. In 2018, the flaring penalty was increased to 2 USD per mscf for concessions producing more than 10,000 bpd of oil.[29] Still, enforcement of this penalty is uneven, not least because companies under-report flaring volumes.[30] Figure A7 presents flaring on all active Nigerian oilfields over time. Average annual flaring between 2012-2020 was 380 million mscf, equivalent to roughly 20.8 million tons of CO_2 emissions, or 17.2% of Nigeria's 2016 total annual emissions according to World Bank data. The trend displays a slight dip followed by a recovery in recent years.

In Table 2, I estimate the effect of localization on flaring in the panel of oilfields from 2012-2019, the years for which gas flaring data is available. Note that gas flaring data is sometimes recorded in years before a field first enters the oil production data. For robustness, I consider estimation both in a balanced panel of fields (columns 1-2) and in the sample of field-years after the field first appears in the oil production data (columns 3-4). Regardless of the sample selection criteria, the results are clear: local ownership increases gas flaring by 0.53-0.62 million mscf on average, 52-64% of the control group mean. This results in an additional 29.5-36.2 thousand tonnes of CO_2 emissions per field annually. Furthermore, this increase is substantially larger than what would be accounted for simply by increased oil production.[31] The gas flaring data imply that local firms are indeed more prolific polluters. However, the negative environmental externalities of local production are consistent with both local technical disadvantage *or* political advantage. In the former case, local firms' greater costs reduce the economic viability of transporting and selling natural gas, while in the latter, local firms use political connections to evade environmental regulation.

TABLE 2 HERE

5.4 Falsification and robustness tests

5.4.1 Parallel trends

I test for divergent pre-trends using a standard event-study model, described in Section 4, estimating the model for output, malfunctions, oil theft, and violence. For each regression, I omit $\tau = -1$ as the pre-event reference year, and estimate the specification without field-level

[29] Current flaring regulations are here: https://ngfcp.dpr.gov.ng/media/1120/flare-gas-prevention-of-waste-and-pollution-regulations-2018-gazette-cleaner-copy-1.pdf

[30] In 2018, for example, the DPR reported 321 million mscf flared, 32% lower than the 472.4 estimated by the Nigeria Gas Flare Tracker using satellite data.

[31] In the 2012-2016 overlapping sample, I estimate that oil output is not significantly associated with flaring, conditional on field and year fixed effects. However, flaring operates on the extensive margin. In the overlapping sample, shut-in fields produce reduce flaring by 0.27 million mscf annually, and local fields reduce shut-ins by 19 p.p, conditional on two-way fixed effects. Therefore, through this extensive margin mechanism the upper bound on the production-driven effect is $0.27 \times 0.19 = 0.05$ million mscf, or only 7.5% of the estimated treatment effect in Table 2, columns 3-4.

controls.[32] Figure3 presents the results. Overall, the pre-trends for all outcomes appear relatively parallel across treated and control fields for all outcomes. For output, the ψ_τ for $\tau < 0$ are all negative, and significantly smaller than $\tau = -1$. Despite this, the event-study coefficients jump significantly in the immediate post-treatment period, remaining elevated more than 15 years after the treatment. For malfunctions, the ψ_τ for $\tau < 0$ are rarely significant and typically near zero. The post-takeover coefficients are generally positive although imprecisely estimated. The coefficients do not contain a jump – in contrast to the pattern for output – but rather increase more or less steadily over the years.

FIGURE 3 HERE

For both theft and violence, the event-study coefficients are insignificant and near zero in the pre-period. The post-event coefficients for oil theft incidents are negative and significant in 16 out of the 17 post-event periods. They display an initial small drop, followed by a long and sustained decline in oil theft for treated oilfields over time. Violence outcomes also witness a small initial drop, but take substantially longer to improve, with a sustained impact only emerging 6-7 years after a localization. Despite the noise of some estimates, these findings generally support that the identification assumptions necessary for DD to deliver a causal effect are likely to be satisfied.

In Figure A8, I assess whether pre-trends are parallel for the gas flaring outcome. Due to the different date range and smaller sample of the flaring data, I consider $\tau \in [-4, 5]$, collapsing all $\tau \geq 5$ into the final post-period dummy. The results generally support parallel pre-trends. None of the pre-period coefficients are significantly different from zero, while those in the post-period are consistently positive and significant. The dynamic path of the coefficients provides some evidence that the treatment effect decays to zero by $\tau \geq 5$.

5.4.2 Asset type

Offshore assets have higher technology requirements and equipment costs but are less susceptible to direct theft and violence. However, offshore platforms are susceptible to maritime piracy attacks. Onshore assets, with their accessible, unprotected pipelines, are a soft target for oil theft and attacks by criminal groups, but also comparatively easy to operate for the firm. We should therefore expect to see that reductions in criminality are concentrated in onshore assets, which benefit most from local takeover, with the exception of piracy, which should show up on offshore fields. If oil theft in turn drives the local output advantage, then we should expect that localization gains are driven primarily by the onshore assets prone to theft. In contrast, given the greater technological requirements of offshore extraction,

[32] Results are similar, though noisier, with controls.

the multinational efficiency advantage in operational malfunctions should be concentrated in more complex offshore assets.

This argument finds suggestive support in the preferences revealed by firms sorting across asset types, as shown in Figure A3 – local companies have grown 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 sorting pattern suggests that the ability to mitigate crime gives local firms a comparative advantage in onshore production, while superior technology gives multinationals the advantage offshore.

TABLE 3 HERE

Table 3 supports these hypotheses. This table replicates Panel A of Table 1, 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 onshore assets (14.9 pp vs. 6.0 pp). Similarly, the onshore output effect is roughly 2.2 million barrels, while offshore it is negative and insignificant. As predicted, local operators cause substantially more spills in offshore sites. Local takeover of an onshore field increases malfunctions by only 1.4-1.5, insignificant, while for offshore fields this number rises to 2.5-2.6, significant at 5%. The greater technological requirements of offshore extraction result in greater efficiency costs of local ownership. At the same time, the political risks of onshore extraction give rise to a comparative advantage for local firms, highlighted by the concentration of output gains onshore.

Patterns of heterogeneity in crime effects across asset types in Table 4 mirror those of output. For theft, I find that the effect of localization in offshore fields is very close to zero. In contrast, the effect of indigenization on theft is entirely concentrated in offshore fields, where the coefficient ranges from 2.8-3.4, significant at 5%. The same pattern holds for oil-related conflict deaths.[33] As expected, the maritime piracy effect is concentrated on the offshore assets which are vulnerable to this threat: local ownership reduces piracy by 0.34-0.52 attacks annually, significant at 5%.

TABLE 4 HERE

The results are summarized in Figure 4, which compares standardized localization coefficients in the onshore and offshore subsamples across all 6 outcome variables. The figure reveals that the subsample coefficients are statistically significantly different from each other for all of the outcomes with the exception of piracy – where the estimates are noisy – and malfunctions. The standardized coefficients also allow comparison of effect sizes across out-

[33] In the case of oil-related conflict deaths, it should be noted, this is by construction since very few attacks take place in offshore waters during the sample period. This is also true to a lesser extent for pipeline sabotage; 12.3% of offshore fields experience sabotage at some point, compared with nearly 87.7% of onshore fields.

comes. In standard deviations, the largest effects of localization are the onshore reduction in shut-ins, the onshore reduction in violence, and the offshore reduction in piracy.

FIGURE 4 HERE

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.

5.4.3 Divestments, transitions, and terminations

The Nigerian administrative data only contains the identity of the operating firm for active fields. This excludes a wide range of transactions that may provide valuable information. In this section I show that the main effects in Table 1 are robust to measuring local ownership directly from the DI corporate transactions data,[34] rather than inferring it from potentially incomplete data on operator identity. Furthermore, the transactions data allow for falsification tests using fields that experienced non-localizing changes in ownership or localizations terminated for exogenous bureaucratic reasons.

In Table A4, I re-estimate the main difference-in-differences equation using a treatment measure derived from the DI data. I define as treated by "divestment" all fields with a transaction in which the seller is multinational and the buyer is Nigerian in all years after the transaction was completed. Unlike the administrative data, this treatment includes all transactions in which a local firm purchased a multinational's ownership stake, even if the local firm did not become the operator. Divestment as measured by transactions data increases output by 1 million barrels per year on average, significant at 1% (columns 3-4, Panel A). This corresponds to a 35% increase relative to the control group mean, in absolute terms 54-59% as large as the coefficients estimated in Table 1.

The efficiency costs of local ownership are also robust to the transactions-based divestment measure. In Panel A columns (5)-(6), I estimate that divestment increases operational failure oil spills by 1.4-1.9 annually, or 21-28% of the control group mean, though the smaller estimate is only significant at 10%. These magnitudes are similar to those in Table 1. In Panel B I find that divestment is accompanied by a reduction in crime: oil theft falls by 4.2-4.5 incidents annually (43-46%), significant at 1%, and piracy attacks fall by 0.16-0.17 events, significant at 5%. These estimates are slightly larger than the coefficients in Table 1. However, the effects on shut-ins and oil-related violence are no longer significant, though the latter still retain the correct sign. For the primary outcomes of interest – output and oil theft – I find that

[34] For a detailed description of the DI data, see Appendix B.

parallel trends still hold (Figure A9). Indeed, the pre-period coefficients are flatter and closer to zero than the corresponding trends in Figure 3.

The DI data also allow for two distinct placebo tests. First, transactions that do not change assets' local status can rule out spurious "transition effects" whereby the observed effects are not driven by localization per se, but rather by *any* new owner revitalizing fields. In this case, we might expect to observe output and criminality effects for such transactions. Second, for 43 fields, multinational-to-local divestments were planned but either delayed or terminated for bureaucratic reasons typically exogenous to field characteristics or trends.[35] 77% of these fields were ultimately divested to local firms. If differential trends on unobservables account for the effect of localization, then this selection "effect" should also appear in fields targeted for divestment at time $\bar{t}$ but not actually divested until $\bar{t} + k$, if at all. In the absence of selection based on unobserved differential trends, we should not observe any "effect" from such terminated transactions in the k periods before actual divestment.

The results of these placebo tests are in Table 5 for the two key outcomes, oil output (columns (1)-(4)) and oil theft (columns (5)-(8)). Columns (1) and (5) test whether local-to-local transitions generate effects similar to those observed in MNC-to-local divestments, while columns (2) and (6) test MNC-to-MNC transitions. Across both outcomes, none of the placebo coefficients are significant, while the coefficients on the "true" MNC-to-local divestment indicator remain large and significant. Columns (3) and (7) include all of the transaction indicators in the same regression, with similar results.

TABLE 5

Columns (4) and (8) show that terminated and delayed divestments do not significantly affect output or criminality. Here, the treatment indicator equals one for all periods after a terminated or delayed MNC-to-local divestment is announced but before that divestment is eventually consummated, if at all. It is worth nothing that the placebo transitions and terminated divestments are somewhat rare and may be underpowered.[36] As a result, these coefficients are less precisely estimated than the main divestment coefficients. Still, the magnitudes indicate that spurious transition effects and selection based on unobserved differential trends are unlikely to be driving the results.

[35]For example, in 2011 the Shell-Total-Agip joint venture put several large oil blocks up for divestment, with the Nigerian firm Conoil the winning bidder. Subsequently, the NNPC exercised its legal right to take over operatorship and the divestment was withdrawn. The blocks were later sold separately in 2012-13 to several different Nigerian firms

[36]While there are 58 fields divested from MNC to local firm in the DI data, there are only 43, 27, and 23 affected by terminated, local-to-local, MNC-to-MNC, and divestments, respectively.

5.4.4 Additional robustness tests

In Appendix D, I test the robustness of the main results in Table 1 to additional potential sources of bias in the estimate of ψ. The effects might be driven by the state-owned oil company rather than local private firms. Table B2 shows that the effects are driven relatively more by local private companies, rather than divestment to the state-owned oil company. In addition, the state oil company drives the observed effect on malfunctions; local private fields do not appear to incur any efficiency disadvantage whatsoever.

Localized fields may respond differently to oil price fluctuations than non-localized ones. In Table B3 I interact a time-invariant treatment indicator with oil prices. Localization may also be correlated with the effects of the Niger Delta amnesty, which increased both oil theft and oil output in the region (Rexer and Hvinden 2020). I include controls for amnesty in Table B4; the effects of the amnesty, while often significant, are largely orthogonal to localization. It is also possible that the parsimonious set of interacted controls in the main results does not fully capture the full extent of differential trends across field-specific characteristics. For example, Table A1 shows that localized fields have 5 fewer wells on average, and smaller fields may be trending differently for reasons unrelated to localization. I test robustness to field-specific characteristics in Table B5, including the number of wells, field age, onshore dummy, and maximum well depth interacted with year fixed effects. I also show that the main results are unlikely to be affected by measurement error induced by multiple output data sources (Figure B1, B2) and are robust to randomization inference (Figure B6).

It is possible that while the effect of localization on theft incidents is negative, the effect on quantities stolen may not be. This might be the case if local firms experience different patterns of predation by oil theft gangs (i.e. fewer small thefts, but an increase in large thefts). We do not observe theft quantities, but as a proxy, I disaggregate total theft incidents into theft on individual asset types. Certain oil pipelines, such as "trunk" lines, are larger in diameter and lead to bigger losses. If this story is true, we would expect theft to increase on these larger assets while reductions are concentrated on smaller assets. I find no evidence that this is the case in Table B7. In Table B8, I control for location specific differential time trends in outcomes over the sample period by interacting the year dummies with state or municipality fixed effects. In effect, this restricts the difference-in-differences comparisons to within locality comparisons. I find that the results generally hold, though the effects on violence and operational failure are somewhat weaker.

Lastly, recent results from Callaway and Sant'Anna 2019, Chaisemartin and D'Haultfoeuille (2019), Goodman-Bacon (2019), and Abraham and Sun 2018 demonstrate potential sources of bias in the TWFE estimator in staggered-adoption designs that arise from aggregation of heterogeneous treatment effects. Applying these methods, I test robustness of the main estimate to decomposition of the main treatment effect (Table B9), time and unit-specific treatment effect heterogeneity (Figure B7, Table B11, Figure B11), accounting for "negative" weights (Figure B8, Table B11, Figure B11), cohort-specific heterogeneity (Figure B12), and a "stacked" difference-in-differences estimation (Table B10, Figure B10, Figure B9). Despite the presence of dynamic and cohort-specific heterogeneous effects, I find that these issues do not materially change the main TWFE and event-study results.

Figure 5 summarizes the results of 162 different robustness specifications for output (Panel A) and theft (Panel B). The specification set features all possible combinations of 3 interacted controls, 3 simultaneous shocks, 3 interacted fixed effects, 3 different treatment measures, and 2 samples.[37] These combinations are indicated in the footer of each panel.[38] I then plot the coefficients in ascending order with 95% confidence intervals and the main estimate highlighted for reference. In both cases, the main estimate is near the middle of the coefficient distribution, indicating the preferred specification is unlikely to be cherry-picked or spurious. All the estimates are of the right sign; for output, 82% of estimates are significant at 5%, while for theft 56% are significant. The insignificant cases are almost exclusively due to increased noise introduced by unnecessary controls, particularly the LGA fixed effects.

FIGURE 5 HERE

6 Theory

6.1 Model

In this section, I develop a simple model to explain the local advantage in crime reduction. I relegate the details of the model to the Online Appendix C, and here only describe the strategic environment and provide important frictions. In the model, firms bargain with state security agents over protection for assets, where the state maintains the option of accepting criminal bribes and allowing theft. Corruption is considered efficient when it generates the provision of law enforcement, which maximizes social surplus.

The interaction is a simple two-stage, one-shot game between firms, indexed by $f \in F$, gangs, indexed by $g \in G$, and state security agents, which are homogenous. The firm produces a fixed level of surplus $\bar{Q}$, sold at the international oil price p. The game proceeds as follows. Firms and gangsters simultaneously offer bribes b_f and b_g to law enforcement. If law enforcement accepts the gangster's bribe, oil theft is allowed, and enforcement $e = 0$. The

[37] This yields $2 \times 3^4 = 162$ specifications

[38] They are: *i*) with none, spatial, or field controls, *ii*) with none, price, or amnesty controls, *iii*) with none, state-, or LGA-by-year fixed effects, *iv*) with local, local-private, or divested treatment, and *v*) with full sample or onshore field only.

gangster steals a constant quantity $q < \bar{Q}$ at fixed cost $c - \epsilon_g$, where ϵ_g is private information. Theft is inefficient both because gangsters incur costs that firms don't, $pq - c + \epsilon_g < pq$ and because it directly destroys output, denoted by $\kappa > 0$. If law enforcement instead accepts the bribe b_f, then $e = 1$ and they must enforce the law. The enforcement technology reduces the probability of theft to $\alpha < 1$ at cost $c(e)$, where $c(1) = \eta > c(0) = 0$. Furthermore, since most illegal activities along the black-market value chain depend on the actual procurement of stolen oil, then under enforcement the cost of theft is reduced by a factor of α as well. All players are price takers at world oil price p.

Firms may differ in a number of 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 or reputational costs of corrupt payments. In addition, firms only receive a share γ_f of Q, to capture the important role of joint-ventures in Nigeria, as shown in Figure A1. Importantly, law enforcement may internalize firm f's output based on the parameter μ_f. This measures how the strength of political connections determines enforcement behavior. If a firm is unconnected, then $\mu_f = 0$. Note that $\gamma_f + \mu_f \leq 1$

In the base case of the model, under mild assumptions on ϵ_g, the enforcement decision of the security forces is summarized in the following equation:

$$Pr(e = 1) = \frac{1}{c}[(\gamma + \mu - 1)pq + (\gamma + \mu)p\kappa] - \frac{\lambda_f + \eta}{(1 - \alpha)c} + 1$$

Government enforcement increases in μ, the measure of political connections, as well as γ, the ownership share. In both cases this is because participants – government and the firm, respectively – internalize a larger share of the losses from theft. An increase in q reduces enforcement, since it increases the value of the outside option of criminal bribes more than it increases the collective costs of theft to government and firm. Enforcement also declines in η and increases in government effectiveness $1 - \alpha$. Importantly, the prediction on prices is ambiguous: if $\frac{\kappa}{q} > \frac{(1-\gamma-\mu)}{(\gamma+\mu)}$, a positive price shock will only increase enforcement if the resulting theft losses are sufficiently large relative to the increase in value of the criminal bribe. Given the enforcement technology, the comparative statics for theft are identical with signs reversed. The central point is that inefficient theft may occur in equilibrium given the substantial frictions in this bargaining interaction. However, political connections induce government to internalize the losses from theft, lowering the cost of bribery and increasing the scope for efficient corruption.

However, the base case (in Appendix C.2), perhaps optimistically, that all parties can perfectly commit. In reality, however, law enforcement may accept bribes from firms and allow theft nonetheless, given the cost of enforcement and the lure of black-market rents. I show in Appendix C.3 that commitment problems substantially increase the importance of political connections. Government will only enforce the law when the firm is sufficiently connected:

$$\mu > \bar{\mu} = \frac{(1-\alpha)(pq-c)+\eta}{(1-\alpha)p(q+\kappa)}$$

Which is strictly positive as long as theft is profitable for criminals. Commitment problems imply that bribes alone will not suffice; enforcement will not be provided in absence of political connections that give state security agents a stake in the profits of the firm. This stark result motivates the subsequent focus on political connections in Section 7. Finally, in Appendix C.4 I show that a repeated interaction restores some of the scope for efficient corruption with bribes when $\mu < \bar{\mu}$, as long as the parties are sufficiently patient. In this case, law enforcement must be provided with an additional rent to maintain commitment. The key comparative statics remain.

7 Mechanisms

The security forces in Nigeria are known to be involved in organized crime and facilitating black market activity (SDN 2019a, SDN 2019b, Asuni 2009). The model in Section 6 shows that if local firms have better political connections, this will enable security agencies to internalize the losses from theft, resulting in better law enforcement protection, lower crime, and ultimately improved output performance. In this section, I present several pieces of evidence for this mechanism. I first show that local takeover leads to large gains in law enforcement provision, suggesting that local firms do experience preferential state provision of security. Next, I show that $i)$ local firms are better connected politically, and $ii)$ these political connections are associated with lower theft, with the largest returns for connections to the security forces. I focus primarily on political connections because this is the binding constraint to enforcement in a no-commitment environment. However, I also provide that lower corruption penalties, λ, and greater ownership shares γ, provide additional advantages to local firms.

7.1 Law enforcement corruption

When black markets generate large rents, law enforcement agencies to have incentives to accept bribes in exchange for protection of criminal activity. What differentiates our context from other black markets, however, is that private firms likewise have strong incentives to pay security agents for protection. Law enforcement corruption therefore presents a novel mechanism by which political connections can affect firm outcomes – connections align the incentives of law enforcement to privately provide protection from organized criminal activity. This corrupt marketplace can explain local advantage – better-connected local firms obtain superior law enforcement protection, leading to lower black market losses and ultimately better output performance.

To test the hypothesis that local takeover increases enforcement activity, I use data on anti-oil theft law enforcement actions taken from Nigerian news media reports, described in Section 3.4. I estimate the staggered-DD divestment model with the count of all anti-oil theft law enforcement activities undertaken by the Nigerian state in a given field-year as the outcome variable. The results of this estimation are in Table 6. Column (1) estimates the baseline TWFE specification, while column (2) includes interacted geographic controls, column (3) adds controls for field characteristics, column (4) includes field-specific linear trends, and (5) and (6) include state- or municipality-by-year fixed effects. Local takeover leads to a large and statistically signifiant increase in enforcement provision. Anti-oil theft law enforcement actions increase by 2.9-5.1 events (roughly 83-146% greater than the control group mean), significant at 5 or 10% in all except one specification. Nigerian security forces provide preferential law enforcement protection to local firms.

TABLE 6 HERE

The time-path of enforcement activities also provides additional support for the mechanism. Figure 6 plots event-study coefficients using several different enforcement outcomes. The top left panel estimates the model for all anti-oil theft enforcement. The dynamic coefficients indicate parallel trends in enforcement the years leading up to divestment. Enforcement activity spikes in the first year of divestment and then falls gradually thereafter, with positive coefficients until 10 years after divestment and zeroes thereafter. Combining this plot with the event-study results for oil theft in Figure 3, bottom left panel, we can conclude that an immediate spike in enforcement upon local takeover precedes a long-run reduction in crime, which in turn leads to an eventual tapering of enforcement activity as crime falls.

FIGURE 6 HERE

In the remaining panels of Figure 6, I consider several additional subcategories of anti-oil theft enforcement: seizures of stolen oil, raids of illegal refineries, and raids on illegal export activities. I disaggregate these outcomes because enforcement patterns may plausibly differ along the black market value chain. For example, export activities are controlled by powerful militant groups with connections to organized crime, while refineries are primarily the purview of smaller-scale entrepreneurs. However, all these types of oil crime exhibit very similar dynamics – an initial spike in enforcement, followed by a gradual tapering. In FigureA10 I plot DD coefficient estimates for these and several other types of illegal activities. Coefficients are positive and significant for oil seizures, illegal refining, and illegal export, all highly visible theft activities that thrive under state protection. Estimates are also positive – though smaller and less significant – for piracy, an adjacent illegal activity often carried out by similar criminal groups. I also consider two placebo outcomes – enforcement on militancy and on non-oil crime (for example, armed robbery or fraud). Consistent with oil sector law enforcement corruption as the key mechanism driving local advantage, I find that that localization has no effect on these placebo enforcement outcomes that are either unrelated to the oil sector or less manipulable by security agents.

7.2 Political connections

To obtain greater law enforcement protection, local firms must have better political connections. Table A5 demonstrates that local firms are indeed more likely to cultivate political connections, particularly among the security forces. Local firms are less constrained by meritocratic hiring practices, legal requirements, or reputational risks. In addition, while multinationals are subsidiaries owned by foreign firms that are publicly traded, local firms are much more likely to be privately held by politically-connected individuals and families.

To provide additional evidence for the political connections channel, I first test whether political connections are indeed important determinants of theft. Then, I disaggregate political connections to identify whether effects are heterogeneous across types of connections – in particular concentrated among the security forces. Table 7 contains the results of the TWFE regression of theft on political connections.[39] 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. "Technocrats" are those associated with ministerial posts or regulatory agencies, typically the NNPC, DPR, or Ministry of Petroleum Resources. "Cabinet-level politician" indicates that the company operating field i in state s is connected to a politician who at some point served in a ministerial post. Lastly, "security forces" are those linked to the military or police forces.

TABLE 7 HERE

Each pair of columns in Table 7 indicates the impact of a specific type of connection, estimated with and without control variables. Having any political connection reduces field-level theft by 2.9 incidents per year on average (columns (1)-(2)), although these effects are only significant at 5% without controls. The effects of technocratic politicians (columns (3)-(4)) are also negative and significant. More prominent politicians provide a larger advantage; cabinet-level connections result in between 3.9-6 fewer theft incidents, the latter of which is significant at the 10% level. However, by far the most pronounced effects are for connections to the security forces (columns (7)-(8)), at 7.9-9.7, respectively, all of which are significant at 1%. These effect sizes are summarized in Figure 7, which plots the coefficients and confidence

[39] The sample here is only 3236 field-years because of the three firms with missing political connections data.

intervals for each category of political connection. Inclusion of controls only slightly affects the estimates.

FIGURE 7 HERE

Clearly, politically connected firms are better able to mitigate the activities of organized crime. Importantly, the effects are largest for high-level politicians and security forces. This suggests that connections to high-level politicians give firms access to the 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 by security forces to enforce the law or collude with organized crime is the primary determinant of the viability of theft, connected firms are able to divert theft away from their assets.

7.3 Additional mechanisms

The model also implies that local advantage may arise if local firms face lower corruption costs λ, or hold greater profit shares γ. In this section, I show suggestive evidence that both of these forces also play a role in local advantage. I also provide evidence that the results are driven, at least partially, by spatial spillovers of crime to nearby un-divested fields. In Online Appendix G.1, I also provide additional results on heterogeneous treatment effects consistent with the predictions of the model.

Ownership consolidation: Joint ownership drives a wedge between the losses to the operating firm and criminal profits; operators with larger ownership stakes γ internalize a greater share of the losses from theft. Because of indigenization policies and natural consolidation of stakes during divestments, local firms typically hold more equity in their assets. Multinationals are 33.5 p.p. more likely to be in joint ventures and 43 p.p. less likely to obtain sole-risk licenses. As a result, the average multinational concession has a government stake roughly 85% higher than the average Nigerian independent operator. In Appendix E.1, I use oil block licensing data to show that divestment causally increases the concession-level ownership Herfindahl Index by 16.7%, and increases the ownership stake of the operating firm by 12.8% p.p., or 20.1% (see Table B12). This is consistent with higher γ providing local firms with stronger incentives to bargain with criminal groups.

Corruption penalties: Multinational firms may face higher expected costs λ of engaging in corrupt behavior because of home-country anti-corruption statutes, such as the Foreign Corrupt Practices Act (FCPA) in the United States, that prohibit improper payments to foreign officials. 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. In Appendix E.2, I test whether exposure to an international corruption law affects outcomes, using variation in timing of passage of these laws for identification. I find that passage of a home-country corruption law is associated with a statistically significant annual increase in theft of 2.7-6.7 incidents, or 24.4-58.6% of the mean, in the sample of multinational-operated fields (see Table B13).[40] Despite the small sample of law changes, this analysis provides suggestive evidence that anti-corruption laws affect multinationals' ability to use corrupt payments to deter crime.

Spatial spillovers: In general equilibrium, gangs may may not operate as local monopolists but rather choose targets for theft across all oil fields. As such, localization could increase targeting of surrounding multinational fields if local fields are politically protected but their multinational neighbors are not. Since spillovers may bias the treatment effect by violating the stable unit treatment value assumption (SUTVA) (Rubin 2005), it is important to estimate the treatment effect purged of spillovers. Furthermore, spillovers have important implications for the aggregate productivity effects of localization. In Appendix F.1, I estimate substantial negative output, theft, and conflict spillovers for fields between 30-40 km from the treated field (see Figure B13). However, even after accounting for these spillovers the main treatment effects remain significant.

7.4 Alternative explanations

Local employment spillovers: Localization may generate positive economic spillovers to local labor markets. If these spillovers improve employment opportunities for young men, this may bid up the opportunity cost of joining gangs and therefore increase the gangster's cost c. In this case, the crime-reduction effects might be driven by higher labor costs in the criminal sector. To test this hypothesis, I use data from three rounds of Nigeria's General Household Panel Survey on 16,211 working-age Nigerians in 500 villages from 2010-2016, linking each village to its nearest oilfield. In Appendix F.2, I estimate the effect of localization of nearby fields on employment and consumption outcomes using TWFE. I find no evidence of employment spillovers from local ownership (see Tables B14, B15 and Figures B14, B15).

Targeted CSR investment: The most visible local benefits of oil extraction are typically not jobs but rather host community investments in the form of corporate social responsibility (CSR). Oil companies may prefer to provide CSR benefits to troubled areas to dissuade militancy and theft than to negotiate with organized crime directly. If local firms are more likely to engage in this practice, this could account for localization benefits. In 2016, oil companies' expenditures on CSR projects in host communities totaled 92.6 million USD, 72% of which was spent by multinationals. Since this expenditure is a small fraction of the annual profits from oil theft, these projects are unlikely to meaningfully dissuade crime and violence. However, in Appendix F.3, I use cross-sectional data on oil company CSR projects in 2016 to

[40] Restricting the sample to multinationals also allows me to remain agnostic about the content, quality, and enforcement of Nigeria's own anti-corruption laws. This is preferable to assessing the effectiveness of these laws, which legal analysis suggest are basically ineffective (Aigbovo and Atsegbua 2013).

test whether local firms are better at targeting their investments toward volatile communities. Though these cross-sectional correlations are somewhat speculative, I find that multinational CSR investments are more responsive to recent levels of local oil-related militant violence (see Figure B16), suggesting that local advantage is not driven by this mechanism.

Differences in discount rates: The local advantage in production may be driven not by organized crime and differential bargaining frictions, but rather by different optimal extraction profiles given underlying time preferences. This is a plausible mechanism if local companies have shorter time horizons. This is unlikely to be the case, as oil output is difficult to adjust along the intensive margin in the short-run for a given stock of fixed capital. Most of the increase in output we observe from localization can be attributed mechanically to bringing more fields into production and reducing the quantity lost to theft. However, the discount rate mechanism has a clear empirical implication: longer time horizons should dampen the short-run elasticity of production to increases in the stock of available oil reserves once exploration costs are sunk, so local companies should exhibit greater short-term responses to a positive reserve shock. I test this hypothesis using firm-level data on 49 new reserve discoveries across 23 firms from 2001-2016 in a stacked-DD event-study model with a 5-year symmetric event-window. I regress log output on log size of discovery interacted with pre-and-post event dummies so that the coefficients are easily interpreted as elasticities and comparable across companies of differing sizes. Figure A11 shows that neither multinationals nor local firms are particularly responsive to new discoveries, with post-event dynamic elasticities very near to zero.

Grievance toward multinationals: Criminal and militant activity may be driven by grievance rather than economic motives (Buhaug, Cederman, and Gleditsch 2014). Niger Deltans retain longstanding, justified grievances against multinationals due to a long history of corporate malfeasance and environmental pollution (Obi and Rustad 2011). Sentiments toward local companies may be considerably better, resulting in reduction in grievance-driven attacks and productivity gains. If so, we should expect to observe a reduction in community protest, the most direct expression of grievance. Protests against oil companies – generally peaceful but occasionally riotous – are common in host communities, affecting 26% of all fields during the sample period. In Table A6, I re-estimate the main specification using the number of protests (columns 1-2), oil-related protests (columns 3-4), and riots (columns 5-6) within 15 kilometers of the field as the outcome variable. The point estimates are, if anything, positive, but generally insignificant. There is no evidence of a change in grievance as a result of localization.

8 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.

In Nigeria's oil sector, where militant groups and organized crime are ever-present threats to firm operations and corruption buys law enforcement protection for assets, local companies posses distinct advantages. Using data on Nigerian 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 increases output by 1.6 million barrels per year, a 60% gain. Local firms accomplish this feat in part by reviving moribund fields: the likelihood that a field is nonproducing falls dramatically upon local takeover. However, consistent with a technical efficiency disadvantage, local fields show evidence of lower operational standards. They experience more oil spills from equipment malfunctions and flare more natural gas. Both of these practices entail substantial environmental costs.

The key to the local output advantage is in dealing with the multi-billion dollar black market for stolen oil. I find that local takeovers reduce oil theft and militant violence substantially. 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 on offshore fields with high technological requirements. This further underscores that while multinationals have a technology advantage, the black market generates a much larger local advantage. Placebo tests using data on the universe of Nigerian oil and gas corporate transactions rule out spurious transition effects or unobserved differential trends.

A model of the bargaining interaction between firms and law enforcement shows that better political connections, lower corruption penalties, and greater ownership stakes allow local firms to expand the scope for efficient corruption, purchase law enforcement protection, and outperform multinationals despite their quality disadvantage. In the no-commitment case, political connections are a binding constraint to law enforcement provision, and unconnected firms are left to the mercy of rapacious armed groups.

To test whether this drives local advantage, I first show that local firms indeed receive preferential state protection: localization leads to an 80% increase in law enforcement activity. Using data on firm connections to politicians, technocrats, and the security forces, I demonstrate substantial returns to political connections in reduced theft. Consistent with the model, connections to the security forces – the agencies at the center of corruption in the black market for stolen oil – are particularly beneficial. Local firms are more likely to possess the connections to security agents required to protect assets. In addition, I find evidence that exposure to a home-country anti-corruption law is associated with greater multinational predation, suggesting that the cost of legal liability generates a wedge in the bargaining process. Finally, I find that empirical patterns of heterogeneity with respect to oil prices and local militant capacity are consistent with the predictions of the model.

The findings suggest that when political and social conflict in the natural resource sector is extreme, localization gains in output may be large enough to outweigh the loss of multinational productivity and therefore justify indigenization policies on efficiency grounds. More broadly, the results support the notion that efficient corruption – in which local firms have a comparative advantage – may indeed be welfare-improving given a particular set of second-best institutional constraints. However, I find no evidence that indigenization of oil assets improved local employment prospects in oil-producing communities, tempering optimism that local ownership will fundamentally alter the enclave nature of oil extraction. Still, since negative spillovers to unconnected multinational firms are possible and the environmental costs of local ownership are substantial, the results cannot rule out that localization simply represents a transfer of rents to connected local elites rather than a true welfare gain. I leave a thorough welfare analysis of indigenization policies for future work.

The Nigerian oil and gas sector may well be a representative rather than an extreme case. In extractive sectors across the globe – from Congolese minerals to Colombian gold – firms face a complex political economy characterized by black markets, organized crime, armed groups, and corrupt politicians. The conventional economic wisdom on multinational productivity from manufacturing and service firms in middle-income countries does not apply to natural resource sectors in poor ones, where resource rents comprise 12% of GDP. Instead, we must seriously consider the productivity gains and the environmental costs of indigenization.

References

Abraham, Sarah and Liyang Sun (2018). "Estimating Dynamic Treatment Effects in Event Studies With Heterogeneous Treatment Effects". In: *Working Paper*.

Aigbovo, Osaretin and Lawrence Atsegbua (2013). "Nigerian anti-corruption statutes: an impact assessment". In: *Journal of Money Laundering Control* 16.1, pp. 62–78.

Aitken, Brian and Ann Harrison (1999). "Do Domestic Firms Benefit from Direct Foreign Investment? Evidence from Venezuela". In: *American Economic Review* 89.3, pp. 605–618.

Akcigit, Ufuk, Salome Baslandze, and Francesca Lotti (2018). "Connecting to Power: Political Connections, Innovation, and Firm Dynamics". In: *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*. Ed. by Francisco L Rivera-Batiz, Mariana Spatareanu, and Can Erbil. Singapore: World Scientific.

Alfaro, Laura and Maggie X. Chen (2018). “Selection and Market Reallocation: Productivity Gains from Multinational Production”. In: *American Economic Journal: Economic Policy* 10.2, pp. 1–38.

Aragon, Fernando M., Punam Chuhan-Pole, and Bryan Christopher Land (2015). “The Local Economic Impacts of Resource Abundance: What Have We Learned?” In: *World Bank Policy Research Working Paper* 7263.

Aragon, Fernando M. and Juan Pablo Rud (2013). “Natural Resources and Local Communities: Evidence from a Peruvian Gold Mine”. In: *Journal of International Money and Finance* 5.2, pp. 1–25.

— (2011). “Polluting Industries and Agricultural Productivity: Evidence from Mining in Ghana”. In: *Economic Journal* 126.November, pp. 1980–2011.

Arnold, Jens Matthias and Beata S. Javorcik (2009). “Gifted kids or pushy parents? Foreign direct investment and plant productivity in Indonesia”. In: *Journal of International Economics* 79.1, pp. 42–53.

Asuni, Judith Burdin (2009). “Understanding the Armed Groups of the Niger Delta”. In: *Council on Foreign Relations Working Paper*.

Autor, David H., Christopher J. Palmer, and Parag A. Pathak (2014). “Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts”. In: *Journal of Political Economy* 122.3, pp. 661–717.

Balsvik, Ragnhild (2011). “Is Labor Mobility a Channel for Spillovers from Multinationals? Evidence from Norwegian Manufacturing”. In: *The Review of Economics and Statistics* 93.1, pp. 285–297.

Berman, Nicolas et al. (2017). “This Mine is Mine! How Minerals Fuel Conflicts in Africa”. In: *American Economic Review* 107.6, pp. 1564–1610.

Blair, Graeme, Darin Christensen, and Valerie Wirtschafter (2019). “How Does Violence Shape Investment? Evidence from Mining”. In: *Working Paper*.

Blair, Graeme and Kosuke Imai (2013). “Civilians and the Strategic Use of Information During Conflict in Resource-Rich Territory”. In: *International Growth Centre Working Paper*.

Bloom, Nicholas and John Van Reenen (2010). “Why Do Management Practices Differ across Firms and Countries?” In: *Journal of Economic Perspectives* 24.1, pp. 203–224.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen (2012). “Americans Do IT Better: US Multinationals and the Productivity Miracle”. In: *American Economic Review* 102.1, pp. 167–201.

Bruederle, Anna and Roland Hodler (2019). “Effect of oil spills on infant mortality in Nigeria”. In: *PNAS* 116.12, pp. 5467–5471.

Buhaug, Halvard, Lars-Erik Cederman, and Kristian Skrede Gleditsch (2014). “Square Pegs in Round Holes: Inequalities, Grievances, and Civil War”. In: *International Studies Quarterly* 58.2, pp. 418–431.

Buonanno, Paolo et al. (2015). "Poor Institutions, Rich Mines: Resource Curse in the Origins of the Sicilian Mafia". In: *Economic Journal* 125.August, pp. 175–202.

Burger, Martijn, Elena Ianchovichina, and Bob Rijkers (2015). "Risky Business: Political Instability and Sectoral Greenfield Foreign Direct Investment in the Arab World". In: *World Bank Economic Review* 30.2, pp. 306–331.

Callaway, Brantly and Pedro H. C. Sant'Anna (2019). "Difference-in-Differences with Multiple Time Periods". In: *Working Paper*.

Chaisemartin, Clement de and Xavier D'Haultfoeuille (2019). "Two-way Fixed Effects Estimators with Heterogeneous Treatment Effects". In: *NBER Working Paper* 25904.

Christensen, Darin (2019). "Concession Stands: How Mining Investments Incite Protest in Africa". In: *International Organization* 73.1, pp. 65–101.

Couttenier, Mathieu, Pauline Grosjean, and Marc Sangnier (2017). "The Wild West is Wild: The Homicide Resource Curse". In: *Journal of the European Economic Association* 15.3, pp. 558–585.

Criscuolo, Chiara and Ralf Martin (2009). "Multinationals and U.S. Productivity Leadership: Evidence from Great Britain". In: *The Review of Economics and Statistics* 91.2, pp. 263–281.

Cust, James and Steven Poelhekke (2015). "The Local Economic Impacts of Natural Resource Extraction". In: *Annual Review of Resource Economics* 7, pp. 251–268.

Deshpande, Manasi and Yue Li (2019). "Who Is Screened Out? Application Costs and the Targeting of Disability Programs". In: *American Economic Journal: Policy* Forthcoming.

Diamond, Rebecca and Tim McQuade (2019). "Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development". In: *Journal of Political Economy* 127.3, pp. 1063–1117.

Dube, Oeindrila and Juan F. Vargas (2013). "Commodity Price Shocks and Civil Conflict: Evidence from Colombia". In: *Review of Economic Studies* 80, pp. 1384–1421.

Faccio, Mara (2006). "Politically Connected Firms". In: *American Economic Review* 96.1, pp. 369–386.

Fetzer, Thiemo and Stephan Kyburz (2018). "Cohesive Institutions and Political Violence". In: *Centre for Competitive Advantage in the Global Economy* Working Paper.

Fisman, Raymond (2001). "Estimating the Value of Political Connections". In: *American Economic Review* 91.4, pp. 1095–1102.

García-Jimeno, Camilo (2016). "The Political Economy of Moral Conflict: An Empirical Study of Learning and Law Enforcement Under Prohibition". In: *Econometrica* 84.2, pp. 511–570.

Goodman-Bacon, Andrew (2019). "Difference-in-Differences with Variation in Treatment Timing". In: *Working Paper*.

Gormley, Todd A. and David A. Matsa (2011). "Growing Out of Trouble? Corporate Responses to Liability Risk". In: *Review of Financial Studies* 24.8, pp. 2781–2821.

Guadalupe, Maria, Olga Kuzmina, and Catherine Thomas (2012). “Innovation and Foreign Ownership”. In: *American Economic Review* 102.7, pp. 3594–3627.

Guidolin, Massimo and Eliana La Ferrara (2007). “Diamonds Are Forever, Wars Are Not: Is Conflict Bad for Private Firms?” In: *American Economic Review* 97.5, pp. 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*. Ed. by Dani Rodrick and Mark R. Rosenzweig. Elsevier B.V. Chap. 63, pp. 4040–4214.

Hodgson, R. (2018). “Generating a scalable calibration equation that can be applied to VIIRS Nightfire (VNF) radiant heat calculations to estimate gas flaring volumes in Nigeria”. In: *PhD Thesis, Department of Geography, Birkbeck University*.

Javorcik, Beata S. and Shang-Jin Wei (2009). “Corruption and cross-border investment in emerging markets: Firm-level evidence”. In: *Journal of International Money and Finance* 28.4, pp. 605–624.

Khwaja, Asim Ijaz and Atif Mian (2005). “Do Lenders Favor Politically Connected Firms? Rent Provision in an Emerging Financial Market”. In: *The Quarterly Journal of Economics* 120.4, pp. 1371–1411.

Kyburz, Stephan (2018). “The Local Political Resource Curse”. In: *Working Paper*.

Li, Hongbin et al. (2008). “Political Connections, financing, and firm performance: Evidence from Chinese private firms”. In: *Journal of Development Economics* 87, pp. 283–299.

Lippert, Alexander (2014). “Spill-Overs of a Resource Boom: Evidence from Zambian Copper Mines”. In: *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”. In: *World Development* 84.C, pp. 219–234.

NBS (2017). *Demographic Statistics Bulletin*. Tech. rep. Abuja, Nigeria: Nigeria National Bureau of Statistics.

NEITI (2016). *2016 Oil and Gas Industry Audit Report*. Tech. rep. Abuja, Nigeria: Nigeria Extractive Industries Transparency Initiative.

Nwokolo, Arinze (2018). “Oil Price Shocks and Civil Conflict: Evidence from Nigeria”. In: *HiCN Working Paper* 274.

Obi, Cyril and Siri Aas Rustad (2011). *Oil and Insurgency in the Niger Delta. Managing the Complex Politics of Petroviolence*. Zed Books.

Ologunorisa, Temi E. (2009). “A review of the effects of gas flaring on the Niger Delta environment”. In: *International Journal of Sustainable Development and World Ecology* 8.3, pp. 249–255.

Poole, Jennifer P. (2013). “Knowledge Transfers from Multinational to Domestic Firms: Evidence from Worker Mobility”. In: *The Review of Economics and Statistics* 95.2, pp. 393–406.

Ralby, Ian (2018). *Downstream Oil Theft: Global Modalities, Trends, and Remedies*. Tech. rep. The Atlantic Council.

Rexer, Jonah and Even Hvinden (2020). "Delta boys: Bargaining, war, and black market oil in Nigeria". In: *Working Paper*.

Roth, Jonathan (2019). "Pre-test with Caution: Event-study Estimates After Testing for Parallel Trends". In: *Working Paper*.

Rubin, Donald B. (2005). "Causal Inference Using Potential Outcomes: Design, Modeling, Decisions". In: *Journal of the American Statistical Association* 100.469, pp. 322–331.

SDN (2019a). *Communities not Criminals: Illegal Oil Refining in the Niger Delta*. Tech. rep. Port Harcourt, Nigeria: Stakeholder Democracy Network.

— (2019b). *Economic Dynamics of the Artisanal Oil Industry in the Niger Delta over Five Years*. Tech. rep. Port Harcourt, Nigeria: Stakeholder Democracy Network.

— (2019c). *Pipeline surveillance contracts in the Niger Delta*. Tech. rep. Port Harcourt, Nigeria: Stakeholder Democracy Network.

Sexton, Renard (2019). "Unpacking the Local Resource Curse: How Externalities and Governance Shape Social Conflict". In: *Journal of Conflict Resolution* Forthcoming.

Sierra, Raul Sanchez de la (2020). "On the Origin of States: Stationary Bandits and Taxation in Eastern Congo". In: *Journal of Political Economy* Forthcoming.

Teece, David (1977). "Technology Transfer by Multinational Firms: The Resource Cost of Transferring Technological Know-how". In: *Economic Journal* 87.346, pp. 242–261.

Watts, Michael (2007). "Petro-insurgency or Criminal Syndicate: Conflict and Violence in the Niger Delta". In: *Review of African Political Economy* 34.114, pp. 637–660.

Table 1: The effect of divestment on output and criminality

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Output and efficiency						
Outcome	Shut-in		Output		Malfunctions	
Local operator	-0.167**	-0.163**	1.689**	1.884***	1.519*	0.991
	(0.068)	(0.066)	(0.722)	(0.545)	(0.907)	(0.976)
Control group mean	0.244		2.835		6.864	
Observations	2464	2464	2464	2464	3497	3497
R^2	0.657	0.670	0.861	0.878	0.573	0.631
Panel B: Crime and violence						
Outcome	Theft		Violence		Piracy	
Local operator	-3.390***	-3.483***	-0.699***	-0.732**	-0.135	-0.107
	(1.141)	(1.306)	(0.262)	(0.293)	(0.095)	(0.084)
Control group mean	10.172		0.399		0.150	
Observations	3497	3497	3497	3497	3497	3497
R^2	0.714	0.755	0.130	0.155	0.235	0.313
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

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 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. Violence is the total number of oil-related conflict 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 2: The effect of divestment on gas flaring

Sample	Balanced panel		Post-first year	
	(1)	(2)	(3)	(4)
Local operator	0.539**	0.578**	0.660**	0.662**
	(0.250)	(0.254)	(0.332)	(0.331)
Field FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	No	Yes
Control group mean	1.096		1.096	
Observations	2512	2512	2445	2445
R^2	0.811	0.817	0.807	0.813

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 oilfields from 2012-2019. The sample is either a balanced panel, in columns (1)-(2), or field-years only after a field first appears in the NNPC/DPR production data, columns (3)-(4). Outcome variable is annual gas flaring on the field, measured in million mscf. 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	
Panel A: Onshore fields						
Local operator	-0.198**	-0.149**	2.198***	2.231***	1.543	1.385
	(0.079)	(0.070)	(0.620)	(0.581)	(0.989)	(0.938)
Control group mean	0.293		1.412		6.794	
Observations	1729	1729	1729	1729	2518	2518
R^2	0.658	0.685	0.795	0.809	0.604	0.686
Panel B: Offshore fields						
Local operator	-0.033	0.060	-1.496	-0.560	2.463*	2.660**
	(0.031)	(0.065)	(2.423)	(1.447)	(1.381)	(1.086)
Control group mean	0.134		6.046		7.037	
Observations	735	735	735	735	979	979
R^2	0.628	0.663	0.860	0.899	0.571	0.606
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

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 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. Violence is the total number of oil-related conflict 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		Violence		Piracy	
Panel A: Onshore fields						
Local operator	-3.369**	-2.807*	-1.006***	-1.285**	-0.046	-0.078
	(1.396)	(1.459)	(0.369)	(0.504)	(0.100)	(0.090)
Control group mean	14.234		0.559		0.150	
Observations	2518	2518	2518	2518	2518	2518
R^2	0.707	0.745	0.134	0.206	0.260	0.363
Panel B: Offshore fields						
Local operator	0.027	-0.048	-0.010	0.025	-0.523*	-0.339**
	(0.026)	(0.044)	(0.010)	(0.028)	(0.273)	(0.148)
Control group mean	0.043		0.000		0.151	
Observations	979	979	979	979	979	979
R^2	0.180	0.393	0.093	0.124	0.218	0.321
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

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 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. Violence is the total number of oil-related conflict 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 divestment on output and oil theft, placebo tests

Outcome	Output				Oil theft			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
MNC-to-local divestment	0.987***	0.935***	0.888***		-4.659***	-4.438***	-4.599***	
	(0.335)	(0.274)	(0.280)		(0.932)	(0.898)	(0.927)	
Local-to-local sale	0.272		0.334		1.072		1.030	
	(0.873)		(0.866)		(1.090)		(1.098)	
MNC-to-MNC sale		0.626	0.649			-0.461	-0.376	
		(0.724)	(0.720)			(1.243)	(1.255)	
Terminated divestment				0.146				0.086
				(0.284)				(1.521)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2464	2464	2464	2464	3497	3497	3497	3497
R^2	0.878	0.878	0.878	0.877	0.756	0.756	0.756	0.754

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016 for which output information is available. "MNC-to-local divestment" is an indicator that equals one in all periods after a field is sold from a multinational to a Nigerian buyer, as derived from DrillingInfo transactions data. Outcome variable is indicated in the table header. 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. 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 6: The effect of local ownership on law enforcement activity

Outcome	Anti-oil theft law enforcement					
	(1)	(2)	(3)	(4)	(5)	(6)
Local operator	5.138**	3.681**	3.525	3.858**	4.585**	2.940*
	(2.186)	(1.845)	(2.160)	(1.942)	(2.085)	(1.547)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	No	No
Field FE × linear trend	No	No	No	Yes	No	No
State × Year FE	No	No	No	No	Yes	No
Locality × Year FE	No	No	No	No	No	Yes
Controls × Year FE	No	Yes	No	Yes	No	No
Field controls × Year FE	No	No	Yes	No	No	No
Control group mean	3.496					
Observations	3497	3497	3316	3497	3497	3497
R^2	0.543	0.657	0.582	0.765	0.640	0.876

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016 for which political connections data is available. Outcome variable is anti-oil theft law enforcement, the total count of events within 20 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 7: The effect of political connections on oil theft

Outcome	Oil theft							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Any politician	-2.901**	-2.872*						
	(1.322)	(1.505)						
Technocrat			-3.613**	-4.649**				
			(1.746)	(1.881)				
Cabinet-level politician					-3.993	-6.020*		
					(3.020)	(3.499)		
Security forces							-7.926***	-9.745***
							(0.918)	(2.896)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls $\times$ Year FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	3236	3236	3236	3236	3236	3236	3236	3236
R^2	0.713	0.754	0.713	0.754	0.713	0.753	0.713	0.753

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 1: Indigenization and offshoring

(a) Onshore share

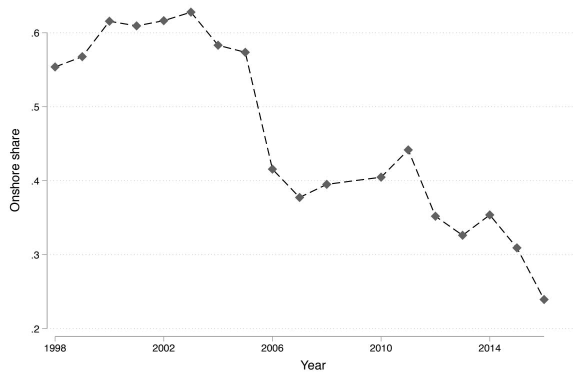

(b) Output share by company type

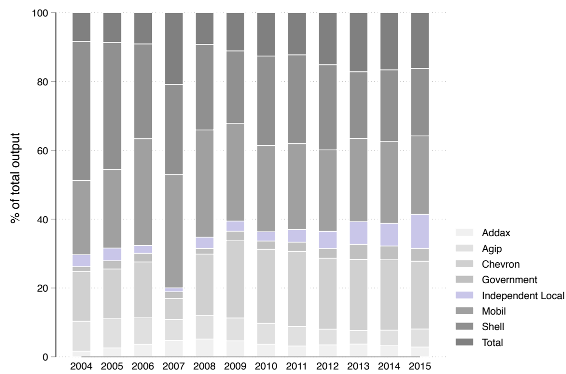

Note: Figure shows the share of total oil output produced onshore (Panel A), and by different companies (Panel B), over time. Onshore share is all output produced in onshore assets. Company categories are Addax, Agip, Chevron, Mobil, Shell, Total, the state-owned oil company, and independent private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

Figure 2: Map of treatment status

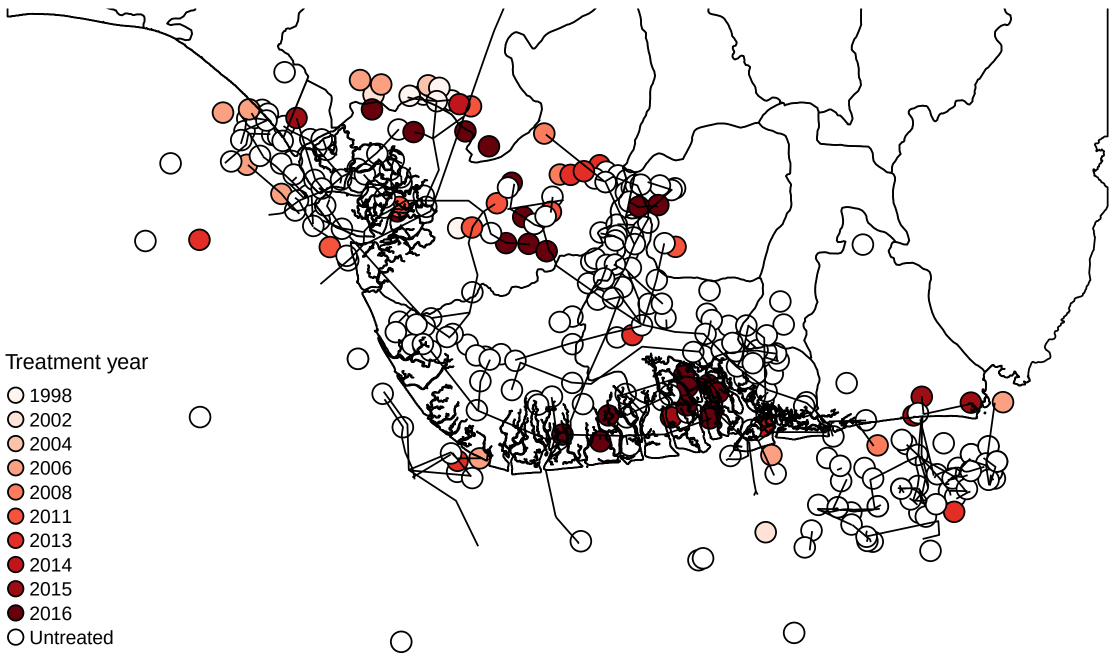

Note: Figure maps the centroids of 314 active Nigerian oilfields. Marker color indicates the year of local takeover of the field. White markers are never-treated fields. Basemap is Nigerian states of the Niger Delta region, while lines indicate oil pipelines.

Figure 3: Event study plots, main DD specification

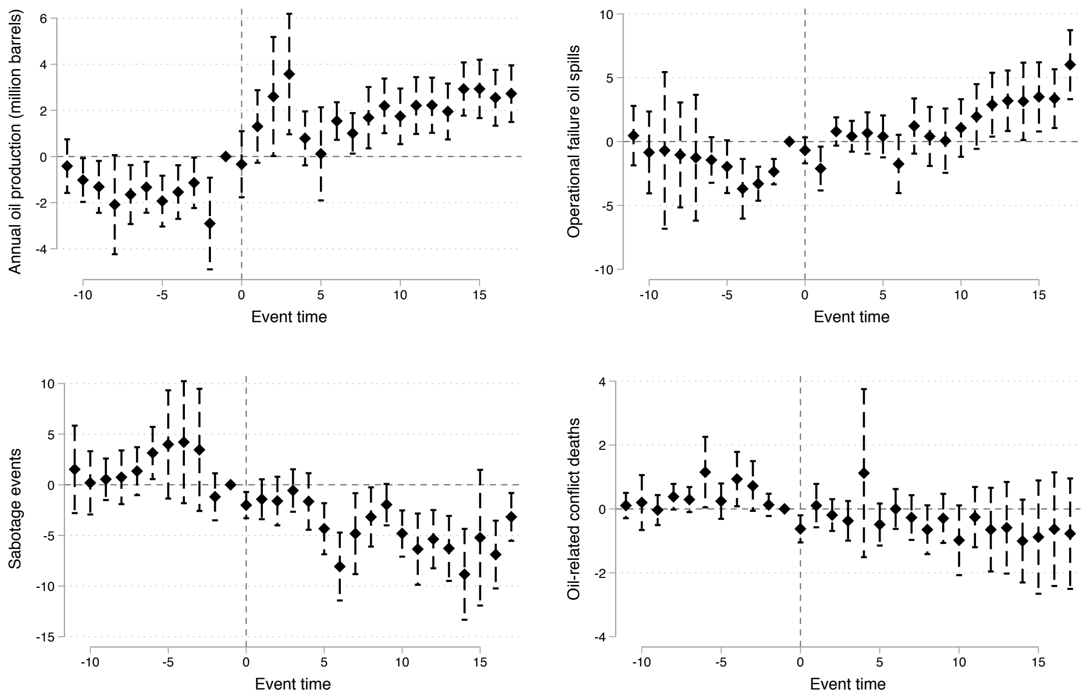

Note: Figure displays coefficients of event-study regressions of outcomes on pre-and-post treatment indicators, conditional on unit and year fixed effect and controls interacted with year dummies. Malfunctions are the total number of non-sabotage spills within 15 km of the field. Theft is the total number of sabotage spills within 15 km of the field. Violence is the total number of oil-related conflict deaths 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: The effect of divestment by asset type

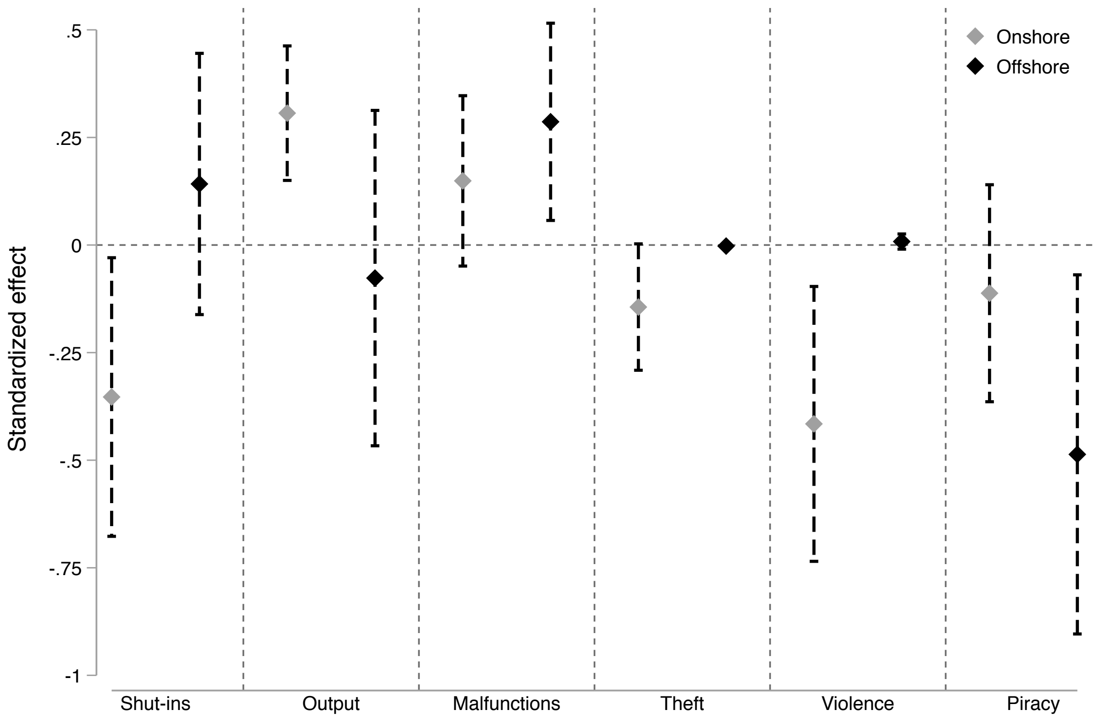

Note: Figure plots the estimates from the difference-in-differences regressions in Tables 3 and 4. Sample is the unbalanced panel of either onshore or offshore oilfields from 2006-2016, as indicated. Outcome variable is given on the categorical axis. Shut-in is defined as a field registering zero output 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. Violence is the total number of oil-related conflict 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.

Figure 5: Robustness plots

(a) Output

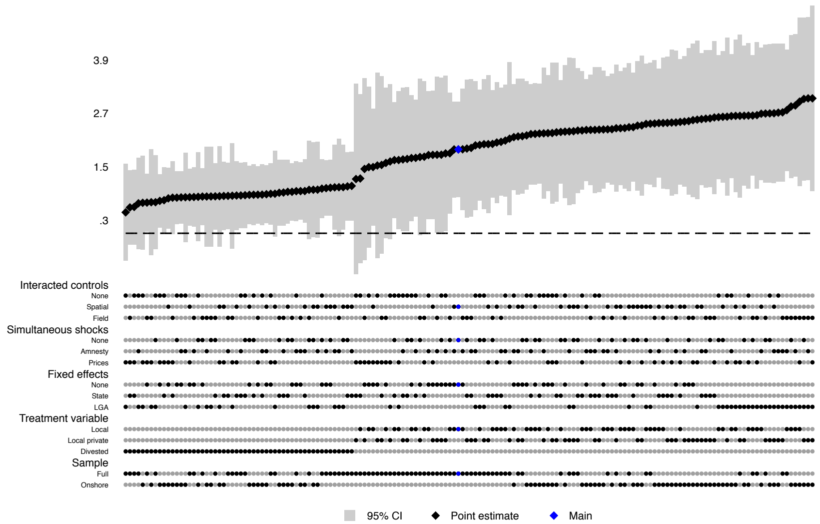

(b) Theft

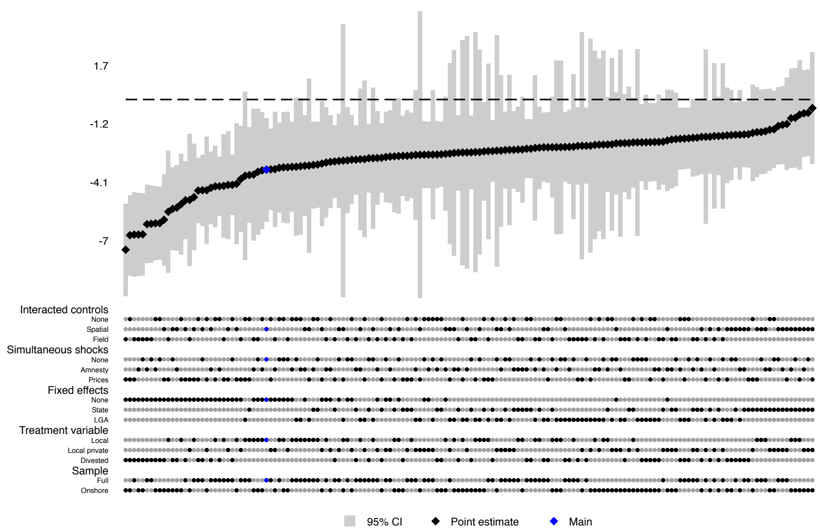

Note: Figure displays estimated coefficients on $local_{it}$ for robustness tests across specifications, for oil theft (Panel A) and oil production (Panel B) outcomes. Specification is indicated by points in the bottom of the figure. "Main" specification is that of Table 1, column (2). Spatial controls 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 B3. Amnesty controls indicate inclusion of controls for the 2009 Niger Delta amnesty. LGA indicate local government area fixed effects. Local private indicates from Table B2, while divested is the treatment indicator from Table A4. Sample is either all or all offshore nonmissing observations for the outcome in question.

Figure 6: Event study plots: enforcement outcomes

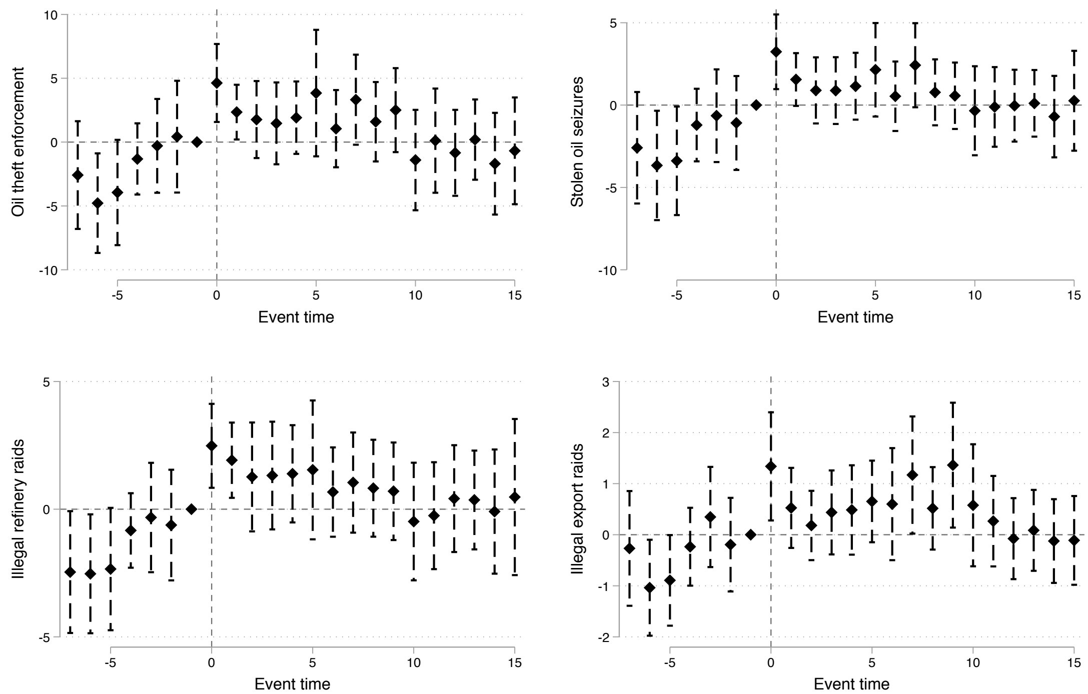

Note: Figure displays coefficients of event-study regressions of outcomes on pre-and-post treatment indicators, conditional on unit and year fixed effect and controls interacted with year dummies. All enforcement outcomes are the total count of enforcement events of each subtype within 20 kilometers of a field centroid. 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: Political connections and theft

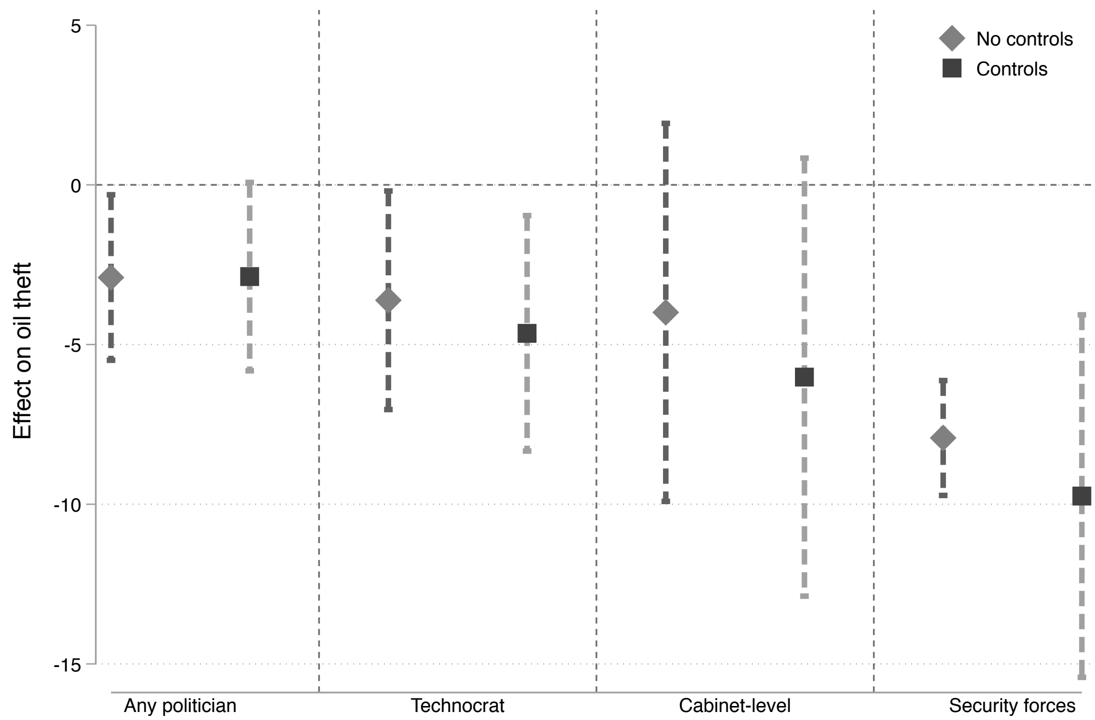

Note: Figure plots the estimates from Table 7. Sample is the panel of 314 oilfields from 2006-2016 for which political connections data is available. Outcome variable is oil theft, the total number of sabotage spills within 15 km of the field. Controls are latitude of the field centroid, distance to 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.

A Supporting exhibits

Table A1: Summary statistics

	Untreated	Treated	Full Sample
Covariates			
Field latitude	4.94	5.28	5.02
	(0.60)	(0.69)	(0.64)
Distance to coast (km)	33.47	31.87	33.10
	(29.72)	(31.05)	(29.99)
Distance to Niger River (km)	81.00	68.65	78.15
	(75.34)	(61.15)	(72.39)
Distance to state capital (km)	87.06	80.52	85.55
	(50.84)	(53.93)	(51.55)
Distance to militant camp (km)	30.36	37.45	32.00
	(24.47)	(32.74)	(26.72)
Number of wells	20.65	15.00	19.34
	(33.12)	(19.98)	(30.65)
Year of first well	1973.77	1975.18	1974.10
	(11.88)	(13.02)	(12.15)
Onshore field	0.69	0.82	0.72
	(0.46)	(0.38)	(0.45)
Max well depth (m)	2694.10	2789.82	2716.24
	(819.28)	(982.54)	(858.97)
Outcomes			
Sabotage events	10.33	5.32	9.33
	(20.85)	(11.74)	(19.47)
Oil-related conflict deaths	0.42	0.50	0.44
	(3.54)	(2.94)	(3.43)
Piracy attacks	0.11	0.13	0.11
	(0.56)	(0.76)	(0.61)
Shut-in field	0.16	0.15	0.16
	(0.36)	(0.36)	(0.36)
Annual oil production (million barrels)	3.44	2.66	3.30
	(7.31)	(4.95)	(6.93)
Operational failure oil spills	7.01	3.30	6.27
	(9.79)	(6.13)	(9.29)
Any politician	0.33	0.43	0.35
	(0.47)	(0.50)	(0.48)
Technocrat	0.33	0.24	0.31
	(0.47)	(0.43)	(0.46)
Elected politician	0.00	0.03	0.01
	(0.00)	(0.17)	(0.08)
Security forces	0.00	0.15	0.03
	(0.00)	(0.35)	(0.17)
Cabinet-level politician	0.14	0.17	0.15
	(0.35)	(0.37)	(0.35)
Number of clusters	244	70	314

Table displays means of variables with standard deviations in parentheses. Sample is a panel of 314 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 1998-2016.

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	173.896***	135.501***	143.360***	203.280***	1.071***	1.141***	1.095***	1.248***
	(42.799)	(39.560)	(49.039)	(42.362)	(0.226)	(0.230)	(0.243)	(0.303)
Treated × Oil price (USD/barrel)			0.434				-0.004	
			(0.977)				(0.004)	
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	2464	2464	2464	2464	1881	1881	1881	1881
R^2	0.835	0.852	0.852	0.876	0.766	0.779	0.779	0.857

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 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 by type of violence

Outcome	All violence		Oil violence		Oil militant		Oil non-militant	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Local operator	-0.799	-1.137	-0.699***	-0.732**	-0.020	0.033	-0.679***	-0.765***
	(0.663)	(0.762)	(0.262)	(0.293)	(0.126)	(0.134)	(0.209)	(0.244)
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	No	Yes	No	Yes	No	Yes
Observations	3497	3497	3497	3497	3497	3497	3497	3497
R^2	0.238	0.314	0.130	0.155	0.173	0.235	0.106	0.122

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 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 A4: The effect of divestment on output and criminality, DI transactions data

	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Output and efficiency					
Outcome	Shut-in		Output		Malfunctions	
MNC-to-local divestment	0.081	0.094	1.003***	1.022***	1.894**	1.424*
	(0.059)	(0.057)	(0.336)	(0.321)	(0.781)	(0.838)
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	2464	2464	2464	2464	3497	3497
R^2	0.655	0.669	0.861	0.878	0.573	0.631
	Panel B: Crime and violence					
Outcome	Theft		Violence		Piracy	
MNC-to-local divestment	-4.239***	-4.504***	-0.270	-0.245	-0.170**	-0.159**
	(0.761)	(0.906)	(0.192)	(0.201)	(0.070)	(0.068)
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	3497	3497	3497	3497	3497	3497
R^2	0.715	0.756	0.129	0.154	0.236	0.315

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. "MNC-to-local divestment" is an indicator that equals one in all periods after a field is sold from a multinational to a Nigerian buyer, as derived from DrillingInfo transactions data. 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. Violence is the total number of oil-related conflict 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: Political connections at the firm-level

	MNC	Local
Any politician	2	26
Any technocrat	2	15
Any cabinet-level politician	1	10
Any elected politician	0	6
Any security forces member	0	7
Number of firms	5	33

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 38 firms for which political connections data on boardmembers, managers, and shareholders is available.

Table A6: The effect of divestment on riots and protests

Outcome	All protests		Oil protests		Riots	
	(1)	(2)	(3)	(4)	(5)	(6)
Local operator	0.345*	0.262	0.144	0.085	0.655	0.589
	(0.181)	(0.191)	(0.133)	(0.144)	(0.449)	(0.448)
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	3497	3497	3497	3497	3497	3497
R^2	0.317	0.366	0.145	0.189	0.445	0.490

Standard errors, in brackets, are clustered at the field level. Outcome variable given in table header, and is defined is the total number of incidents within 15 km of the field. 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 number of wells, initial year, onshore dummy, and maximum well depth. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure A1: Ownership mixes in 2016

(a) HHI

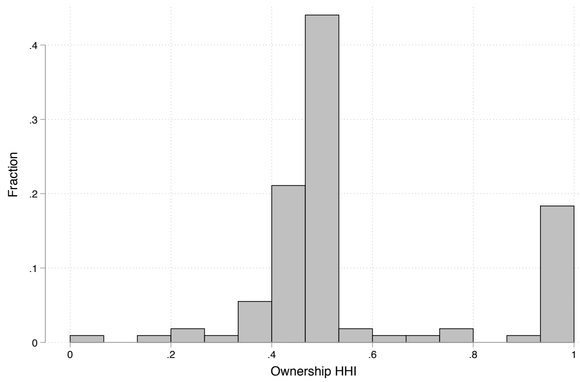

(b) Operator share

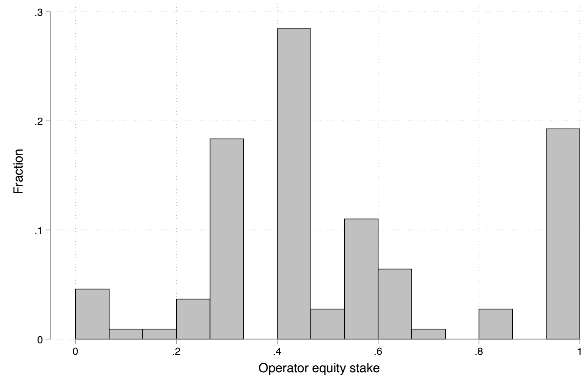

Note: Figure shows histograms of ownership concentration, measured as the Herfindahl index (Panel A), and the stake owned by the operating company (Panel B). Sample is a cross-section of 106 active oil blocks (licenses) in 2016.

Figure A2: Pipeline sabotage and operational malfunction over time

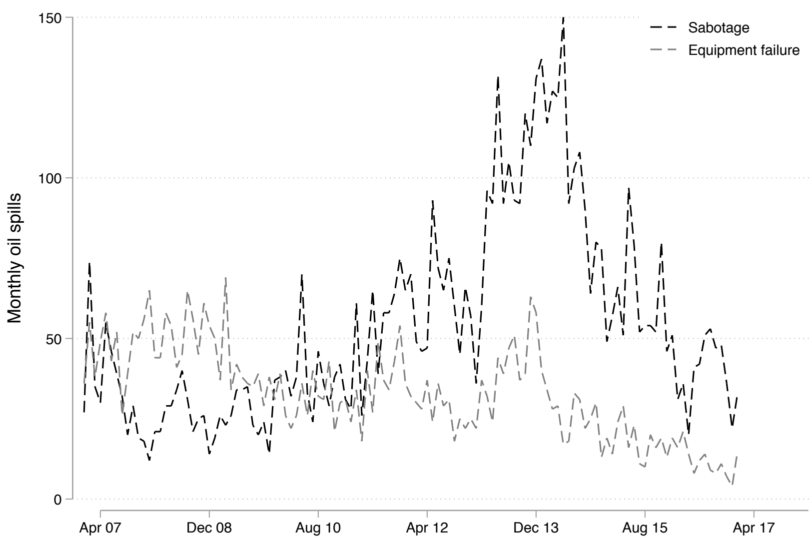

Note: Figure shows monthly totals of oil spills due to sabotage and non-sabotage (equipment failure) over time. Data come from 11,587 oil spills recorded by the NOSDRA OSM from 2006-2017. Vertical lines indicate the beginning of the federal amnesty program for ex-combatants, the end of the initial amnesty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

Figure A3: Indigenization

(a) Fields

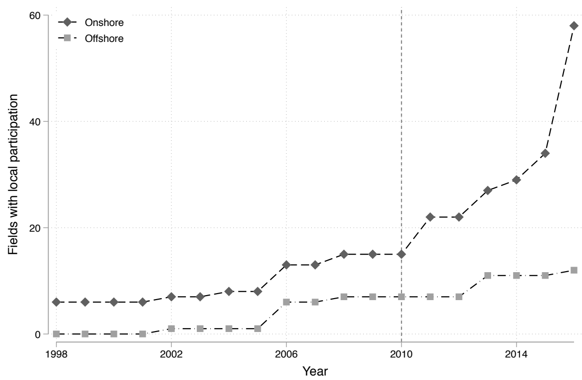

(b) Output shares

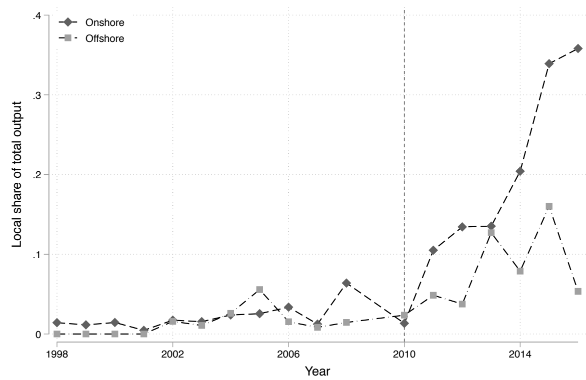

Note: Figure shows number of fields (Panel A) and output share (Panel B) of local Nigerian operators over time by type of asset (onshore vs. offshore). Vertical line indicates the 2010 passage of the Nigerian Local Content Act. Sample is an unbalanced panel of 314 oilfields from 1998-2016. Oil production data and output shares are missing for 2009.

Figure A4: Theft over time

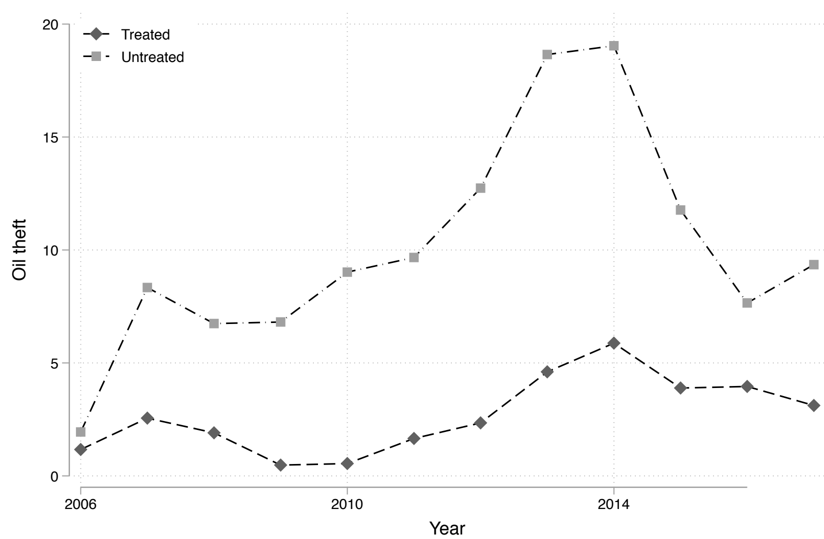

Note: Figure shows mean annual field-level sabotage incidents over time for a sample of 70 ever-treated and 244 never-treated oilfields.

Figure A5: Local ownership and malfunctions, output adjustment

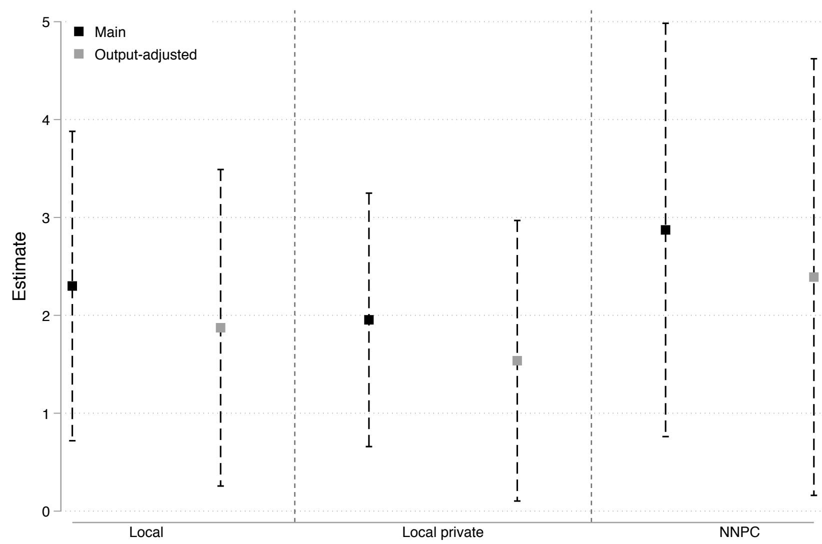

Note: Figure displays coefficients of difference-in-differences regressions of malfunctions on local ownership, local private ownership, and NNPC ownership, as indicated in the categorical axis, conditional on unit and year fixed effects. Output-adjusted estimates are the main estimate, minus the output-malfunctions elasticity times the effect of ownership on output. Malfunctions are the total number of non-sabotage spills within 15 km of the field. Sample is the subset of 180 oilfields that have a fully balanced panel of oil output from 2006-2016.

Figure A6: Local ownership and violence by type of violence

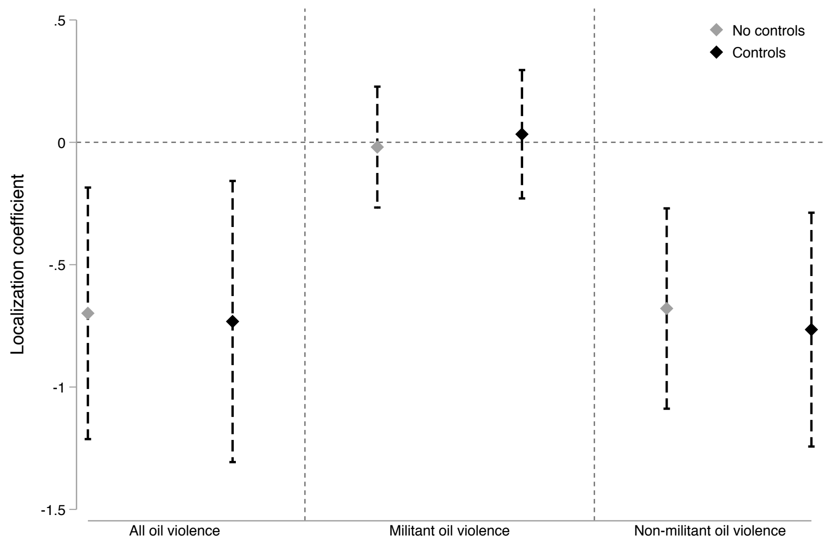

Note: Figure displays coefficients of difference-in-differences regressions of conflict on local ownership as well as time and unit fixed effects. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. Sample is all oilfields from 2006-2017 for which the outcome is non-missing. Militant oil violence is oil-related violence attributed to any organized rebel or militant group by ACLED, while non-militant oil violence is not attributed to any group. All outcome variables are measured in annual number of fatalities.

Figure A7: Gas flaring volumes over time

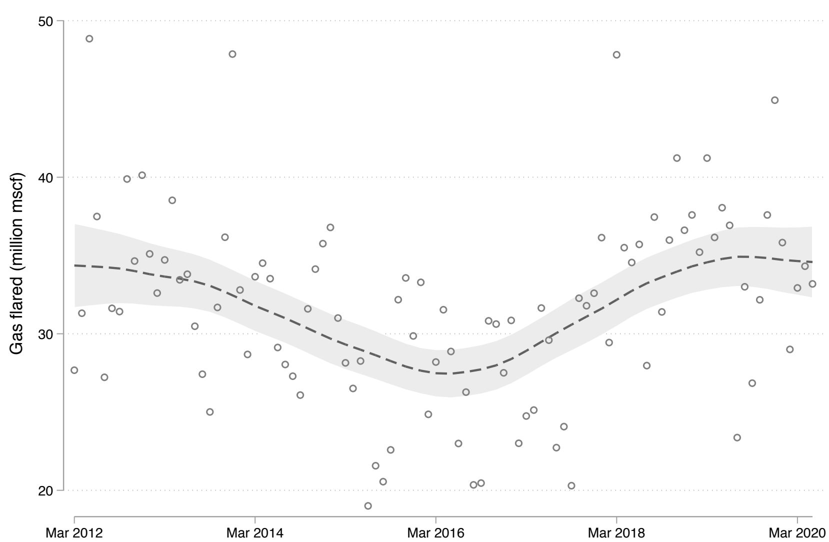

Note: Figure displays total volume of gas flaring on 314 active Niger Delta oilfields over time from March 2012-May 2020.

Figure A8: Gas flaring event-study

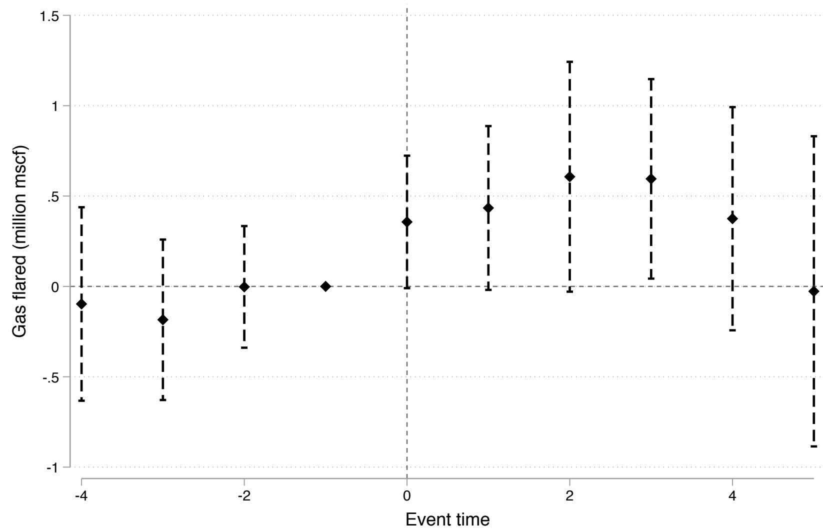

Note: Figure shows coefficients from event-study regressions of outcomes on pre-and-post treatment indicators, conditional on unit and year fixed effects and controls interacted with year fixed effects. Outcome is total volume of flared natural gas in millions of mscf. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. Sample is balanced panel of 314 active Niger Delta oilfields from 2012-2020.

Figure A9: Divestment event-study

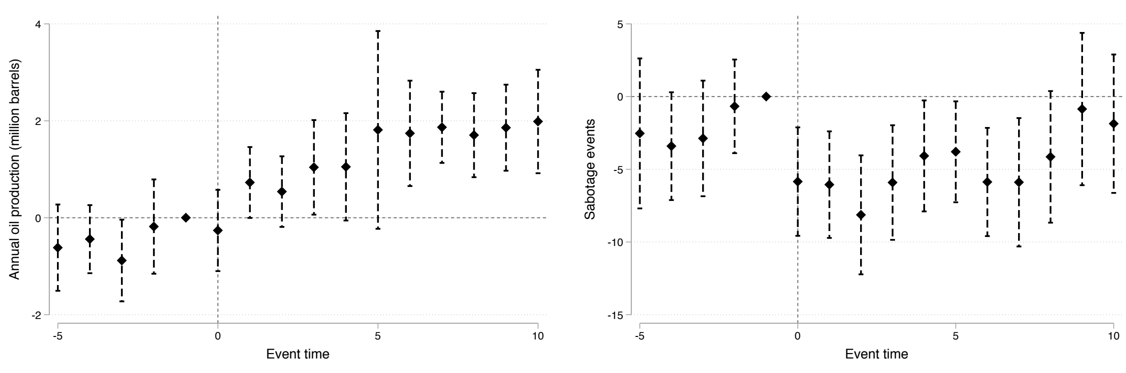

Note: Figure shows coefficients from event-study regressions of outcomes on pre-and-post treatment indicators, conditional on unit and year fixed effects. Treatment is defined as MNC-to-local divestment as indicated in the DrillingInfo corporate transactions data. Output is annual oil production in millions of barrels. Oil theft is the total number of sabotage spills within 15 km of the field. Sample is all nonmissing observations for the outcome in question.

Figure A10: Coefficient estimates for different enforcement outcomes

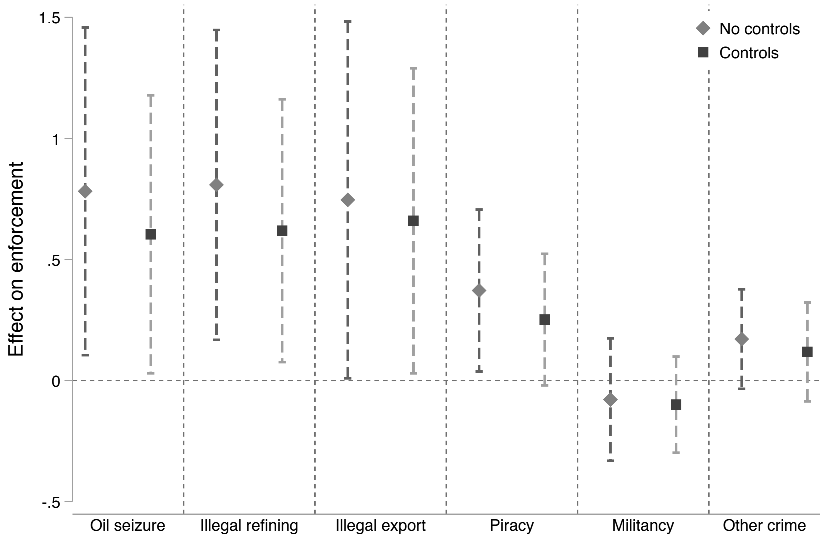

Note: Figure shows coefficient estimates from two-way fixed effects regressions using various different anti0crime enforcement activities as the outcome variable. Outcomes are defined as the count of all enforcement activities of a given subtype – indicated in the x-axis – within 20 kilometers of the field centroid. Standard errors are clustered at the field level. Sample is all active fields from 2006-2017.

Figure A11: The response of output to new discoveries

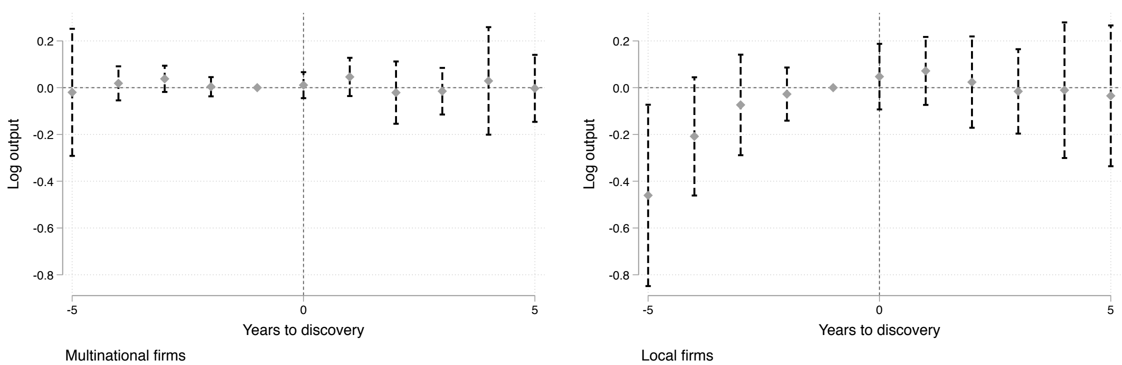

Note: Figure shows coefficients from firm-level stacked-DD event-study regressions of log oil output on log size of oil discovery interacted with dummies for years before and after the discovery date. Event-study models are estimated separately for the sample of multinational and local firms. Each cohort stack includes as treated all firms that experienced a discovery in that year, and includes as controls all firms that did not experience any discovery within 5 years before or after the stack year. All regressions use symmetrical 5 year windows. Standard errors are clustered at the firm-by-event-cohort level.

ONLINE APPENDIX
— For Online Publication Only —

B Data appendix

Oil production and infrastructure data: Information on 314 active Nigerian oilfields forms the core of the data. These field-level data come from Annual Statistical Bulletin of the NNPC, augmented with confidential data from the Department of Petroleum Resources (DPR)[41] 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.[42] 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 source. A "shut-in" field is defined as a field that is nonproducing in a given time period.

There are significant reporting format and content differences between the DPR and NNPC data. DPR data, which is the "official" record, covers a larger number of fields and companies. NNPC reports, in contrast, are provisional, and may aggregate across neighboring fields for smaller operators, or even exclude them entirely. Unfortunately, DPR data only available for four years of the sample: 2006-2008 and 2016, none of which overlap with years in which NNPC data is available. To validate the comparability of the two series', I estimate AR(1) regressions for each pair of consecutive years in the sample. The resulting R^2 and autocorrelation coefficient ρ for these regressions are plotted in Figure B1. Year-to-year correlation is generally high and similar across both data sources, and remains high in year-pairs when the data source changes. Figure B2 plots the log of output in year t against year $t-1$ for years in which the dataset switches from NNPC to DPR (2006 and 2016). These correlations are not noticeably different from those of the previous year.

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 124 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

[41] The DPR is Nigeria's primary petroleum sector regulatory body.
[42] Unfortunately, disaggregated data are unavailable for 2009.

Figure B1: Year-to-year correlations in oil output

(a) ρ (b) R^2

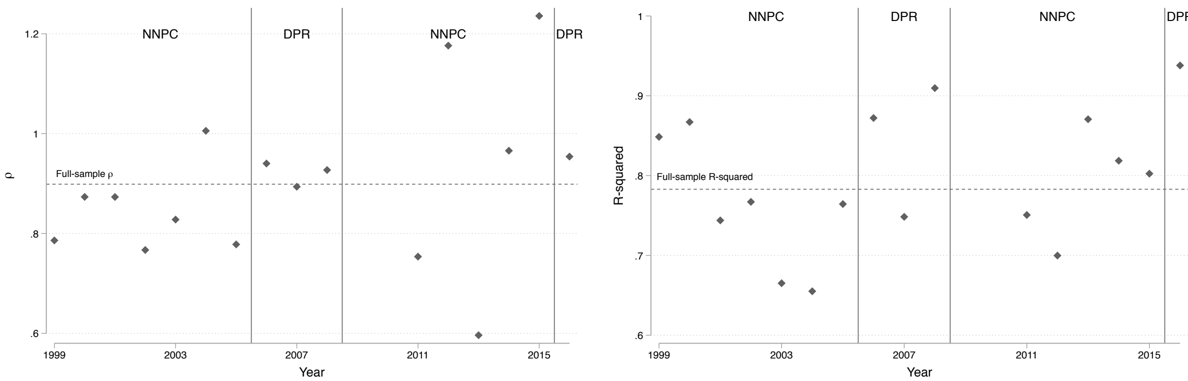

Note: Figure shows coefficient estimates (Panel A) and R^2 (Panel B) from separate AR(1) regressions of oil output for each consecutive year pair in the data. Horizonal line indicates the coefficient or R^2 from an AR(1) regression on the pooled full sample. Vertical lines indicate points at which the data source for oil production changes, with the source indicated in the Figure. Sample is an unbalanced panel of 314 oilfields from 1998-2016. Oil production data are missing for 2009.

result are robust to excluding these observations.[43]

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 block ownership data: Concessions – large blocks of territory, typically containing several oilfields – are the primary unit of ownership in the Nigerian oil market. The exceptions to this rule are 30 "marginal" fields, which are independently-owned fields awarded to local operators that do not belong to larger concession blocks. Concessions are typically jointly owned by several partners, often including an equity stake for the Nigerian National Petroleum Corporation (NNPC).

Detailed data on 113 concessions for the years 2013-2018 comes from the DPR and the Nigerian Extractive Industries Transparency Initiative (NEITI). These sources contain the concession size, location, operator, license type, and detailed equity breakdown. Licenses fall into the following categories: sole risk, joint venture, production sharing, and service contracts. These 113 concessions cover 304 of the 314 fields in the main field-level data, or 97%. From

[43]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.

Figure B2: Oil output, selected consecutive years

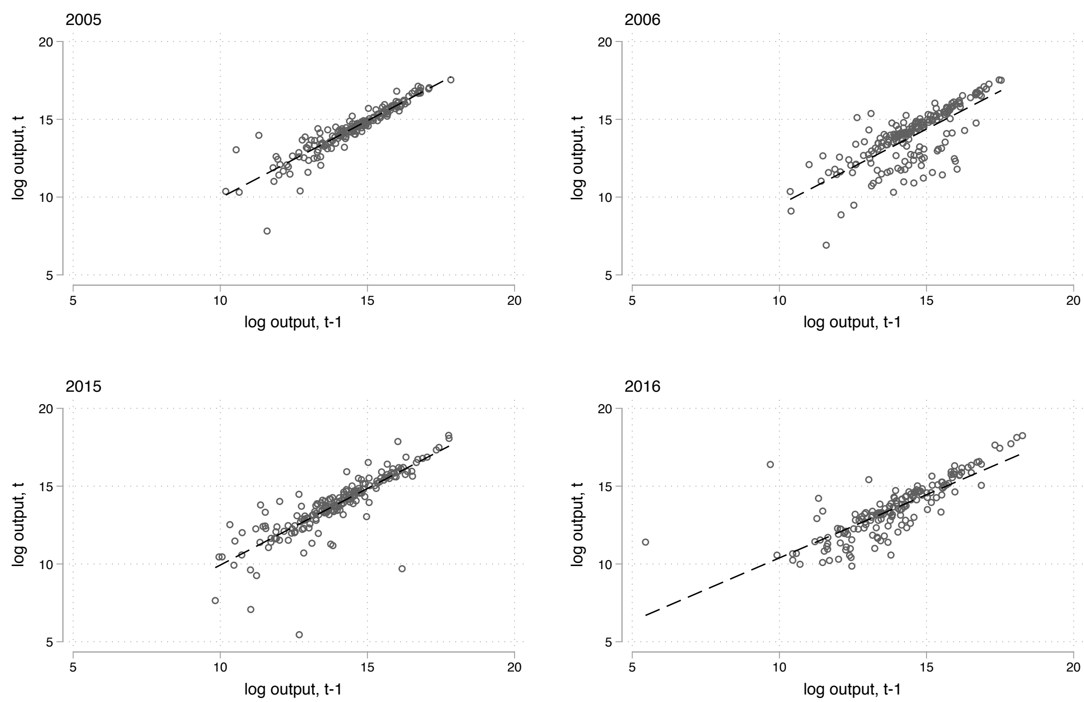

Note: Figure shows field-level scatterplots and linear fits of the log of output in t against $t-1$ for selected consecutive year-pairs in which the data source for oil production changes. Sample is an unbalanced panel of 314 oilfields from 1998-2016. Oil production data are missing for 2009.

this data I obtain the ownership shares of all partners for all active oil mining leases, as well as the operating firm. Ownership data is only available from 2012-2018. I therefore exclude it from the main analysis and use it only to test mechanisms.

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.[44] 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.

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 oil-related violence, I use all conflict events that contain the following oil-industry-related strings: petroleum, petro, Agip, Shell, Eni, drilling, rig, well, pipeline, ndv, flow, NNPC, NPDC, exxon, mobil, total, addax, or gas. This captures attacks on the oil sector perpetrated by any armed groups. I then further distinguish between onflict events perpetrated by organized rebel or political militia groups, which I call "militant" attacks, and those perpetrated by unknown or unorganized groups, which I call "non-militant" attacks. 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

[44] Rexer and Hvinden (2020) for a discussion about measuring oil theft.

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.[45] 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.[46] 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 incompletion 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 contains any personnel ever satisfying some criterion.

Data on militant groups: Finally, I use data on militant camps, described in detail in Rexer and Hvinden (2020). 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 Rexer and Hvinden (2020)), 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 construct 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

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

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

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 Rexer and Hvinden (2020) which identifies the strongest camps based on the number of local allies along the pipeline network. The logic behind this measure is that groups with a greater number of local allies along their pipeline are better able to coordinate and carry out large scale infrastructure attacks because of strategic complementarities. Figure B3, reprinted from Rexer and Hvinden (2020), plots the percent change in output between 2005-2009 against the number of allies within 10 kilometers along the pipeline at the camp-level, conditional on state fixed-effects and camp-level controls for slope, altitude, average temperature, average precipitation, latitude, and distance to the nearest pipeline, state capital, and Atlantic coast. I choose 2005-2009 as this corresponds to the period of greatest violence in the Niger Delta conflict. The plot shows a robust negative correlation – camps with more local allies see substantially larger declines in oil output during the height of the conflict.

DrillingInfo corporate transactions data: Data on corporate transactions comes from DrillingInfo (DI), a paid-subscription database on the oil and gas sector. From DI I obtain a list of 171 corporate transactions in the Nigerian oil and gas sector from 2006-2020. I then download and digitize corresponding PDF files for each transaction which contain, among other information, the announcement and closing dates, name of buyers, sellers, and assets, deal value, deal status at the time of reporting (closed, terminated, or in progress), and the typeof transaction (corporate M&A, new discoveries not yet developed, exploration blocks previously awarded, fields under development, producing fields, and new awards). From this dataset I extract an asset-level cross-section containing, for each asset, the nationality of buyers, sellers, and the transaction date, for all transactions concerning that asset. I define the transaction date as the closing date where available; if unavailable, I use the announcement date. Many transactions contain information on both fields and block, since the former is typically, though not always, contained in the latter. If field-level information is available, I use that, since some fields within a block may be divested while others are not; otherwise I take the block-level information. Following this procedure, the 171 DI transactions produce a datasets of 126 individual fields and 69 distinct blocks. I then merge these data to the main field-level dataset. In total, 19 blocks show up in the block-level merge, and 50 fields in the field-level merge. Since only 21% of transactions cover assets that are actively producing at the time of the transaction, these match rates are reasonable.

I define several variables from the DI data. Firstly, I define a MNC-to-local "divestment" indicator which equals one for a given field in all years after that field experienced any transaction in which any buyer was Nigerian and any seller was multinational. Because this measure incorporates transactions unobserved in the DPR/NNPC administrative data[47] it is not

[47] For example, cases when a local firm acquires a non-operating stake in a given asset.

Figure B3: Local alliance density and damage

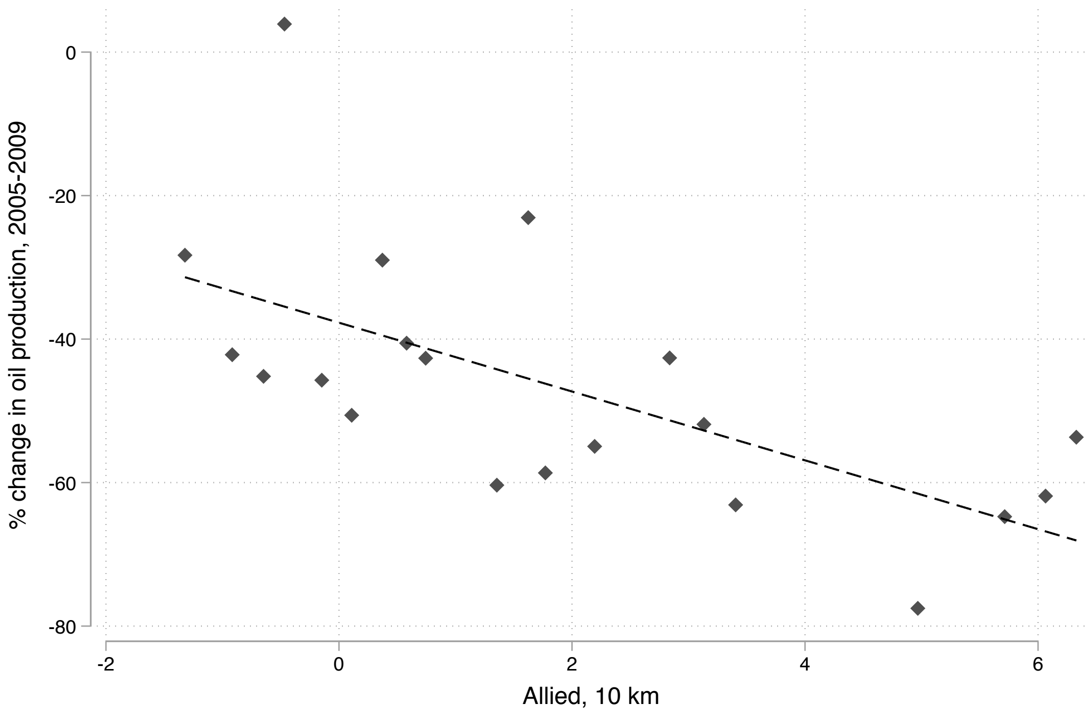

Note: Figure is reprinted from Rexer and Hvinden (2020). This figure shows a binned scatterplot of the militant-camp-level relationship between the number of allied connections along the pipeline and damage inflicted during the height of the Niger Delta Crisis. Damage is measured as the % change in onshore oil production within 20 km of the militant camp location between 2005 and 2009. with the number of allied camps within 10 km along the pipeline as the independent variable. Correlations are conditional on state fixed-effects and camp-level controls for slope, altitude, average temperature, average precipitation, latitude, and distance to the nearest pipeline, state capital, and Atlantic coast.

perfectly correlated with the treatment indicator defined by operatorship; the within-R^2 of a regression of one treatment measure on the other, conditional on field and year fixed effects, is 0.16. I define similarly variables that measure a field's exposure to local-to-local and MNC-to-MNC transactions. 58, 27, and 23 fields are ever-exposed to MNC-to-local, local-to-local, and MNC-to-MNC transactions, respectively, during the sample period. Finally, I define an indicator of terminated or delayed divestments which equals one for all years after the announcement of an MNC-to-local divestment but before it's consummation. In some cases, these are terminated/nullified transactions, while in others, this reflects a delay between the announcing and closing dates. This indicator equals zero if and when the field eventually becomes "treated" according to the divestment measure. 43 fields are exposed to a delayed or terminated divestment in the sample period.

Gas flaring data: Data on gas flaring volumes comes from the Nigeria Gas Flare Tracker,[48] a joint project by NOSDRA and the NGO Stakeholder Democracy Network. I download monthly panel data on total gas flaring volume from March 2012 to May 2020, measured in thousands of cubic feet (mscf), for 210 flare sites. I then georeference these sites manually by cross-referencing the map interface of the Gas Flare Tracker against a Google maps layer containing Nigeria's oil and gas infrastructure. I then match flares to fields using a spatial merge process. 119 flare sites fall directly within the boundaries of an identifiable field. A further 73 are matched to their nearest field within 10 kilometers. The remaining 18 flare sites either fall on the Cameroonian side of the maritime border ($n = 9$), are far from the Niger Delta ($n = 2$), or are not near any identifiable field ($n = 7$). In total, these 192 final flare sites cover 143 fields. Lastly, I merge to the production data; 180 out of 192 flare sites occur in fields actually actually contained in the DPR/NNPC output data. These matched fields account for 93.4% of the flared gas volume over the period.

Law enforcement activity: Data on law enforcement activity comes from the text of Nigerian news media reports. We begin by assembling a comprehensive collection of plausibly relevant news articles covering topics of oil theft, law enforcement, and crime in Nigeria by searching relevant keywords in the Dow Jones Factiva media database. We collect all articles that satisfy each of the following criteria: i) mention the word "oil", ii) mention at least one of a set of enforcement-related keywords[49] , iii) mention at least one of a set of exact oil crime-related phrases.[50] Some examples of relevant articles are shown in Figure B4.

[48] https://nosdra.gasflaretracker.ng/

[49] These are: "raid, raids, raided, seize, seized, seizes, seizure, seizures, destroy, destroys, destroyed, operation, capture, captures, captured, arrest, arrested, arrests, kill, killed, kills, apprehend, apprehended, apprehends, burn, burns, burned, invade, invaded, invades, search, searches, search"

[50] These are: artisanal refineries, artisanal refinery, artisanal refining, bunkerers, bunkering camp, bunkering gang, bunkering site, ex toru, illegal bunkering, illegal diesel, illegal fuel, illegal oil, illegal refineries, illegal refinery, illegal refining, illegally refined, joint task force, Nigerian military, Nigerian Navy, oil bunkerers, oil bunkering, oil smugglers, oil theft, oil thief, oil thieves, oil vandals, operation 777, operation awase, operation crocodile smile, operation delta safe, operation eagle eye, operation pulo shield, operation python dance, operation restore hope, operation river sweep, operation safety check, operation tsare teku, pipeline sabotage, pipeline vandal,

Figure B4: Relevant articles

(a)

(b)

(c)

Note: This figure shows screenshots from relevant articles in *The Vanguard*, *The Sun*, and *The Guardian Nigeria*, all local Nigerian newspapers.

This procedure yields 17146 total articles potentially related to oil theft enforcement.[51] We then hired Nigerian research assistants to first identify all articles that are relevant to law enforcement activity in Nigeria, yielding a total of 3932.[52] From this set of relevant articles we then manually extract all *law enforcement events*, where an event is defined as a unique interaction between law enforcement and suspected criminals that occurs in a specific location. For each event, we code the following variables: *i*) the location of the event, typically a neighborhood, village, oil asset, or local government area (municipality) *ii*) the law enforcement agency, *iii*) the illegal activity committed, selected from a pre-coded list,[53] *iv*) the items seized or destroyed in the law enforcement action, selected from a pre-coded list,[54] *v*) the total number of arrests, and *vi*) the total number of fatalities. Extensive manual quality checks were conducted on weekly researcher submissions.

We consider any two articles as duplicates if they are published in the same calendar week and their headlines exceed a string similarity score threshold as defined by Levenshtein edit distance. After grouping duplicates into unique articles, we take the union of all events identified by the researchers to allow that duplicate articles may contain both repeated events as well as independent information.[55] In total, we obtain 5682 law enforcement events for which the location can be reliably geocoded, of which 3261 are related to oil theft. These events cover 3379 unique articles. 89% of all locations mentioned in relevant events were successfully geocoded. We then merge these enforcement events to villages in our sample using 5 kilometer rings, the same criteria used for oil theft. Figure B5 plots quarterly total enforcement actions for anti-militancy and anti-oil theft actions separately.

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.[56] 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.[57]

---

pipeline vandalism, pipeline vandals, pirate, pirates, stolen crude, stolen, diesel, stolen oil, swamp buggy.

[51] Of course, these search terms are unlikely to be exhaustive, but they were derived from substantial reading of these articles. Also, note that this figure may be inflated because the same story is sometimes published by multiple different media outlets.

[52] We excluded articles about unrelated conflicts such as Boko Haram in Northern Nigeria, but included articles about non-oil illegal activities such as armed robbery, gang activity, and fraud

[53] These are: oil theft, piracy, illegal refining, pipeline vandalism, transportation of stolen oil, kidnapping, cultism/gang activity, militancy, and other illegal activities.

[54] These are: no item, boats, stolen oil, arms, illegal refineries, trucks, oil theft equipment, and other items.

[55] For example, if there are two articles about the same raid, one may mention a second event, while the other does not.

[56] I do not observe the cause of missingness. I therefore assume this data is missing at random. Table B1 shows that outcomes and covariates are very similar across these samples, supporting this assumption.

[57] I choose not to restrict the sample for all estimation in order to make full use of available data for non-

Table B1: Summary statistics by estimation sample

	Output (1)	Full (2)
Sabotage events	9.77	9.33
	(20.87)	(19.47)
Operational failure oil spills	7.09	6.27
	(10.10)	(9.29)
Piracy attacks	0.15	0.14
	(0.78)	(0.70)
Oil-related conflict deaths	0.26	0.37
	(1.92)	(3.09)
Local operator	0.11	0.12
	(0.31)	(0.32)
Field latitude	5.01	5.01
	(0.63)	(0.62)
Distance to coast (km)	33.76	32.22
	(29.93)	(29.10)
Distance to Niger River (km)	76.93	78.98
	(73.78)	(72.44)
Distance to state capital (km)	87.42	84.27
	(50.95)	(50.72)
Number of observations	2476	3497

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 314 fields between 2006-2016.

Figure B5: Enforcement

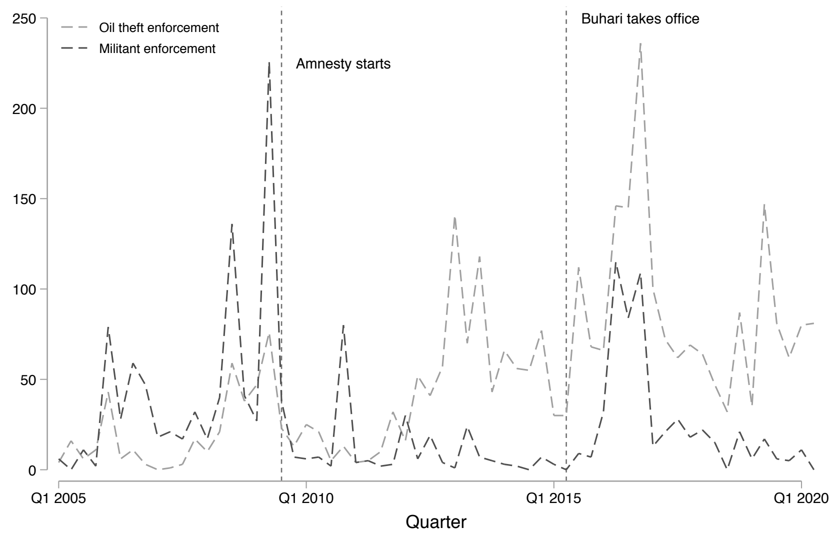

Note: This figure shows total quarterly militancy and oil theft-related law enforcement actions. Important dates in the amnesty process are indicated by labeled vertical lines.

production outcomes.

C Model

C.1 Set-up

The interaction is a simple sequential, one-shot game between firms, indexed by $f \in F$, gangs, indexed by $g \in G$, and state security agents, which are homogenous. The firm produces a fixed level of surplus $\bar{Q}$, sold at the international oil price p. The game proceeds as follows. Firms and gangsters simultaneously offer bribes b_f and b_g to law enforcement. Law enforcement observes these bribe offers and decides which to accept and which to reject. If law enforcement accepts the gangster's bribe, oil theft is allowed, and enforcement $e = 0$. The gangster steals a constant quantity $q < \bar{Q}$ at fixed cost $c - \epsilon_g$, where ϵ_g is private information. Theft is inefficient both because gangsters incur costs that firms don't, $pq - c + \epsilon_g < pq$ and because it directly destroys output, denoted by $\kappa > 0$. If law enforcement instead accepts the bribe b_f, then $e = 1$ and they must enforce the law. The enforcement technology reduces the probability of theft to $\alpha < 1$ at cost $c(e)$, where $c(1) = \eta > c(0) = 0$. Furthermore, since most illegal activities along the black-market value chain depend on the actual procurement of stolen oil, then under enforcement the cost of theft is reduced by a factor of α as well. All players are price takers at world oil price p.

Firms may differ in a number of 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 or reputational costs of corrupt payments. In addition, firms only receive a share γ_f of Q, to capture the important role of joint-ventures in Nigeria, as shown in Figure A1. Importantly, law enforcement may internalize firm f's output based on the parameter μ_f. This measures how the strength of political connections determines enforcement behavior. If a firm is unconnected, then $\mu_f = 0$. Note that $\gamma_f + \mu_f \leq 1$

C.2 Base case

The payoffs under $e = 0$ and $e = 1$ are as follows:

$$
\begin{aligned}
U_f^0 &= \gamma_f p(\bar{Q} - q - \kappa) & \qquad U_f^1 &= \gamma_f p\bar{Q} - \alpha\gamma_f p(q + \kappa) - \lambda_f - b_f \\
U_g^0 &= pq - c + \epsilon_g - b_g & \qquad U_f^1 &= \alpha(pq - c + \epsilon_g) \\
U_s^0 &= b_g + \mu p(\bar{Q} - q - \kappa) & \qquad U_s^1 &= b_f + \mu p\bar{Q} - \alpha\mu p(q + \kappa) - \eta
\end{aligned}
$$

Definition 1. ***Efficient corruption.*** *Law enforcement corruption is efficient when the equilibrium bargaining outcome maximizes total surplus.*

Total surplus $S(e)$ is defined as $S(1) = p(\bar{Q} - \alpha k) - \alpha c - \eta$ and $S(0) = p(\bar{Q} - k) - c$.[58] If $S(1) > S(0)$, then $e = 1$ is efficient corruption, else $e = 0$. This leads to our first assumption:

Assumption 1. ***Enforcement cost.*** *Enforcement costs are sufficiently low, enforcement sufficiently productive, or crime is sufficiently wasteful, that stopping crime is always socially optimal, so $S(1) > S(0)$. This yields*

$$\eta \leq (1 - \alpha)(pk + c)$$

Of course, if enforcement is costless this condition is always satisfied. I restrict focus to cases when crime deterrence is efficient in order to rule out pathological cases where crime is socially valuable. This assumption does not affect any of the subsequent analysis, but it does affect the welfare implications of different equilibrium outcomes.

Definition 2. ***Bargaining Range.*** *The bargaining range B is the set of firm bribes b_f for which enforcement can be sustained in equilibrium, defined as the interval $[\bar{b}_g, \bar{b}_f]$.*

$\bar{b}$ are the reservation points of gangster and firm. If $b_f < \bar{b}_g$, then the gangster is willing to pay more than the firm offers, and crime occurs with probability one. Similarly b_f must be individually rational and therefore cannot exceed $\bar{b}_f$. Using the utilities for f and g yields the reservation points

$$\bar{b}_g = (1 - \alpha)(pq - c + \epsilon_g) \qquad \bar{b}_f = (1 - \alpha)\gamma_f p(q + k) - \lambda_f$$

Note that government rents stem directly from their partial monopoly of violence. When enforcement is ineffective, $\alpha = 1$, neither party has any incentive to bribe the security forces.

The government prefers to enforce whenever $U_s^0 < U_s^1$. This yields the reservation point

$$b_g + \mu(\bar{Q} - q - k) = b_f + \mu\bar{Q} - \alpha\mu(q + k) - \eta$$

Definition 3. ***Bribe offers.*** *Assume that law enforcement extracts all of the surplus from gangsters, so that $b_g = \bar{b}_g$.*[59] *Then the threshold bribe for which government enforces is given by:*

$$b^* = (1 - \alpha)(pq - c + \epsilon_g) + (\alpha - 1)\mu p(q + \kappa) + \eta$$

This expression gives us our first key prediction: since $\alpha - 1 < 0$ an increase connections μ reduces the bribe required for security agents to enforce the law. Note also that setting $\mu = \alpha = \eta = 0$ reflects the situation where firm and gangster bargain directly with each other

[58]Note that I consider enforcement and theft costs as surplus destruction, because these are real resource costs. However, labor costs in both of these are plausibly transfers to local communities. As such, we can think of this more narrowly as the surplus from oil production that can be appropriated by the agents in our model.

[59]Note that this is without loss of generality. We could allow some fraction of the surplus to be retained by gangsters, in which case we would simply have another fractional parameter to carry around.

and gangs receive a take-it-or-leave-it offer. μ introduces a friction in favor of firms, while costs of enforcement η and its incomplete nature α introduce wedges in favor of theft.

Assumption 2. ***Information structure.*** *Assume that the firm does not observe ϵ_g until the bargaining phase, so it is stochastic in the output choice stage. Assume ϵ_g is distributed uniformly on the interval $[0, c]$.*

Efficient corruption occurs whenever $b^* < \bar{b}_f$. Using the uniform distribution of ϵ_g, the probability of enforcement is

$$\begin{aligned} Pr(e = 1) &= Pr(b^* < \bar{b}_f) \\ &= \frac{1}{c}[(\gamma + \mu - 1)pq + (\gamma + \mu)p\kappa] - \frac{\lambda_f + \eta}{(1 - \alpha)c} + 1 \end{aligned}$$

Proposition 1. ***Comparative statics: enforcement and theft.*** *Given $\gamma_f + \mu_f \leq 1$ and A2, the likelihood of enforcement is decreasing in η, λ, q, α, and increasing in μ, γ, κ. Enforcement is increasing in p whenever $\frac{\kappa}{q} > \frac{1}{(\gamma+\mu)} - 1$. Since theft is just $\alpha Pr(e = 1) + (1 - Pr(e = 1))$, it has the same predictions in the opposite direction.*

Proof:

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

Note the condition for the comparative static on prices. $\frac{\partial Pr(e=1)}{\partial p} > 0$ whenever

$$(\gamma + \mu - 1)q > -(\gamma + \mu)\kappa$$

$$\frac{\kappa}{q} > \frac{(1 - \gamma - \mu)}{(\gamma + \mu)}$$

$$\frac{\kappa}{q} > \frac{1}{(\gamma+\mu)} - 1$$

□

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

C.3 No commitment

The environment in the base case makes an important implicit assumption: that all contracts can be perfectly enforced. In violent, anarchic environments like the Niger Delta, a no-commitment assumption is more plausible. As before, I hold the behavior of the gangs fixed and assume that government extracts all surplus from that interaction, in order to focus on the interaction between government and firm. Assume further that $\epsilon_g = 0$ for all g, so WLOG we focus only on a single equilibrium instead of a continuum of equilibria. Now, the security agent has a third action available: accept a bribe from both parties and renege on the agreement with the firm.[60] In this case, political connections become a binding constraint to enforcement.

For illustration, consider the case when $\mu = 0$. Then $b_g + b_f > b_f - \eta$ and $b_g + b_f > b_g$. So accepting both bribes and allowing theft is a dominant strategy for the government, for any bribe levels. As such, the firm's will always to obtain payoff U_f^0 for any bribe. Therefore, setting $b_f > 0$ and incurring the cost of corruption λ can never be optimal for the firm, since $U_f^0 - \lambda_f - b_f < U_f^0$. Therefore, without commitment mechanisms, bribes are ineffective and political connections are a necessary condition to sustain enforcement. This leads us to a more general proposition.

Proposition 2. ***Enforcement without commitment***. *Assume a no-commitment environment, i.e., that law enforcement maintains the option to renege on a deal with the firm, and assume that the behavior of the gangster is fixed at $b_g = \bar{b}_g$. Then there are two possible outcomes of the stage game, each a unique Nash equilibrium. Let $\bar{\mu} = \frac{(1-\alpha)(pq-c)+\eta}{(1-\alpha)p(q+\kappa)}$. When $\mu \geq \bar{\mu}$, the government accepts any firm bribe offer and sets $e = 1$, and the firm sets $b_f = 0$. When $\mu < \bar{\mu}$, the government accepts both firm and rebel bribe offers and sets $e = 0$, and the firm sets $b_f = 0$.*

Proof: When the firm is politically connected, the incentives can align for sufficiently large μ. In particular, for $e = 1$ to be a dominant strategy, the payoff to the security forces

[60] I ignore the deviation where government accepts a bribe from both but still enforces the law, since the government's is always willing to enforce in this case if it is willing to enforce without the additional bribe from the gang. In our setting, this omission seems well since gangs can easily punish the police with violence to enforce commitment in that interaction.

from following the agreement must exceed that of reneging and accepting both bribes:

$$U_s^1 \geq U_s^0 + b_f > U_s^0$$

Which yields the condition

$$\mu \geq \frac{(1-\alpha)(pq-c)+\eta}{(1-\alpha)p(q+\kappa)} = \bar{\mu}$$

When this condition is met, the government has a dominant strategy. For any $b_f \geq 0$, accepting the bribe and enforcing is a best response, since μ is such that that the government sufficiently internalizes theft losses. Knowing this, the firm will set $b_f = 0$ to maximize its payoff. Therefore, the Nash equilibrium is unique.

Clearly, when $\mu < \bar{\mu}$ then we have $U_s^1 < U_s^0 + b_f$ and of course $U_s^0 + b_f \geq U_s^0$. So $e = 0$ is a dominant strategy for the government for any b_f. Again, the firm must set $b_f = 0$ because $U_f^0 - \lambda_f - b_f < U_f^0$. Finally, note that the profitability of theft is a sufficient condition for political connections to be a binding constraint on enforcement. Then $pq - c > 0$ and so $\bar{\mu} > 0$. □

C.4 Dynamic bargaining

Of course, the assumption of a one-shot game without commitment makes sustaining cooperation very difficult, and may be too extreme. In reality, law enforcement and firms interact in a repeated setting. Consider the game with no commitment, repeated infinitely. Let the players have a common discount factor δ and for simplicity let $\alpha = 0$, so enforcement is perfect. Then efficient corruption may occur even when $\mu < \bar{\mu}$.

Proposition 3. ***Dynamic enforcement***. *Let $\mu < \bar{\mu}$. Then for sufficiently large δ, law enforcement provision can be sustained in a subgame perfect equilibrium of the infinitely repeated game where government cannot commit in the stage game.*

Proof: First note that when $\mu \geq \bar{\mu}$, enforcement can be trivially sustained in subgame perfect equilibrium by playing the Nash equilibrium of the stage game in every period. When $\mu < \bar{\mu}$, enforcement is no longer a Nash equilibrium of the stage game. Nevertheless, it can be restored with a simple trigger strategy profile: the firm begins by offering $b_f = b^*$ and continues to do so in every period until the cooperative outcome is not played, after which the firm sets $b_f = 0$ forever. The government accepts all bribes $b_f \geq b^*$ and responds with $e = 1$. After any period in which the cooperative outcome is not played, government sets $e = 0$ forever.

First note that since $\mu < \bar{\mu}$, the punishment is the stage game Nash and so is subgame perfect after a deviation. The value to the security forces of playing the punishment equilibrium is:

$$r_s = \sum_{t=0}^{\infty} (pq - c + \mu p(\bar{Q} - q - \kappa))^{\delta} = \frac{pq - c + \mu p(\bar{Q} - q - \kappa)}{1 - \delta}$$

Given a bribe b_f, security forces are willing to enforce the law rather than deviate and allow theft whenever:

$$b_f + (1 - \delta) r_s + \delta r_s \leq \frac{b_f + \mu p \bar{Q} - \eta}{1 - \delta}$$

Solving for b_f gives us the minimal bribe that the government is willing to accept for the equilibrium to be sustained.

$$b^* = \frac{1}{\delta}(pq - c - \mu p(q + \kappa) + \eta)$$

Note that this is similar to the minimum bribe in the base case. However, in the dynamic game, the minimum per-period rent transferred to the state must be inflated by a factor of $\frac{1}{\delta}$ relative to the minimal transfer in the one shot game with commitment, since now it must be enforced with dynamic incentives. The firm's value of punishment:

$$r_f = \sum_{t=0}^{\infty} (\gamma_f p(\bar{Q} - q - \kappa))^{\delta} = \frac{\gamma_f p(\bar{Q} - q - \kappa)}{1 - \delta}$$

The firm must be willing to set $b_f > b^*$ rather than set $b_f = 0$ and induce punishment. So the firm's incentive condition is

$$r_f \leq \frac{\gamma_f p \bar{Q} - \lambda_f - b_f}{1 - \delta}$$

Yielding the same maximal willingness to pay as the base case:

$$b_f = \bar{b}_f = \gamma_f p(q + \kappa) - \lambda_f$$

Importantly, note that this condition is identical because the firm has no commitment problem given the sequential structure of the stage game. To illustrate why this matters, assume briefly that the firm can deviate in the stage game and enjoy a single period of bribe-free enforcement. Then the incentive condition becomes:

$$\gamma_f p \bar{Q} + \delta r_f \leq \frac{\gamma_f p \bar{Q} - \lambda_f - b_f}{1 - \delta}$$

Yielding a maximal willingness to pay:

$$\bar{b}_f = \delta \gamma_f p(q+\kappa) - \lambda_f$$

As before, efficient corruption occurs whenever $\bar{b}_f \geq b^*$. This implies the following condition:

$$\delta \geq \bar{\delta} = \frac{(pq - c - \mu p(q+\kappa) + \eta)}{\gamma_f p(q+\kappa) - \lambda_f}$$

Note that $\mu < \bar{\mu}$ implies that $\bar{\delta} > 0$, so the incentive constraint binds. □

Now we can slightly revise the predictions of Proposition 1 to say that the enforcement equilibrium becomes *more likely* and theft becomes *less likely* as $\bar{\delta}$ falls.

Proposition 4. ***Comparative statics: dynamic enforcement***. *Let $\mu < \bar{\mu}$. Say that the likelihood of enforcement is decreasing in $\bar{\delta}$. Then the comparative statics from Proposition 1 all hold in the no-commitment dynamic bargaining game.*

The proof is immediate, since $\delta > \bar{\delta} \iff \bar{b}_f > b^*$. Similarly, Proposition 1 relies on the condition that $\bar{b}_f > \delta b^*$ and $\delta > 0$. □

D Difference-in-differences robustness tests

D.1 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 B2 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, violence, 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.

D.2 Oil prices

I also test robustness of the main results to differential oil price effects in Table B3. To do this, I include the interaction between the time-invariant localization treatment indicator and the time-varying oil price series p_t. I find no evidence that differential responses to oil price changes by localized fields are driving the results.

D.3 Amnesty policy

Rexer and Hvinden (2020) show that the 2009 amnesty for Niger Delta militants reduced violence and increased oil theft differentially in amnestied regions. If multinationals divested of onshore oilfields in militant-controlled areas during and after the conflict period, then it may be the case that the amnesty policy is contaminating our estimate of the effect of localization on violence and theft. I test robustness to this concern in Table B4 by including the interaction between indicators for post-amnesty and amnestied area[61] in the main TWFE model. The results are unaffected. Consistent with Rexer and Hvinden (2020), the coefficient on the amnesty interaction term is positive and significant for theft, and negative and significant for piracy and violence.

D.4 Field-level covariates

A key threat to identification is that there may be selection into field takeover based on field characteristics. Table A1 demonstrates that localized fields are younger, smaller, and

[61] This variable equals one if the field is within 30 km of an amnestied militant camp.

Table B2: The effect of divestment on output and criminality, public and private

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Output and efficiency						
Outcome	Shut-in		Output		Malfunctions	
Private local operator	-0.277***		2.181***		-0.216	
	(0.075)		(0.659)		(1.193)	
Government operated		-0.005		2.060**		3.248***
		(0.077)		(0.838)		(0.927)
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	2464	2464	2464	2464	3497	3497
R^2	0.673	0.668	0.878	0.878	0.631	0.632
Panel B: Crime and violence						
Outcome	Theft		Violence		Piracy	
Private local operator	-3.053**		-0.718**		-0.128	
	(1.192)		(0.358)		(0.115)	
Government operated		-4.670**		-0.452**		-0.016
		(1.892)		(0.191)		(0.039)
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	3497	3497	3497	3497	3497	3497
R^2	0.754	0.755	0.154	0.154	0.313	0.313

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 oilfields 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. Violence is the total number of oil-related conflict 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 B3: The effect of divestment on output and criminality, robustness to prices

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Output and efficiency						
Outcome	Shut-in		Output		Malfunctions	
Local operator	-0.172**	-0.167**	1.485	1.487**	0.640	-0.141
	(0.069)	(0.068)	(0.921)	(0.671)	(0.992)	(1.083)
Treated $\times$ Oil price (USD/barrel)	-0.000	-0.000	-0.011	-0.022	-0.049***	-0.064***
	(0.001)	(0.001)	(0.015)	(0.014)	(0.013)	(0.015)
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	2464	2464	2464	2464	3497	3497
R^2	0.657	0.670	0.862	0.879	0.574	0.633
Panel B: Crime and violence						
Outcome	Theft		Violence		Piracy	
Local operator	-4.311***	-4.848***	-0.439**	-0.445**	-0.110	-0.104
	(1.032)	(1.349)	(0.199)	(0.223)	(0.086)	(0.076)
Treated $\times$ Oil price (USD/barrel)	-0.052*	-0.077***	0.015***	0.016***	0.001	0.000
	(0.028)	(0.029)	(0.004)	(0.005)	(0.002)	(0.002)
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	3497	3497	3497	3497	3497	3497
R^2	0.714	0.755	0.131	0.156	0.235	0.313

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 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. Violence is the total number of oil-related conflict 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 B4: The effect of divestment on output and criminality, robustness to amnesty policy

	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Output and efficiency					
Outcome	Shut-in		Output		Malfunctions	
Local operator	-0.169**	-0.165**	1.757**	1.884***	1.472	0.858
	(0.067)	(0.066)	(0.732)	(0.551)	(0.908)	(0.952)
Post-amnesty × Amnestied	-0.024	-0.011	0.796*	-0.003	-1.280	-1.980*
	(0.037)	(0.068)	(0.413)	(0.545)	(0.821)	(1.088)
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	2464	2464	2464	2464	3497	3497
R^2	0.657	0.670	0.862	0.878	0.573	0.632
	Panel B: Crime and violence					
Outcome	Theft		Violence		Piracy	
Local operator	-3.343***	-3.290**	-0.711***	-0.741**	-0.152*	-0.134
	(1.136)	(1.280)	(0.259)	(0.287)	(0.092)	(0.081)
Post-amnesty × Amnestied	1.254	2.864**	-0.326	-0.128	-0.454***	-0.400***
	(1.269)	(1.390)	(0.209)	(0.364)	(0.105)	(0.120)
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	3497	3497	3497	3497	3497	3497
R^2	0.714	0.755	0.130	0.155	0.254	0.319

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. 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. Violence is the total number of oil-related conflict deaths within 15 km of the field. Piracy is the number of pirate attacks within 15 km of the field. Amnestied is a dummy for being within 30 kilometers of an amnestied militant camp, and post-amnesty is a dummy for post-2009. 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 B5: The effect of divestment on output and criminality, robustness to field-level covariates

Outcome	Shut-in	Output	Malf.	Theft	Violence	Piracy
	(1)	(2)	(3)	(4)	(5)	(6)
Local operator	-0.148**	1.957***	1.697*	-3.605***	-1.207***	-0.100
	(0.060)	(0.630)	(0.922)	(1.355)	(0.459)	(0.089)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Spatial controls $\times$ Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Field controls $\times$ Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2374	2374	3316	3316	3316	3316
R^2	0.679	0.885	0.651	0.760	0.206	0.327

Standard errors, in brackets, are clustered at the field level. Outcome variable given in table header. 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. Violence is the total number of oil-related conflict deaths within 15 km of the field. 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 number of wells, initial year, onshore dummy, and maximum well depth. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

more likely to be be onshore. If multinationals abandoned fields with these characteristics because they were experiencing differential trends in output and theft over the sample period, this could contaminate the results. In Table B5, I test robustness to including interactions between fixed field characteristics and time dummies in the main TWFE equation. Note that the sample size falls to 2,374 field-years for output and 3,316 for other outcomes because 15 fields have missing characteristics. Despite this, the results are unchanged.

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 B6, 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.

D.5 Randomization inference

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 B6. 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.

Table B6: The effect of divestment on output and revenue, measurement error

Sample	Single-operator				No shut-in	
Outcome	Q	$\log(Q)$	R	$\log(R)$	Q	R
	(1)	(2)	(3)	(4)	(5)	(6)
Local operator	0.800	1.014***	65.912*	1.014***	2.574***	156.791***
	(0.550)	(0.290)	(37.356)	(0.290)	(0.766)	(58.552)
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	2386	1823	2386	1823	1881	1881
R^2	0.881	0.771	0.854	0.778	0.887	0.862

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.. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure B6: Randomization inference

(a) β (b) t-statistic

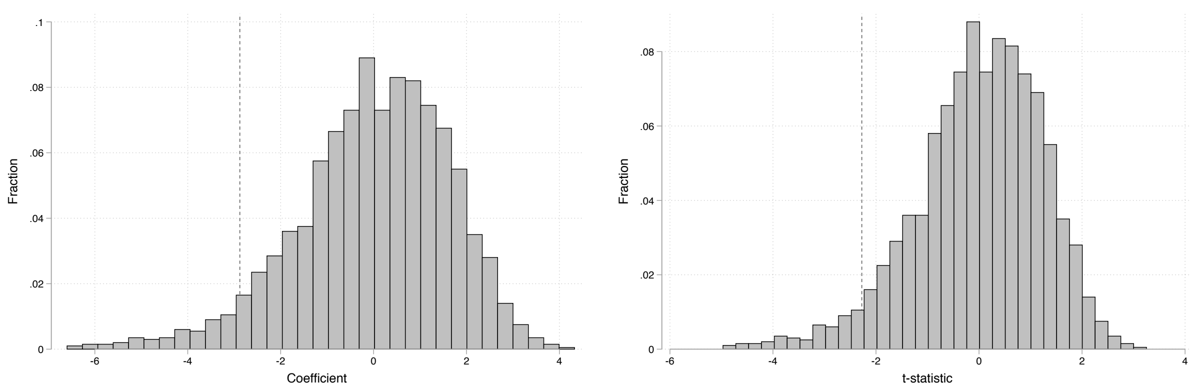

Note: Figure shows histograms of coefficient estimates (Panel A) and t-statistics (Panel B) for 2000 draws of a randomization inference routine. Outcome variable is theft, the total number of sabotage spills within 15 km of the field. Vertical line indicates the estimate for the observed data.

Table B7: The effect of divestment on oil theft by pipeline type

Asset type	Trunkline		Flowline		Delivery line		Wellhead	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Local operator	-1.100**	-0.745	-0.067	-0.070	-1.371*	-1.982**	-0.504***	-0.620***
	(0.540)	(0.566)	(0.222)	(0.272)	(0.746)	(0.864)	(0.186)	(0.237)
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	No	Yes	No	Yes	No	Yes
Observations	3497	3497	3497	3497	3497	3497	3497	3497
R^2	0.394	0.452	0.453	0.504	0.770	0.808	0.422	0.526

Standard errors, in brackets, are clustered at the field level. Outcome variable given in table header. Theft measured as the total number of sabotage spills within 15 km of the field on a particular asset type. 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 number of wells, initial year, onshore dummy, and maximum well depth. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

D.6 Economic value of theft

In the main results of Table 1, I measure theft as the number of sabotage incidents within 15 kilometers of the oilfield. This variable does not directly measure the economic value of losses due to theft. As such, a reduction in theft incidents may not correspond to a reduction in quantity losses if the localization affects the type of theft, for example, by incentivizing fewer but larger thefts. Unfortunately, we lack detailed data on the size of thefts. However, a reasonable proxy can be derived by exploiting information on the type infrastructure targeted by the theft. In particular, the most lucrative assets are trunklines, delivery lines, flow lines, and wellheads. This is because trunklines are large pipelines that aggregate flows from multiple fields to funnel toward export terminals, while the other pipelines move smaller volumes of oil between or within fields.

In order for total quantity stolen to rise even as aggregate incident counts fall, it must be the case that thefts on larger assets rise enough to more than offset the reduction in theft on smaller targets. I test this in Table B7, re-estimating the main DD specification for oil theft, using thefts on a particular asset type as the outcome variable. I find that the point estimates for each of the asset types is negative, and significant for all except flowlines. There is no evidence that thefts increase as a result of localization for any of the asset types. It is therefore highly unlikely that stolen quantities would increase despite an overall reduction in aggregate theft incidents.

D.7 Location-specific time trends

Outcomes may have evolved differently in localities that have relatively more indigenized fields. For example, localities where many fields were localized may also have had an improving security situation over the sample period for reasons unrelated to localization per se. To control for differential location-specific time trends, I include locality-by-year interacted fixed effects, using both states and local government areas as larger and smaller geographic areas. The resulting specification essentially compares fields within a given locality, with the final estimate a weighted average across localities of these within-locality comparisons. Table B8 presents the results. The impacts of localization on shut-ins, output, and theft remain significant and are similar in magnitude to the main results. The impacts on malfunctions, oil-related violence, and piracy are now smaller and no longer significant. However, the point estimates are of the correct sign.

D.8 TWFE diagnostic tests

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 (Chaisemartin and D'Haultfoeuille 2019, Goodman-Bacon 2019, Callaway and Sant'Anna 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.

Chaisemartin and D'Haultfoeuille (2019) provide some guidance derive a general formula for the unit-time-specific weights of the treated observations, which allows diagnostic testing on the share of negative weights. Figure B7 displays histograms of estimated weights for each of the 6 outcomes in Table 1. 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 B8. 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 B9 presents weights and average treatment effect estimates for each 2×2 DD comparison

Table B8: The effect of divestment on output and criminality, robustness to fixed effects

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Output and efficiency						
Outcome	Shut-in		Output		Malfunctions	
Local operator	-0.182***	-0.232***	1.644**	1.983***	1.101	0.530
	(0.068)	(0.084)	(0.665)	(0.650)	(0.940)	(1.347)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
State $\times$ Year FE	Yes	No	Yes	No	Yes	No
Locality $\times$ Year FE	No	Yes	No	Yes	No	Yes
Observations	2476	2476	2476	2476	3497	3497
R^2	0.684	0.769	0.867	0.893	0.629	0.749
Panel B: Crime and violence						
Outcome	Theft		Violence		Piracy	
Local operator	-3.228***	-3.339*	-0.460*	-0.150	-0.148	0.026
	(1.085)	(1.795)	(0.235)	(0.142)	(0.095)	(0.154)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
State $\times$ Year FE	Yes	No	Yes	No	Yes	No
Locality $\times$ Year FE	No	Yes	No	Yes	No	Yes
Observations	3497	3497	3497	3497	3497	3497
R^2	0.754	0.859	0.520	0.740	0.315	0.687

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. 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. Violence is the total number of oil-related conflict deaths within 15 km of the field. Piracy is the number of pirate attacks within 15 km of the field. Localities are local government areas, the lowest level administrative unit in Nigeria. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure B7: Histogram of weights from Chaisemartin and D'Haultfoeuille (2019)

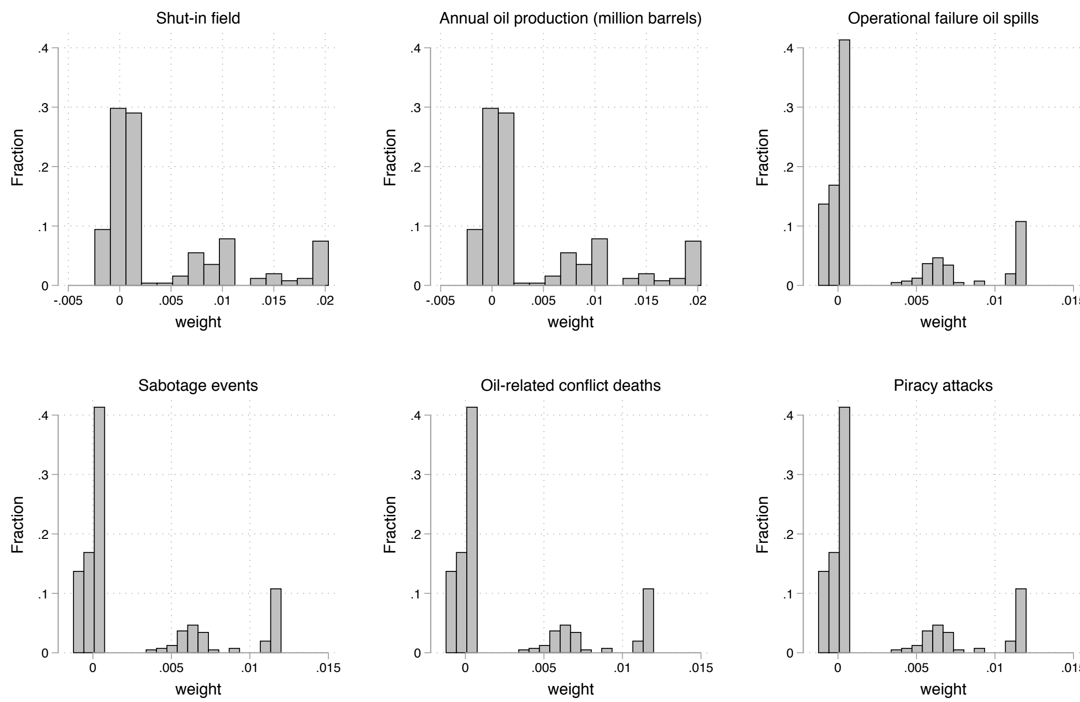

Note: Figure shows implied weights for unit-and-time-specific average treatment effects, as derived in the Chaisemartin and D'Haultfoeuille (2019) decomposition results. I display histograms of the weights for each of the six key outcomes analyzed in Table 1. Sample is the panel of 314 oilfields from 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. Violence is the total number of oil-related conflict deaths within 15 km of the field. Piracy is the number of pirate attacks within 15 km of the field.

Figure B8: Dynamic effects from Chaisemartin and D'Haultfoeuille (2019)

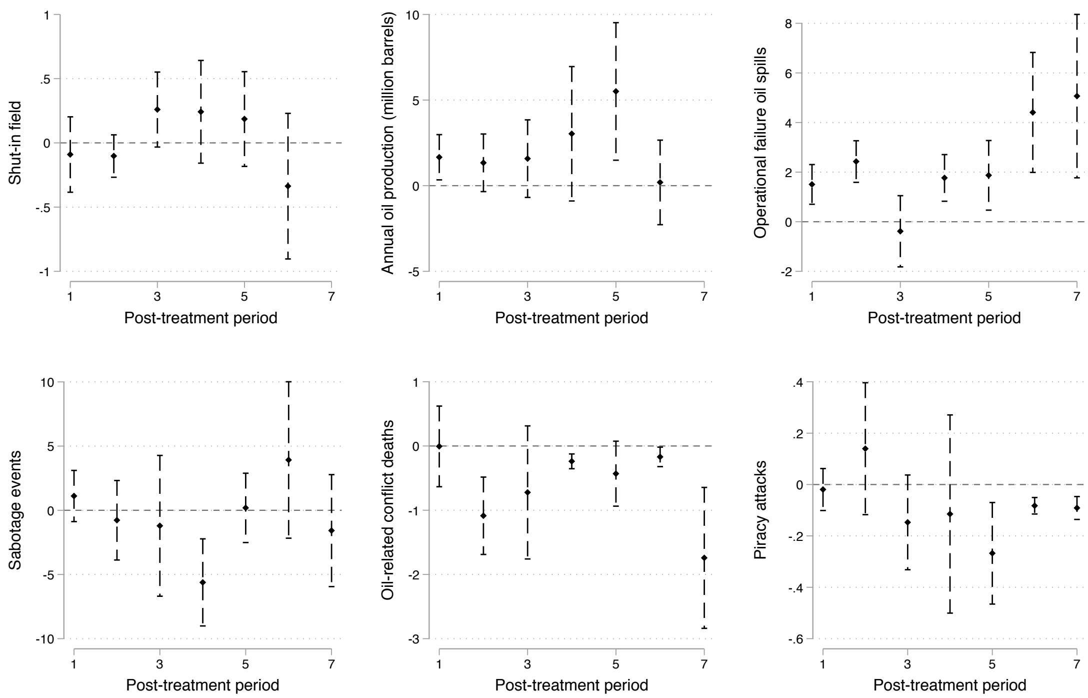

Note: Figure shows dynamic effects among switchers using the Wald-type estimator for staggered adoption designs described in Chaisemartin and D'Haultfoeuille (2019). Standard errors are clustered at the field level and computed using a bootstrapping procedure with 1000 replications.

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.

Table B9: Goodman-Bacon (2019) 2×2 DD weights

DD Comparison	Weight	DD Estimate
Earlier T vs. Later C	0.025	-6.782
Later T vs. Earlier C	0.010	-7.526
T vs. Never treated	0.866	-3.108
T vs. Already treated	0.099	-2.152
TWFE estimate		-3.146

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.

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.[62] One way to address this issue is to run event-study regressions as in Figure 3.[63]

An alternative estimation method is the stacked DD (see Gormley and Matsa 2011, Deshpande and Y. 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, 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 be-

[62] This is identical to the "negative weights" problem identified in Chaisemartin and D'Haultfoeuille (2019).

[63] 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 B12

tween 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 B10. 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 1. 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 B9.

In columns (4)-(6) of Table B10 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 shutins are all robustly significant. I also estimate event-studies in the stacked format, the results of which are displayed in Figure B10. 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 B9. As desired, they are clustered around large negative and positive values, respectively.

Callaway and Sant'Anna (2019) propose a semi-parametric DD estimator to address the TWFE issues of "negative weights" problem (use of treated observations as controls in post-treatment periods) and the down-weighting early and late-treated groups in the presence of cohort-specific heterogeneity. The estimator computes propensity-score-weighted ATT effects for each cohort-period, and then aggregates these estimates using various weighting schemes. It is similar in spirit to the stacked model in that it emphasizes cohort-specific variation and uses only the untreated as controls. However, it does not rely on a linear parametric specification, and allows for more flexible re-weighting in the aggregation of cohort-and-time-specific ATT parameters.

I present the aggregate ATT estimates in Table B11 using four different weighting strategies, explained in detail in Callaway and Sant'Anna (2019).[64] In general, the estimates are of the correct sign, significant at 5%, and similar in magnitude across aggregation methods. One notable exception is the effect on oil theft under the selective timing aggregation method (column 2), which weights based on cohort-size rather than length of treatment, and produces a positive and insignificant coefficient. However, in the presence of dynamic effects that grow over time – which we observe in Figures 3 and B10 – up-weighting a large cohort with very few post-treatment periods could generate misleading estimates. Since our largest cohort is indeed treated in 2016 and has just one post-treatment period, this may bias the estimate.

[64] The sample in this estimation is only the balanced panel of 256 fields from 2006-2017.

Table B10: Stacked Differences-in-Differences estimates

Sample	Full			Ever-treated		
	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A: Theft			
Local operator	-5.630***	-3.438***	-3.133***	-9.971***	-5.774***	-6.077***
	(1.080)	(0.982)	(1.142)	(1.568)	(1.395)	(2.151)
Observations	30962	30962	30962	2992	2992	2990
R^2	0.684	0.715	0.715	0.683	0.721	0.723
			Panel B: Violence			
Local operator	-0.936***	-1.180***	-1.267***	0.637	0.219	0.091
	(0.353)	(0.356)	(0.451)	(0.468)	(0.424)	(0.637)
Observations	30962	30962	30962	2992	2992	2990
R^2	0.178	0.203	0.203	0.429	0.461	0.465
			Panel C: Piracy			
Local operator	-0.066	-0.053	-0.086	0.282***	0.194*	0.395**
	(0.078)	(0.079)	(0.095)	(0.103)	(0.103)	(0.165)
Observations	30962	30962	30962	2992	2992	2990
R^2	0.219	0.249	0.250	0.256	0.323	0.329
			Panel D: Output			
Local operator	2.477***	2.471***	2.255***	2.950***	3.124***	3.601***
	(0.587)	(0.591)	(0.582)	(0.645)	(0.860)	(1.148)
Observations	22234	22234	22168	2194	2194	2184
R^2	0.870	0.870	0.870	0.687	0.696	0.708
			Panel E: Malfunctions			
Local operator	1.222	2.204***	1.857*	1.222	1.834*	2.984**
	(0.950)	(0.838)	(0.980)	(1.022)	(0.984)	(1.448)
Observations	30962	30962	30962	2992	2992	2990
R^2	0.532	0.570	0.570	0.562	0.589	0.596
			Panel F: Shut-ins			
Local operator	-0.186***	-0.196***	-0.176**	-0.387***	-0.402***	-0.472***
	(0.069)	(0.067)	(0.072)	(0.077)	(0.076)	(0.097)
Observations	22234	22234	22168	2194	2194	2184
R^2	0.673	0.682	0.681	0.458	0.539	0.548
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 D. 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure B9: Stacked-DD histogram over event-windows

(a) β, theft

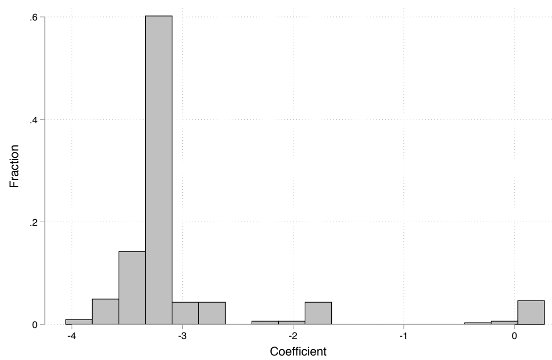

(b) t-statistic, theft

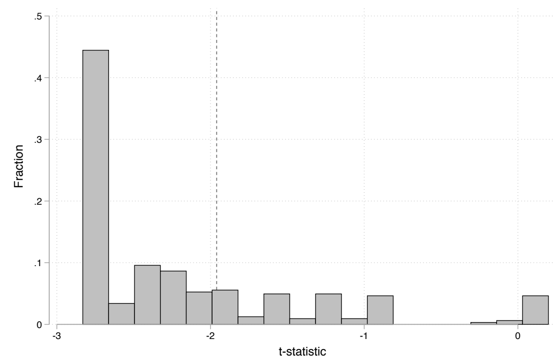

(c) β, output

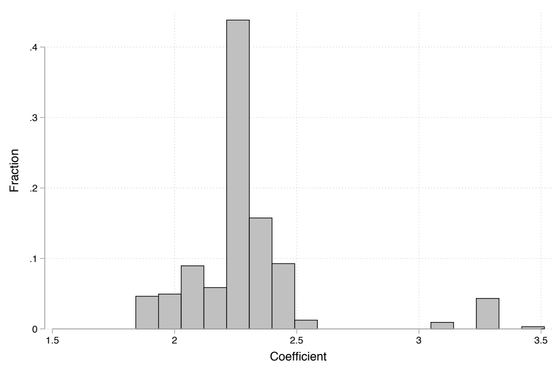

(d) t-statistic, output

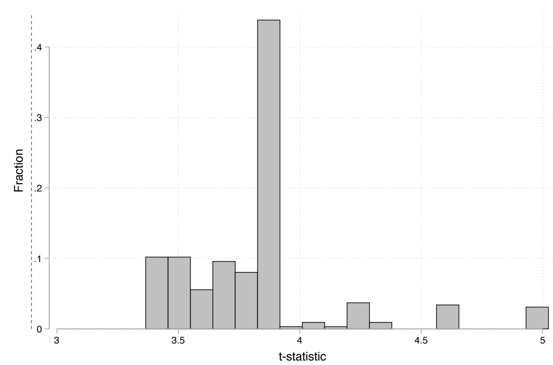

Note: Figure shows histograms of coefficients and t-statistics from the stacked-DD specification described in Section D for oil production and theft outcomes. For each outcome, I estimate treatment effects for all possible combinations of event windows up to 18 years before and 18 years after the event 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 B10: Stacked-DD event-study

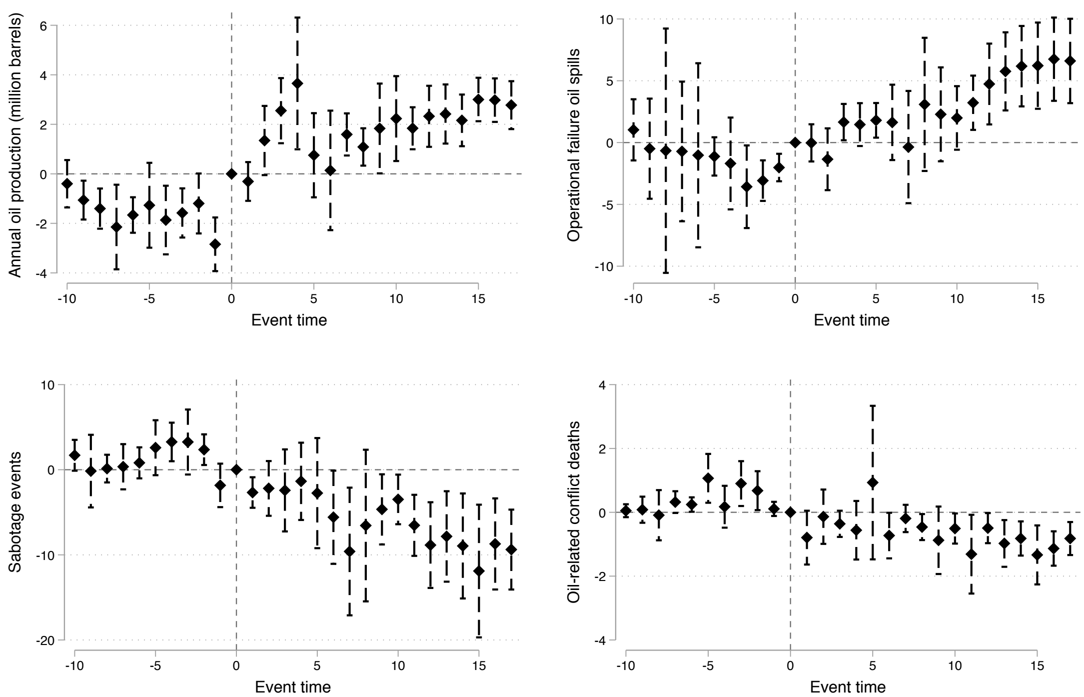

Note: Figure shows coefficients from event-study regressions of the stacked-DD specification described in Section D for oil production, crime, and violence outcomes. Standard errors are clustered at the field-by-event-cohort level. Output is measured in millions of barrels of oil per year. Malfunctions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

Table B11: Re-weighted semi-parametric DD estimates

Outcome	Simple (1)	Selective (2)	Dynamic (3)	Calendar (4)
Output	2.20	1.69	1.75	1.50
	(1.04)	(0.50)	(0.78)	(0.60)
Malfunctions	2.41	2.00	1.05	1.87
	(0.65)	(0.67)	(0.81)	(0.49)
Oil theft	-0.50	1.47	-3.98	-2.06
	(1.21)	(1.48)	(2.01)	(1.57)
Violence	-0.55	-0.41	-0.44	-0.30
	(0.26)	(0.19)	(0.15)	(0.11)
Piracy	-0.08	-0.02	-0.08	-0.05
	(0.05)	(0.05)	(0.05)	(0.04)

Bootstrapped standard errors in parentheses. Table displays difference-in-differences estimates using the semiparametric reweighted estimator described in Callaway and Sant'Anna 2019. Estimation sample is a fully balanced panel of 256 fields from 2006-2017 (3,072 field-year observations). Re-weighting procedure in the aggregation of cohort-and-time-specific ATTs is given in table heading, while outcome variable is given in the leftmost column.

The preferred specifications are the dynamically-weighted estimates in column (3), which are highly consistent with the main results of Table 1. Figure B11 presents post-period estimates re-weighted to account either only for dynamic heterogeneity (left panels) or for dynamic and cohort-specific heterogeneity (i.e., selective timing that changes the cohort-composition of the treatment group in any given post-period, right panels). These estimates are almost always of the right sign and give similar dynamic patterns, though the selective timing estimates are substantially more precise.

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

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

And then form the re-weighted event-study treatment effect $\tilde{\psi}_\tau = \sum_c \hat{\psi}_\tau^c \omega_\tau^c$, where $\omega_\tau^c =$

Figure B11: Re-weighted dynamic effects

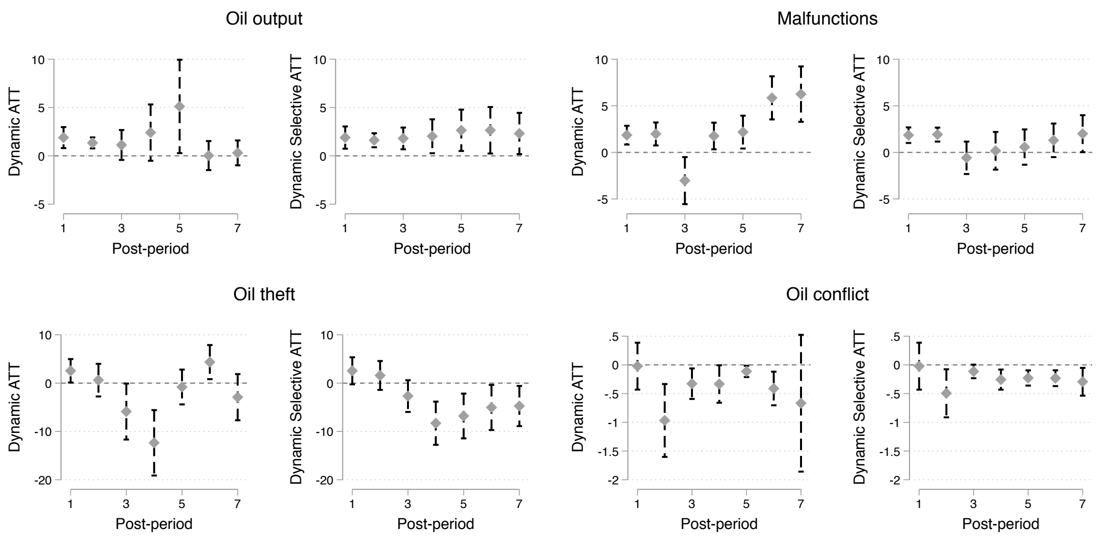

Note: Figure displays dynamic difference-in-differences estimates using the semiparametric re-weighted estimator described in Callaway and Sant'Anna 2019. Estimation sample is a fully balanced panel of 256 fields from 2006-2017 (3,072 field-year observations). Re-weighting procedure in the aggregation of cohort-and-time-specific ATTs is the estimator accounting for dynamic treatment effects alone (left panels) or accounting for dynamic treatment effects and selective timing (right panel). Bootstrapped standard errors.

$\frac{\sum_i L_{it}^{\tau} 1(t_i=c)}{\sum_i L_{it}^{\tau}}$. The results are in Figure B12, 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.

Figure B12: Event-study, Abraham and Sun (2018) weights

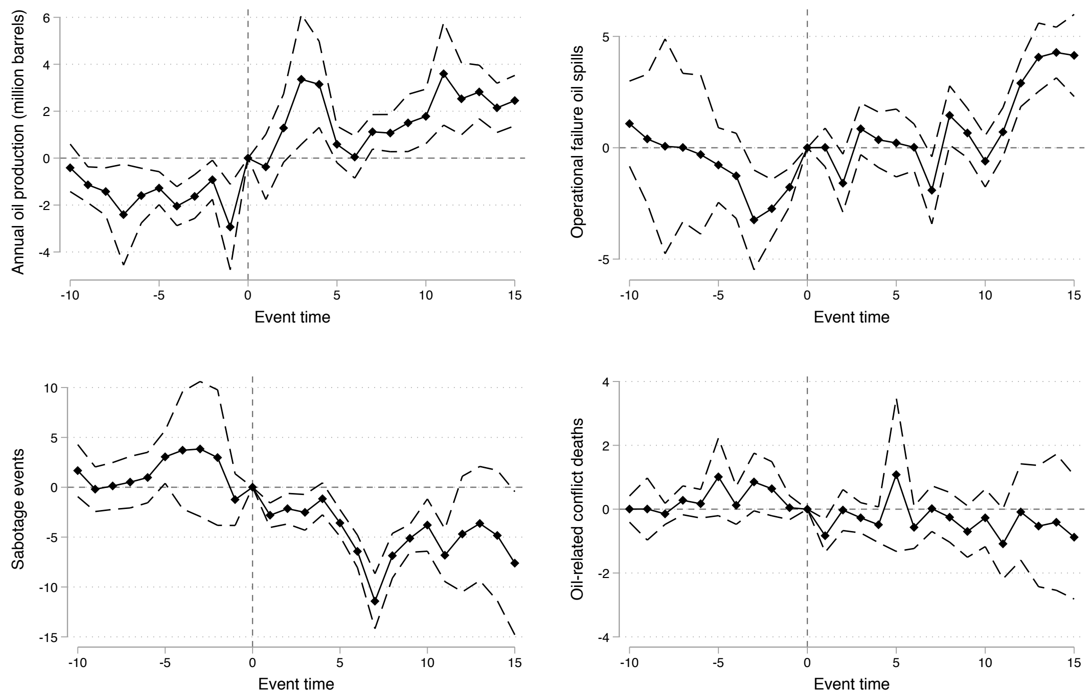

Note: Figure shows re-weighted coefficients from the cohort-specific event-study regression described in Abraham and Sun (2018). These point estimates come from a fully-saturated event-study regression of the outcome of interest on pre-and-post-treatment dummies are interacted with dummies indicating 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.

E Additional mechanisms

E.1 Partial ownership

Partial ownership drives a wedge between the losses to the operating firm and criminal profits; operators with larger ownership stakes γ internalize a greater share of the losses from theft. The Nigerian oil market exhibits substantial variation in ownership agreements (see Figure A1), and local operators may have greater ownership stakes for several reasons. Firstly, multinational divestment may lead to consolidation of stakes in joint ventures. Secondly, because of indigenization policies, local firms are more likely to obtain sole-risk contracts than multinationals, who must provide mandated equity stakes to government. Multinationals are 33.5 p.p. more likely to be in joint ventures and 43 p.p. less likely to obtain sole-risk licenses. As a result, the average multinational concession has a government stake roughly 85% higher than the average Nigerian independent operator.

These descriptive statistics indicate that it is at least plausible that greater ownership stakes allow local firms to more efficiently internalize losses. However, field-level characteristics could be driving these correlations – multinationals own larger fields where government has a greater incentive to increase its stake, or offshore fields where greater financing requirements necessitate joint ventures. To test whether localization causally increases consolidation, I re-estimate the main TWFE regression at the concession-year level, where the outcome variable is either the concession equity Herfindahl-Hirschman Index (HHI), which measures overall consolidation, or the operator's stake, which corresponds directly to γ.[65]

Table B12 presents the results. Columns (1)-(3) estimate the model with the HHI outcome, while columns (4)-(7) use operator share. Columns (1) and (4) give the unconditional relationship, (2) and (5) include year fixed effects, while (3) and (6) include both year and block fixed effects. All specifications control for concession type dummies (joint-venture vs. sole-risk), asset type (onshore vs. offshore), and concession size (area, number of fields, and number of wells). In the full TWFE specification with interacted controls, local divestment increases the HHI by 0.087 p.p., a 16.7% increase on the multinational mean, significant at 5%. Local divestment also increases operator ownership by 12.8% p.p., a 20.1% increase, significant at 1%. The results indicate that divestment substantially increases ownership concentration in the hands local operators. Partial ownership is therefore an important mechanism driving local advantage.

E.2 Corruption penalties

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.[66]

[65] The sample is all concessions observed annually from 2013-2018.

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

Table B12: The effect of local ownership on equity consolidation

Outcome	HHI			Operator stake		
	(1)	(2)	(3)	(4)	(5)	(6)
Local operator	0.047	0.043	0.087**	0.097**	0.093**	0.128***
	(0.034)	(0.035)	(0.039)	(0.039)	(0.040)	(0.045)
Block FE	No	No	Yes	No	No	Yes
Year FE	No	Yes	Yes	No	Yes	Yes
Observations	541	541	538	541	541	538
R^2	0.341	0.352	0.935	0.485	0.495	0.941

Standard errors, in brackets, are clustered at the concession-block level. Sample is the panel of 113 concession blocks from 2013-2018. Outcome variable is indicated in table header; either the block-level equity HHI, or the equity stake of the operating firm. All specifications include dummy controls for joint-venture, sole-risk, and offshore, interacted with year dummies where these are included. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

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.

The results of this estimation for each of the six major outcomes are contained in Table B13. 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.

F Alternative explanations

F.1 Spatial spillovers

Gangs may not operate as local monopolists. In a general equilibrium setting, gangs may optimally choose targets for theft across all possible oil fields, rather than simply facing the

Table B13: The effect of corruption costs on output and criminality

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Output and efficiency						
Outcome	Shut-in		Output		Malfunctions	
Home-country corruption law	-0.029	-0.026	0.802***	0.149	0.639	-1.647**
	(0.031)	(0.037)	(0.273)	(0.305)	(0.562)	(0.748)
Control group mean	0.343		2.151		7.302	
Observations	2262	2262	2262	2262	3148	3148
R^2	0.673	0.683	0.866	0.881	0.569	0.635
Panel B: Crime and violence						
Outcome	Theft		Violence		Piracy	
Home-country corruption law	6.095***	2.966***	0.150*	0.454***	-0.226***	-0.191***
	(0.808)	(0.751)	(0.080)	(0.140)	(0.056)	(0.068)
Control group mean	8.522		0.320		0.247	
Observations	3148	3148	3148	3148	3148	3148
R^2	0.721	0.756	0.133	0.167	0.263	0.344
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

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. Violence is the total number of oil-related conflict 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

binary choice of accepting a bribe or stealing from a single field. As such, localization might generate important spillovers across fields. Localization could increase targeting of surrounding multinational fields if local fields are politically protected but their multinational neighbors are not, as gangs seek to recoup lost income on nearby multinational fields. In contrast, if local firms use their political connections to improve anti-crime enforcement by security forces, this could generate positive enforcement spillovers to nearby multinational firms if security is at least partially non-excludable. In either case, substantial spatial spillovers will severely bias the treatment effect by violating the stable unit treatment value assumption (SUTVA) (Rubin 2005), since nearby untreated fields experience some impact of treatment.

To test for spatial spillovers, I follow the "ring method" common in the urban economics literature (see e.g. Autor, Palmer, and Pathak 2014 and Diamond and McQuade 2019). In the stacked dataset (see Appendix D), for each event date, I identify all untreated fields. For each untreated field, I calculate the distance from that field to the nearest treated field. I then re-estimate the stacked difference-in-differences specification including interactions between the post-treatment indicator and dummy variables for treated fields, as well as dummies for control fields within each ten-kilometer interval from 0 to 100. The result is an estimate of the treatment effect, as well as spillover estimates at each distance "ring" around the treated fields. The omitted group of untreated fields greater than 100 kilometers away from a treated field acts as the "pure" control group. Because the conflict and theft outcomes are defined in a 15 km radius around the field, I omit the spillover coefficients for fields within rings under 30 kilometers, since in these fields there may be overlap which induces a positive spatial correlation in outcomes and therefore spurious spillover effects.

The results are in Figure B13, which plots the treatment effect, as well as coefficients at each ring from 30-40 to 90-100 km, for output, theft, and conflict (Panels A, B, and C).[67] In all cases, the main treatment effects remain strong; this indicates that mechanical spatial correlation in outcomes is not generating spurious treatment effects, since we obtain similar results whether the control group is defined as all untreated or only those further than 100 kilometers away. However, we also observe clear negative spillovers. Consider the results on theft in Panel B. Consistent with the main results, localized fields see a reduction in theft of between 3-4 incidents annually after the divestment. However, this is mirrored almost exactly by a statistically significant increase in theft of similar magnitude on multinational assets 30-40 kilometers away from a treated field. These negative spillovers then taper off to zero as we move to rings further from the treated field. A similar pattern obtains for conflict in Panel C, though these effects exhibit less tapering. As a result, output (Panel A) also exhibits significant negative spillovers, the largest of which equal 42.5% of the treatment effect. Overall, the results are suggestive of localization not only reducing crime for local firms, but also

[67]Here we measure conflict as non-militant oil-related fatalities, since, as shown in Figure A6, non-militant violence is more affected by localization.

redirecting it toward multinationals. As such, the partial and general equilibrium effects of localization may differ substantially.

F.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 16,211 working-age[68] 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

$$y_{ivft} = \alpha + \psi local_{ft} + \delta_t + \xi_f + X'_{ivft}\beta + \mu_{ivft}$$

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 B14. 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.

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

[68] Defined as ages 15-60.

Figure B13: Spillovers

(a) Output

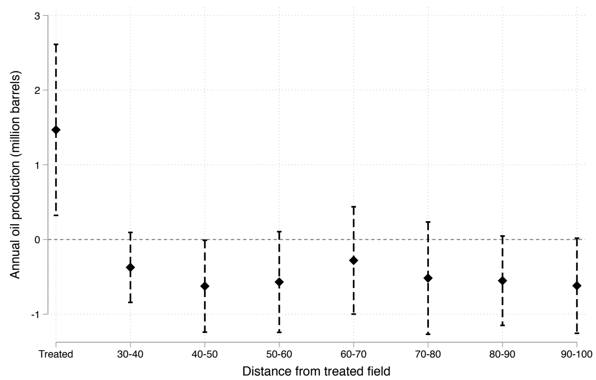

(b) Theft

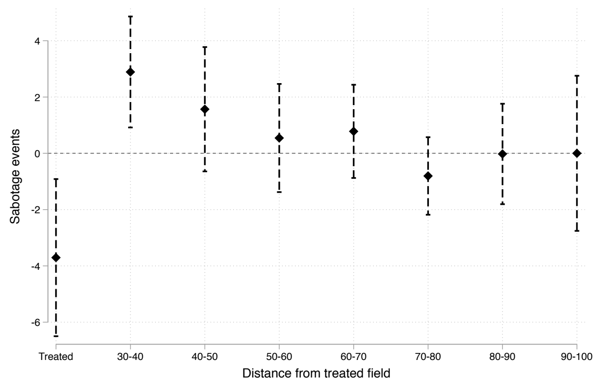

(c) Conflict

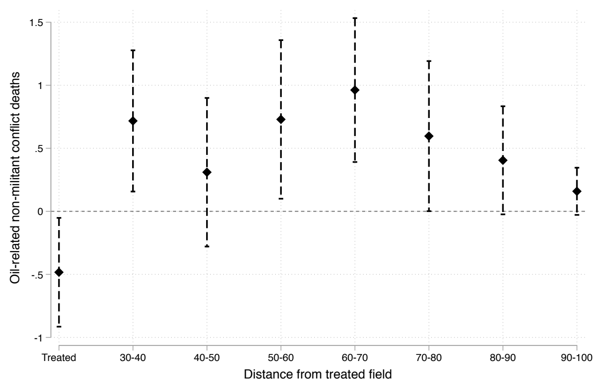

Note: Figure plots coefficient estimates of treatment effect and spillover effects for output (Panel A), theft (Panel B) and non-militant oil-related conflict deaths (Panel C). Estimates are derived from a stacked difference-in-differences regression (described in Section D) of the outcome on a dummy for post-treatment interacted with indicators for "ring" distances from the nearest treated field. Omitted control group is untreated fields further than 100km from the nearest localized field. All specifications include stack, time, and unit fixed effects and their interactions.

Table B14: Local ownership and local employment

Sample	All				Onshore	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Employed						
Local operator	-0.012	-0.019	-0.001	0.003	-0.015	-0.004
	(0.020)	(0.020)	(0.020)	(0.027)	(0.021)	(0.021)
Observations	16211	16211	16211	16211	15616	15616
R^2	0.029	0.030	0.038	0.035	0.029	0.038
Panel B: Employed outside home						
Local operator	-0.007	-0.011	-0.027	0.013	-0.015	-0.033
	(0.026)	(0.025)	(0.026)	(0.024)	(0.026)	(0.025)
Observations	8892	8892	8892	8892	8551	8551
R^2	0.107	0.108	0.160	0.132	0.095	0.140
Panel C: Self-employed						
Local operator	0.081**	0.087**	0.057*	0.064	0.084**	0.062*
	(0.034)	(0.036)	(0.031)	(0.042)	(0.035)	(0.033)
Observations	8892	8892	8892	8892	8551	8551
R^2	0.068	0.069	0.122	0.103	0.066	0.122
Panel D: Employed in household agriculture						
Local operator	-0.059	-0.067	-0.009	-0.033	-0.058	-0.016
	(0.046)	(0.049)	(0.062)	(0.042)	(0.046)	(0.063)
Observations	8892	8892	8892	8892	8551	8551
R^2	0.150	0.151	0.273	0.179	0.150	0.277
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 × State FE	No	No	No	Yes	No	No
Controls × 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 B15: Local ownership and local consumption

Outcome	log(consumption)					
Sample	All				Onshore	
	(1)	(2)	(3)	(4)	(5)	(6)
Local operator	0.138	0.133	0.127	0.006	0.139	0.119
	(0.146)	(0.142)	(0.084)	(0.074)	(0.140)	(0.088)
Observations	4909	4909	4909	4909	4750	4750
R^2	0.242	0.244	0.294	0.268	0.251	0.305
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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

(Panel C) by roughly 6-9 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 B15. 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 B14. All results suggest that pre-trends are essentially flat and insignificant for each outcome considered in Table B14. 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 Chaisemartin and D'Haultfoeuille (2019) and Goodman-Bacon (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 B14 Panel A, despite a small positive effect in some of the event-study coefficients.

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 demo-

Figure B14: Local ownership and local employment, parallel trends

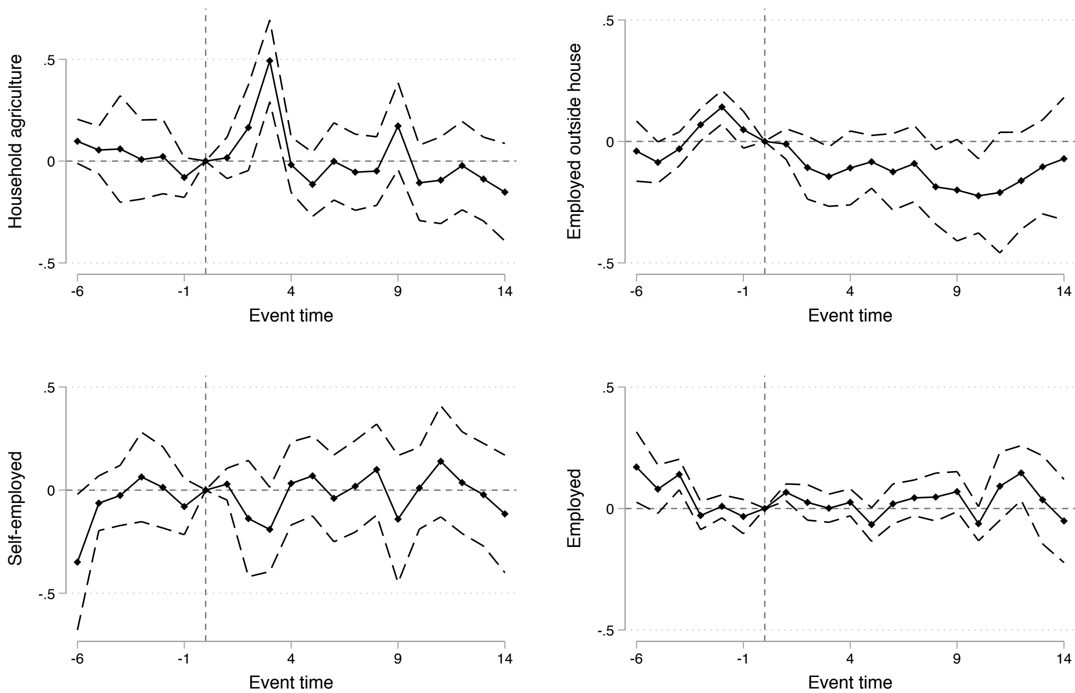

Note: Figure displays coefficients of event-study regressions of employment outcomes on pre-and-post treatment indicators for localization, conditional on unit and year fixed effect and controls interacted with year dummies. Employment outcomes are given in each subfigure. Sample is all individuals in 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 graphic 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 B15, 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.

Figure B15: Local ownership and local employment by age and gender

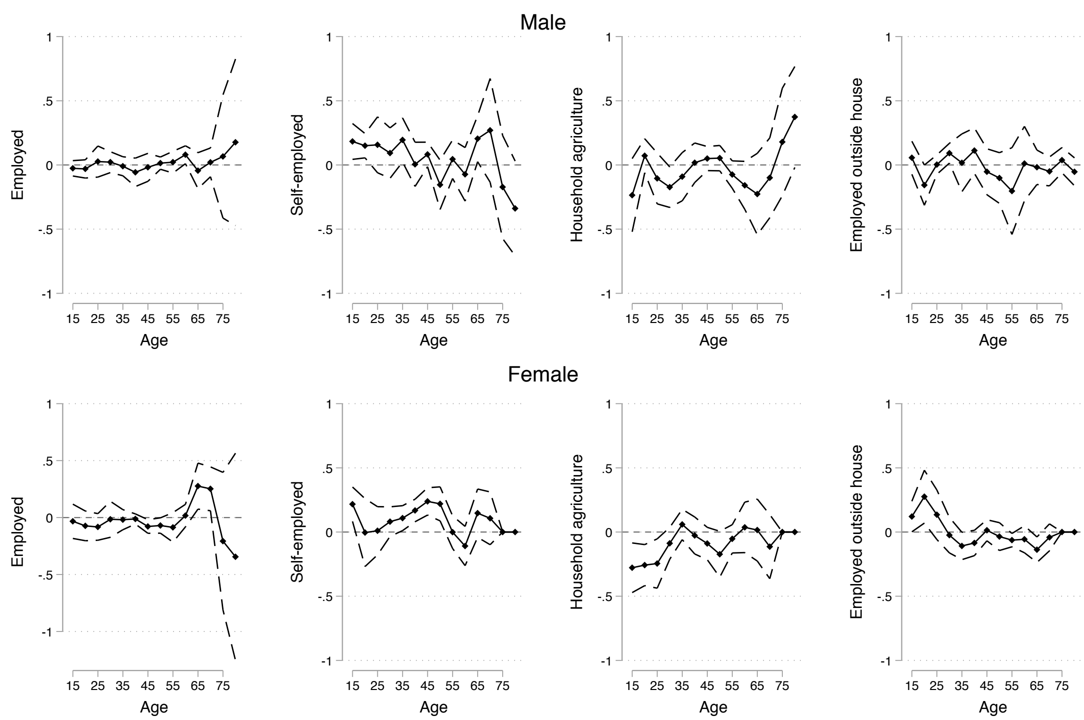

Note: Figure shows coefficients from differences-in-differences regressions of employment outcomes on local ownership of the nearest oilfield. Sample is all individuals in the three waves of the GHS above the age of 10 living in clusters within 50 km of an oilfield. All regressions use household-level 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.

F.3 Targeted CSR investment

The most visible local benefits of oil extraction are typically not jobs but rather host community investments in the form of corporate social responsibility (CSR). It may be the case that positive localization effects on local communities do not show up on average in employment because the benefits are targeted specifically to problem hotspots in the form of CSR investment. It may indeed be more efficient for an oil company to provide CSR benefits to troubled areas to dissuade militancy and theft than to negotiate with organized crime directly.

In 2016, voluntary expenditures on CSR projects by oil companies in host communities totaled 92.6 million dollars, 72% of which was spent by multinationals. This is a tiny fraction of the annual profits from oil theft, suggesting that these projects are unlikely to meaningfully dissuade violence. However, if local firms have a greater propensity to target their investment toward volatile communities, this mechanism could plausibly drive the observed effects. I test this hypothesis using data on 508 community-specific CSR projects in 2016, the only year for which comprehensive data is publicly available. I regress the number and value of multinational or local projects at the village level in 2016 on the level of oil militant conflict in 2015, measured as either the cumulative number of militant attacks from 1997-2015 (measuring long-run conflict) or the number of militant attacks in 2015. I also include state or locality fixed effects for robustness to geography-specific unobserved heterogeneity. Given that we only observe a single cross-section, the results should be taken as purely correlational. Still, if companies follow a targeting policy, we should at minimum observe a reduced form positive correlation between conflict and CSR projects.

Figure B16 plots coefficients from these regression models. In Panel A and B, I use standardized CSR projects as the outcome to account for the fact that local firms are generally smaller and therefore have fewer projects overall, while Panels C and D use total CSR expenditure in millions of USD. Panels A and C use cumulative attacks up to 2015 on the righthand side, while Panels B and D use attacks in 2015, controlling for lagged (2014) attacks. For each specification, I estimate the unconditional bivariate relationship, as well as models with state or locality fixed effects. In general, there is evidence suggestive of targeting – local conflict is positively and significantly correlated with the number and value of CSR projects at the village-level. However, this aggregate relationship obscures substantial differences between local and multinational projects. Across all outcomes and independent variables, the correlation between CSR investments and conflict is much stronger for multinational projects. This suggests that the main results are unlikely to be driven by superior targeting by local firms. If anything, the results are consistent with multinationals leaning more heavily on CSR to mitigate conflict risk that local firms because they face greater costs of bargaining directly with gangs.

Figure B16: CSR projects and local conflict

(a) CSR projects, cumulative attacks

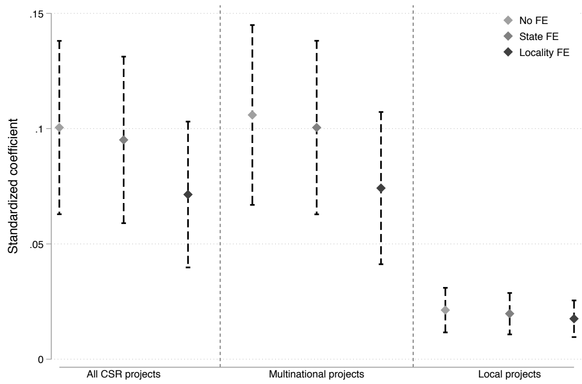

(b) CSR projects, previous year attacks

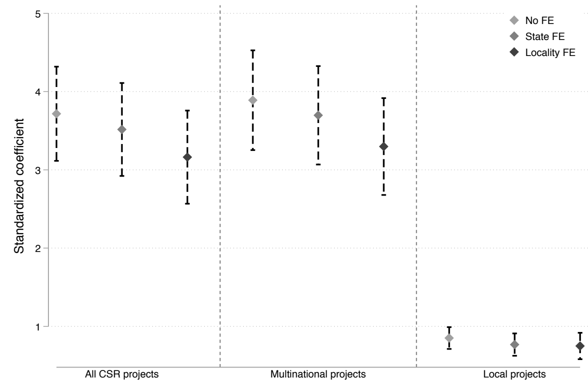

(c) CSR expenditure, cumulative attacks

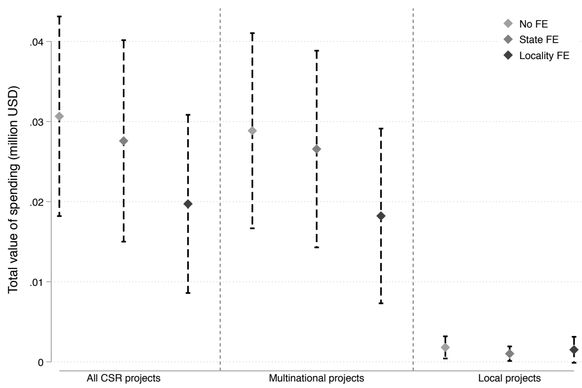

(d) CSR expenditure, previous year attacks

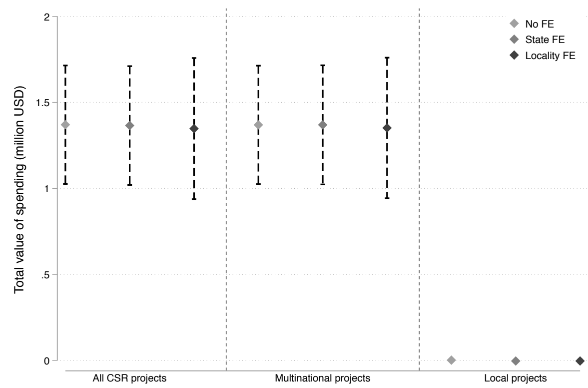

Note: Figure plots coefficient estimates of the village-level correlation between oil company expenditure on corporate social responsibility (CSR) in 2016 and lagged militant activity. The outcome is measured as either the standardized number of CSR projects or total expenditure, either in total or disaggregated by local and multinational projects. The independent variable is measured as the number of oil-related militant attacks in 2015 or the cumulative number oil-related militant attacks from 1997-2015. Model specification is indicated in subfigure headers. Models are either unconditional or include state or locality fixed effects, indicated in subfigure legends.

G Heterogeneous effects

G.1 Heterogeneity: oil prices

For an increase in prices to increase theft, we must have $\frac{\kappa}{q} < \frac{(1+\tau)}{\gamma} - 1$, an expression that depends on the bargaining friction τ. 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

$$y_{it} = \alpha + \theta_0 p_t + \theta_1 local_{it} + \theta_2 local_{it} p_t + \xi_i + X_{it}'\beta + v_{it}$$

Where p_t is the demeaned world price of crude oil, relative to the long-run mean.[69] The empirical implication is that while the sign of θ_0 is ambiguous, θ_2 must be negative, that is, higher prices have increasingly negative effects on theft as bargaining costs fall. Furthermore, if $\theta_0 > 0$, then it is likely that $\frac{\kappa}{q}$ is small relative to $\frac{(1+\tau)}{\gamma} - 1$; oil theft has a low ratio of spillage losses to illicit gains. In order to identify θ_0 in the fixed effects model, I exclude time dummies δ_t in some specifications.

The results are given in Table B16. 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 (1) and (2) include ξ_i. In columns (2) and (4) controls are additionally interacted with p_t to control for potential omitted variables correlated with localization that might respond similarly to oil price trends. Columns (5)-(6) estimate the full TWFE specification.

As predicted, the interaction coefficient θ_2 is robustly negative and significant in all specifications. The estimates imply that the average responsiveness of theft to price is roughly 0.07 to 0.13 incidents lower on locally-operated fields than among multinationals. The coefficients on p_t in columns (1)-(4) show that on multinational assets theft increases in prices, significant at 1% in (1) and (3). Putting these estimates together, the positive effect of prices on theft falls 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 low relative to the quantity stolen q, which leads the reservation price of firms to be less sensitive to oil prices. Finally, since p_t is demeaned, the estimates of θ_1 imply that at long-run average prices, localization effects are large, negative, and significant. 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.

[69] The mean is calculated over the period 2006-2017.

Table B16: The effect of divestment on oil theft by prices

Outcome	Oil theft					
	(1)	(2)	(3)	(4)	(5)	(6)
Local operator	-7.035***	-7.818***	-0.905	-0.702	-4.719***	-5.164***
	(1.127)	(1.456)	(1.033)	(0.971)	(1.184)	(1.403)
Crude oil price (USD/barrel)	0.105***	0.221	0.103***	0.228		
	(0.015)	(0.180)	(0.015)	(0.169)		
Local operator × Crude oil price (USD/barrel)	-0.112***	-0.111***	-0.069***	-0.080***	-0.097***	-0.125***
	(0.023)	(0.030)	(0.018)	(0.020)	(0.020)	(0.026)
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	3497	3497	3497	3497	3497	3497
R^2	0.022	0.225	0.673	0.680	0.715	0.756

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 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-meaned annual average world price, in dollars per barrel. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

G.2 Heterogeneity: capacity for violence

For certain values of the model parameters, it may be the case that gangs are always worth bribing. Groups that can threaten violent retaliation – and not just theft – may fall into this category. Since κ may be interpreted as the additional damage that a gang can inflict on the oil company in the event that protection is not purchased, it is easy to see that rising κ shifts by γp the valid parameter range in A2 of Section 6. In particular, as κ rises, the lower bound on λ does as well, implying a violation of the assumption that some gangsters are not worth bribing. Thus, for large κ it may be that all gangsters are bribed for any value of ϵ_g, so that variation in bargaining frictions has no effect on the margin.

The key implication is therefore that we should expect to see the largest effects of localization on crime among relatively weaker groups. Two pieces of evidence support this prediction. Firstly, as Table A3 makes clear, the effect of local ownership on oil-related deaths is not constant across sub-categories of violence: the entire reduction is driven by events not attributable to an organized militant group. This is consistent with the interpretation that only violence by smaller gangs is affected by changing bargaining costs, since strong militant groups are always bribed.[70]

Secondly, the data show that the benefits of localization are concentrated on assets where nearby gangs have lower capacity for violence. Rexer and Hvinden (2020) show that in the Niger Delta conflict, militant groups with more allies connected locally along the pipeline

[70] The results are also consistent with Rexer and Hvinden (2020), who argue that more organized militant violence needs to be understood in the context of a bargaining interaction with the Federal Government rather than oil companies.

network have greater capacity for output destruction and receive more generous amnesty deals as a result.[71] Using the number of allied connections within 10 km along a pipeline as a proxy for destructive capacity among the nearest group, I find that treatment effects for both oil theft and violence are largest in areas with weaker nearby groups. In Table B17, I interact the main localization regression with this measure of the destructive capacity of the nearest militant group, measured in 2009.

The results in Panel A indicate that the theft reduction of localization is large and significant when the nearest militant group has no local allies. The coefficient then attenuates toward zero with each additional ally, indicating a militarily stronger group. This interaction term is significant in all specifications except the subsample of fields that are geographically distant from militant camps. These results are unaffected by the addition of controls in column (3). The onshore-offshore falsification test in columns (4) and (5) reveals the expected results: both the baseline reduction in theft and the heterogeneity coefficient are large and significant in the onshore sample and zero offshore. Table B17 Panel B shows that similar patterns obtain for the oil-related violence. Finally, note that this interaction effect is not driven by militarily strong groups simply being in areas with a greater density of groups.[72] Including controls for the number of militant camps within 10 kilometers interacted with localization in column (2) does not materially affect the results.

[71] See Appendix B and Figure B3 for a more thorough explanation.

[72] This might be the case if dealing with numerous nearby groups makes it harder to negotiate, coordinate, or enforce bargains)

Table B17: The effect of divestment on criminality by local military strength

Sample	All			Onshore	Offshore
	(1)	(2)	(3)	(4)	(5)
Panel A: Theft					
Local operator	-5.367***	-5.394**	-6.114***	-5.725***	-0.033
	(1.941)	(2.125)	(1.912)	(1.993)	(0.136)
Local operator × Allied camps along pipeline, 10 km	1.382**	1.219*	1.737**	2.024***	0.007
	(0.664)	(0.680)	(0.673)	(0.697)	(0.037)
Observations	3497	3497	3497	2518	979
R^2	0.736	0.740	0.761	0.752	0.442
Panel B: Oil-related violence					
Local operator	-0.879**	-1.249***	-0.837**	-1.394**	-0.069
	(0.393)	(0.402)	(0.396)	(0.553)	(0.081)
Local operator × Allied camps along pipeline, 10 km	0.417**	0.767***	0.406**	0.550**	-0.007
	(0.202)	(0.193)	(0.205)	(0.239)	(0.011)
Observations	3497	3497	3497	2518	979
R^2	0.188	0.200	0.202	0.252	0.124
Controls × Year FE	No	No	Yes	Yes	Yes
Number of camps	No	Yes	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes
Field FE	Yes	Yes	Yes	Yes	Yes

Standard errors, in brackets, are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016 in columns (1)-(3), the sample of onshore fields in (4) and the sample of offshore fields in (5). 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. All specifications include the number of non-allies within 10 km interacted with the local operator indicator. "Number of camps" refers to the inclusion of the number of camps within 10 km interacted with the local operator indicator. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.