



PEDL Research Papers

This research was partly or entirely supported by funding from the research initiative Private Enterprise Development in Low-Income Countries (PEDL), a Foreign, Commonwealth & Development Office (FCDO) funded programme run by the Centre for Economic Policy Research (CEPR).

This is a PEDL Research Paper which emanates from a PEDL funded project. Any views expressed here are those of the author(s) and not those of the programme nor of the affiliated organizations. Although research disseminated by PEDL may include views on policy, the programme itself takes no institutional policy positions.

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

Jonah M. Rexer[†]

Princeton University

August 25, 2022

Abstract

Multinationals in the extractive sectors of weak states face resource theft by armed groups. This criminality is often abetted by state corruption, even though firms are willing to pay for protection. I study indigenization in Nigeria's oil sector, which increased participation by Nigerian firms substantially. Despite evidence that local firms are lower quality, localization increases output and reduces oil theft. A bargaining model illustrates that political connections align law enforcement incentives, solving commitment problems. Data on raids by government forces show that local firms receive preferential law enforcement protection. I find that connections to military elites drive the local advantage.

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 García-Jimeno, Guy Grossman, Santosh Anagol, Arthur van Benthem, Nina Harari, Michael Watts, Jacob Shapiro, Nicholas Ryan, Alexey Makarin, Belinda Archibong, Bruno Baranek, Ameet Morjaria, and Bob Rijkers. 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, the NBER Political Economy Summer Institute, the Working Group on African Political Economy, the World Bank, the BGSE Summer Forum, and the BI Norwegian Business School. I also thank representatives at the Nigerian Bureau of Statistics, the Department of 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, the Centre for Economic Policy Research, and Princeton University. All errors are my own.

[†]School of Public and International Affairs, Princeton University. Email: jrexer@princeton.edu. <https://www.jonahrexer.com/>

1 Introduction

In many weakly-institutionalized states, natural resource extraction is plagued by black markets run by organized criminal and armed groups.¹ These rackets of violence, theft, and extortion impose substantial costs on resource extraction firms (Blair et al. 2022), who should be willing to pay government to protect their assets from predation. Instead, firms spend amply on private security, yet assets remain vulnerable, while criminal groups are often implicitly supported by corrupt law enforcement agencies. These challenges can force multinationals to exit difficult markets entirely (Burger et al. 2015). In 2021, oil supermajor Shell announced that it would leave Nigeria after 65 years as the country’s leading producer, with plans to divest more than \$2 billion in productive assets. This dramatic decision was a reaction to decades of steadily worsening losses from oil theft and widespread pipeline sabotage (Network 2021). Why does the state fail to protect these firms? Such inefficient outcomes suggest a glaring failure of Coasean bargaining (Acemoglu 2003).

In a limited commitment environment with a lucrative illegal sector, local law enforcement may have strong incentives to allow criminal activity.² The large multinational firms that have historically dominated many extractive industries have struggled to incentivize government to protect their assets. Indigenous firms, however, may have a comparative advantage in corruption (Javorcik and Wei 2009). By appointing political and military elites as shareholders and directors, local firms make state agents residual claimants on private resource income. In return, local firms can obtain protection from criminal predation, allowing them to outperform better-resourced multinational peers.

I study the benefits of localness in the context of the Nigerian petroleum sector, an industry fraught with political violence, corruption, and black market oil theft. Following a major policy reform, the share of onshore output produced by local firms grew from just 6.4% in 2006 to 40% in 2016. Using detailed panel data on Nigeria’s active oilfields, I leverage this wave of indigenization to study the effect of local participation on field performance, resource theft, violence, and law enforcement responses. I define a field as “localized” if it is owned or operated by any firm that is not a subsidiary of a foreign multinational.³

To identify the causal effect of local takeover, I employ a difference-in-differences approach that exploits changes in field ownership driven by multinational divestment. Using two-way fixed effects (TWFE), I find that local takeover increases output by roughly 33%. This

¹In Colombia, the gold trade is marred by paramilitary taxation and theft (Zabyelina and van Uhm 2020), while in the DRC, armed militias control substantial mineral extraction (de la Sierra 2020). The global black market for stolen crude oil is worth roughly 100 billion USD annually (Ralby 2018).

²Throughout history, black markets have often been accompanied by law enforcement corruption as state security agents use their position to appropriate illicit rents. For example, prohibition in the U.S. (Garcia-Jimeno 2016) or the current drug trade in Mexico.

³This relatively inclusive definition constitutes a somewhat weak treatment, since even small equity stakes acquired by local firms are considered localizations. As a result, the empirical estimates are likely conservative.

“local advantage” occurs despite local firms exhibiting lower quality. Divested fields experience 10% more oil spills due to mechanical failure, and flare 25% more gas annually, reflecting lower operating standards, management, and maintenance of physical infrastructure.

Local firms increase output despite lower quality by mitigating political risks: local takeover leads to a sustained reduction in incidents of oil theft and violence. On average, localized fields experience 5.7 fewer theft incidents, a 53% decline, and 3 fewer conflict fatalities annually. As a falsification test, I examine heterogeneous effects by asset type. I find that the local advantage in output, theft, and violence is concentrated entirely onshore, consistent with greater political risk exposure. Instead, local firms’ quality disadvantage is more pronounced offshore, where technology and capital requirements are higher.

Divestments to local firms may be correlated with unobserved trends in oilfield quality. Event-study estimates of dynamic effects do not reveal any pre-trends for the main outcomes. I further probe the identification assumption using detailed data on corporate transactions. I exploit the fact that in Nigeria, a weak legal framework implies total ministerial discretion over oil and gas transactions, leading many planned divestments to be stalled or terminated. I show that these terminated transactions do not produce any effects. Since these fields are ostensibly subject to similar unobserved trends and selection effects as successful divestments, this placebo test allays concerns that such biases are driving the results.

The main results are robust to a wide variety of interacted fixed effects, modeling choices, control strategies, measurement assumptions, estimators, sample restrictions, and methods of inference. I run diagnostic tests on the TWFE specifications from Goodman-Bacon (2021), de Chaisemartin and D’Haultfoeuille (2020), and Callaway and Sant’Anna (2021) to assess the extent of bias induced by treatment effect heterogeneity. Decompositions suggest that relying solely on treatment timing produces similar results, while nonparametric and stacked estimators confirm that the results of the TWFE are, if anything, somewhat conservative.

Criminal groups in the sector operate under the protection of the state security forces (SDN 2019a). 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 develop a simple model in which a firm bargains with the state to obtain law enforcement, where purely self-interested government forces maintain the outside option of collecting criminal bribes and allowing theft. Corruption is “efficient” when it incentivizes law enforcement provision, maximizing social surplus.

In the base case with perfect commitment, the bargaining range is eroded by several frictions, such as the state’s cost of enforcement and the firm’s liability cost of making illicit payments. The strength of the state’s monopoly of violence and the size of theft losses drive firms’ willingness to pay bribes. Political connections provide law enforcement with a claim to residual oil income, thereby aligning incentives between the firm and the state, and lowering the threshold bribe that a connected firm must pay. Connected firms therefore experi-

ence more enforcement and less crime, all else equal. With commitment problems, lump-sum bribes are rendered ineffective, since the security forces may renege on any deal. Only private claims to oil surplus can incentivize state agents to fight theft. Political connections are not just valuable, but are indeed 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 appoint influential political and military elites as shareholders and board members in exchange for asset protection.

Using data on law enforcement activity collected from Nigerian news media, I show that local firms indeed receive greater protection. Local participation increases state law enforcement actions against oil theft by 83%, affecting the entire black-market value chain, from illegal refining for domestic consumption to illicit export activity. Placebo tests show null effects for non-oil crime, ruling out unobserved differential trends in policing in divested areas.

I investigate the role of political connections using hand-collected data on the biographies and tenures of all shareholders and board members in Nigeria's oil sector. Descriptive statistics reveal that while overall rates of political connectedness are similar across firms types, local firms are substantially more likely to have political elites with military or police experience among their officers. Using a fixed effects approach, I show that field-level control by politically connected firms is significantly associated with lower oil theft *only* for connections to the armed forces. Connections to regulatory technocrats, traditional leaders, and elected politicians have null effects. Security connections are additionally associated with greater output and more enforcement activity. 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.

Lastly, I find suggestive evidence in favor of additional mechanisms contained in the model. Among multinational firms, exposure to a foreign corruption law increases field-level theft and violence, suggesting that part of the local advantage stems from the lower legal costs of corruption. In addition, I use data on oil licenses to show that indigenization increases the equity share of the operating firm by 17%. With larger ownership stakes, local operators internalize a greater share of theft losses, increasing willingness to pay for protection. Lastly, I consider several alternative explanations for the results – crime and employment spillovers, differences in discount rates, targeted host-community investments, and local grievance toward multinationals – and find no support for any of these alternative mechanisms.

A long-running literature in economics, stretching at least as far back as Hirschman (1969), debates the costs and benefits of multinational investment in developing economies. Multinationals are generally more productive, owing to better human capital, technology, and management practices (Aitken and Harrison 1999, Arnold and Javorcik 2009, Criscuolo and Martin 2009, Guadalupe et al. 2012).⁴ Despite these advantages, it has long been noted that

⁴They also raise aggregate productivity by transferring technology (Teece 1977, Guadalupe et al. 2012) and

so-called “institutional voids” in developing countries may frustrate multinational operations (Palepu and Khanna 1998) and deter investment (Burger et al. 2015, Blair et al. 2022). Evidence suggests that corruption encourages joint ventures as multinationals seek partners to navigate local politics (Javorcik and Wei 2009), and that anti-bribery laws put multinationals at a disadvantage (Chapman et al. 2021).⁵ Nevertheless, evidence on direct comparisons of multinational and local performance in conflict-affected markets is limited. I show that in these markets, local advantage may outweigh the productivity gains from foreign investment. Still, these benefits must be balanced against the welfare costs of increased environmental pollution. This paper therefore illuminates new and important tradeoffs in an old debate.

The results also add to a large empirical literature that studies the relationship between natural resources, criminal economies, and violence. This literature has looked at the effects of natural resource booms on violent conflict (Berman et al. 2017, de la Sierra 2020, Dube and Vargas 2013, Fetzer and Kyburz 2018, Nwokolo 2018) and social unrest (Couttenier et al. 2017, Sexton 2019, Christensen 2019). I join a small but growing body of evidence that demonstrates the centrality of black markets and organized crime to resource curse dynamics (Buonanno et al. 2015, Chimeli and Soares 2017, Gehring et al. 2020). I extend this work in two ways: by analyzing firms as strategic participants, and by modeling endogenous law enforcement. To the best of my knowledge, I am the first to show that local resource ownership can improve law enforcement provision and mitigate the most violent pathologies of the resource curse.

It is well known that political connections are valuable to firms operating in corrupt markets (Fisman 2001, Khwaja and Mian 2005, Faccio 2006, H. Li et al. 2008). However, these studies typically emphasize the negative equilibrium effects of political favoritism: inefficient firms are protected from competition (Akcigit et al. 2018) and charge higher prices (Baranek and Titl 2020). Instead, I emphasize the aggregate benefits of political connections, which mitigate the inefficiencies of black markets and increase output. The results demonstrate a novel mechanism linking political connections to firm outcomes: corrupt networks allow firms to navigate state-run protection rackets (Beckert and Dewey 2017), incentivizing the provision of government law enforcement to combat criminal activity.

More broadly, the results speak to a longstanding debate on the efficiency costs of corruption.⁶ In the spirit of Shleifer and Vishny (1993), I argue that when institutional constraints are limited, the structure of corruption is what determines its efficiency costs. In the Nigerian case, divestment to local firms can force government to internalize some of the costs of illegality, yielding a second-best equilibrium (Lipsey and Lancaster 1956) of efficient corruption.

skills (Bloom and van Reenen 2010, Bloom, Sadun, et al. 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 Harrison and Rodriguez-Clare (2010) and Alfaro and Chauvin (2020) for reviews.

⁵Although Guidolin and Ferrara (2007) show that conflict can be beneficial to multinational firms.

⁶See e.g. Huntington (1968), Bardhan (1997), Kaufmann et al. (1999), and Méon and Weill (2010).

2 Background

2.1 The Nigerian oil sector

Nigeria is the world's 11th largest oil-producer, and the largest in Africa. Rich deposits of crude oil are located onshore and in the waters of the Niger Delta, a region in southern Nigeria where the mouth of the Niger River meets the Gulf of Guinea. The Niger Delta comprises both coastal and inland portions of nine states,⁷ 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, ExxonMobil, Total, and Eni. In 2004, just before the sample period, 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 substantial variation in the share of profits claimed by the operator of a given oilfield.⁸ 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 have a fraught relationship with host communities. Oil spills and gas flares are common, affecting soil, fisheries, and drinking water, and increasing infant mortality (Bruederle and Hodler 2019). For years, myriad armed groups have interacted 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 (Watts 2007, Obi and Rustad 2011). Around 2000, The Niger Delta Crisis saw the emergence of well-armed militant groups capable of launching devastating attacks on the federal government and oil companies. These militants destroyed oil infrastructure and kidnapped staff in an attempt to obtain concessions for themselves and the region (Watts 2007, Asuni 2009). In 2009, the Federal Government provided amnesty to nearly 25,000 combatants, as well as private transfers to top militants in the form of lucrative "pipeline surveillance contracts" (SDN 2019c).

In the aftermath of conflict, an industrial-scale black market in stolen oil siphoned directly from onshore pipelines has flourished (Rexer and Hvinden 2022). Appendix Figure

⁷These are Abia, Bayelsa, Delta, Rivers, Akwa Ibom, Imo, Ondo, Edo, and Cross River states.

⁸Online Appendix Figure A4 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%.

⁹See [this report by the Carnegie Endowment](#) for a survey of corruption issues in Nigeria.

A5¹⁰ charts the growth of the black market, plotting monthly incidents of pipeline sabotage as a proxy measure. In 2016, the black market totaled 4.2 billion dollars, or 15% of Nigeria's total production (NEITI 2016). In this two-tiered market, small-scale 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 Nigerian Federal security forces – the military and police – play an important role in facilitating the smooth functioning of the black market by allowing criminal activities in exchange for kickbacks (SDN 2019a). Protection rackets naturally arise: oil companies must negotiate the demands of gangsters, local communities, and law enforcement in order to safeguard output.

2.2 Offshoring and divestment

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 exploration and production costs for firms. Between 2002 and 2016, the share of Nigerian oil produced from onshore fields fell by half, from 60% to just above 30%. As multinationals face oil theft losses, ESG risks, recurrent militancy, and ballooning private security costs (Network 2021), they have often opted to divest their holdings in the onshore market, or in some cases leave Nigeria entirely. Most strikingly, in 2021, Shell – long Nigeria's largest operator – recently announced plans to divest all of its remaining Nigeria holdings.¹¹

At the same time, the Nigerian government has sought to ensure that the assets of departing multinationals are bought by indigenous Nigerian firms. In 2010, a major policy reform – the Nigerian Local Content Act – legally enshrined a preference for local firms in bidding on both new and divested oil blocks. These underlying trends and policy shocks have led to a dramatic shift in the ownership structure of the upstream market. Figure 1 plots participation by Nigerian oil and gas firms by asset type over time. From 2006-2016, the number of fields with local participation rose fivefold, from 19 to 94. This growth in local participation is concentrated primarily in onshore assets, where the local market share has grown from 6% to 40% over this period. The dotted vertical line in Figure 1 demonstrates that the timing of the law correlates with the upsurge in local onshore participation.¹²

Until 2021, Nigeria's petroleum sector lacked a comprehensive regulatory framework to guide the divestment process, due to longstanding legislative gridlock in passing a Petroleum Industry Bill.¹³ As a result, the divestment process has been governed by the principle of

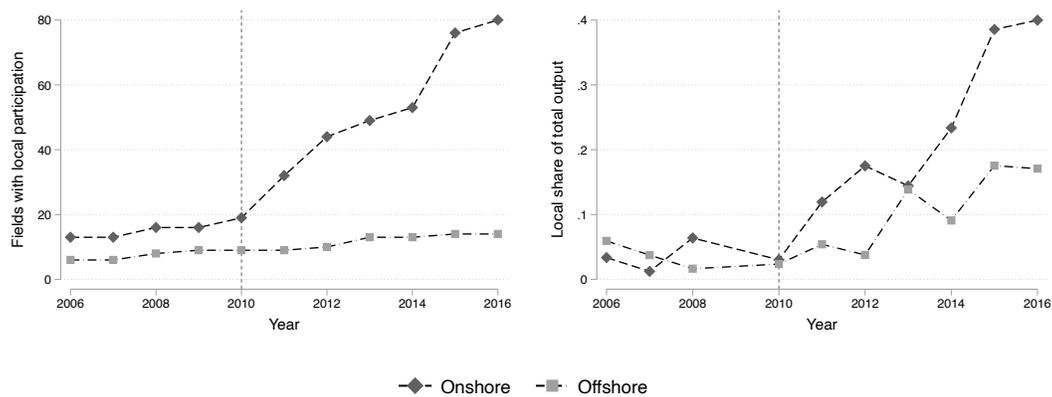
¹⁰Reprinted from Rexer and Hvinden (2022).

¹¹As reported by market intelligence firm Wood Mackenzie.

¹²Local firms have also gained footholds by bidding for the rights to "marginal fields", a category of smaller, undeveloped fields in larger oilblocks reserved exclusively for local companies.

¹³The PIB was passed in 2021, after a 20-year legislative process. [Here](#), KPMG reports on the political process and provisions.

Figure 1: Indigenization over time



Note: Figure shows number of fields (left panel) and output share (right panel) 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 2006-2016. Oil production data and output shares are missing for 2009.

“ministerial consent;” the ultimate authority for approval of any transfer of asset ownership or operatorship lies with the Minister of Petroleum Resources. A set of guidelines released in 2014 by the Department of Petroleum Resources details the process: firms submit proposals to the minister with a list of buyers, assets, bids, and technical details. There is no requirement for competitive bidding. The Minister then evaluates proposed sales on two primary criteria: *i)* the buyers’ capabilities to acquire and/or operate the stake, including access to financing and technology, and *ii)* whether the buyer is an indigenous firm and therefore given first consideration.¹⁴ Failure to properly obtain ministerial consent ex-ante will result in a terminated transaction. In addition, the Minister may reject sales at his or her own discretion.

These rules have three implications for the empirical analysis. First, due to local content, nearly all asset divestments since 2010 have been from multinational to local firms, driving the pattern in Figure 1 and generating a large-scale policy experiment that I leverage to estimate the impact of localization. Second, the primacy of ministerial consent implies that firms have limited control over divestment timing or approval. Third, there are numerous stalled and terminated transactions, which can serve as falsification tests of this policy experiment.

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, the construction of key

¹⁴In cases where the NNPC is a majority owner, then additional consent is required to transfer operatorship. For more details, see [this legal report](#) by the global law firm Chambers and Partners.

variables, and measurement error issues, please refer to Appendix [A](#).

3.1 Oil sector data

3.1.1 Field output and location

An annual panel of 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), covering 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 firm. A field enters the dataset in the year it first appears in these administrative records, and remains thereafter even if output is missing in a subsequent year, creating an unbalanced panel. 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 other outcomes and geospatial control variables using centroid coordinates. I discuss extensive output measurement error checks in Appendix [A.1](#).

3.1.2 Field-level outcomes

Data on oil theft comes from the Nigerian Oil Spill Detection and Response Agency (NOS-DRA), a division of the Federal Ministry of the Environment. I obtain 11,587 reported oil spills from 2006-2017, 68.45 % of which are classified as being caused by “third-party sabotage.” I take this to be my sample of oil theft incidents, since sabotage is a reliable indicator of illegal oil tapping.¹⁵ To measure oil spills caused by regular production activities, I use all field-level spills that are not due to sabotage.¹⁶ Since fields encompass many different types of oil infrastructure (pipelines, wells, flow stations, etc) over large and often heterogeneous areas, I define outcomes as the count of theft incidents or oil spills within 15 kilometers of the field centroid.¹⁷ Appendix Figure [A5](#) shows oil spills and theft over time.

Data on gas flaring from 2012-2020 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).¹⁸ Gas flaring is measured at the field-year in thousand cubic feet (mscf). Field-years in this period for which no flaring is observed are set to zero.

Data on conflict and violence comes from news media sources aggregated by the Armed Conflict Location and Event Dataset (ACLED), covering 1998-2016. I measure conflict intensity using total annual fatalities from events within 15 kilometers of the field centroid. I also

¹⁵See Rexer and Hvinden (2022) for a discussion about measuring oil theft.

¹⁶These are typically caused by equipment malfunction and other unknown causes.

¹⁷I test robustness to this arbitrary threshold in Appendix [C.7](#).

¹⁸For greater detail on the remote sensing methodology, see the [Gas Flare Tracker website](#).

measure violent oil-related activity, specifically, using all conflict events in which the description mentions a set of key words about the oil industry. In some specifications, I disaggregate by media source (local vs. international press) to mitigate reporting biases (Appendix E.2).

3.1.3 Indigenization

I define treatment as a binary indicator of any participation in a field – ownership or operatorship – by an indigenous Nigerian firm. An indigenous Nigerian firm must be headquartered in Nigeria and is not majority-owned by a firm headquartered outside of Nigeria.¹⁹ I combine multiple data sources to precisely identify the year in which a field is “indigenized.”

The DPR-NNPC dataset includes information on the firm operating each field in each year. However, this data has substantial gaps – many fields are missing information for years after they first enter the data. In addition, operatorship information is likely to lag divestment transactions. Furthermore, using the operatorship measure exclusively overlooks cases in which local firms are non-operating shareholders, which may also be important. Therefore, it is difficult from this data alone to determine the exact year in which a given treatment occurs.

From DrillingInfo (DI) – a paid-subscription database on the oil and gas sector – I obtain a list of 117 physical asset transactions in the Nigerian oil and gas sector from 2006-2020. The DI data provides detailed information on a substantially wider set of transactions, covering all cases in which any ownership stake in a given field is transferred from a multinational to a local firm with precise information on transaction timing. However, it does not have information before 2006, so we cannot identify the always-treated fields using only the DI data. Furthermore, it does not distinguish ownership and operatorship.

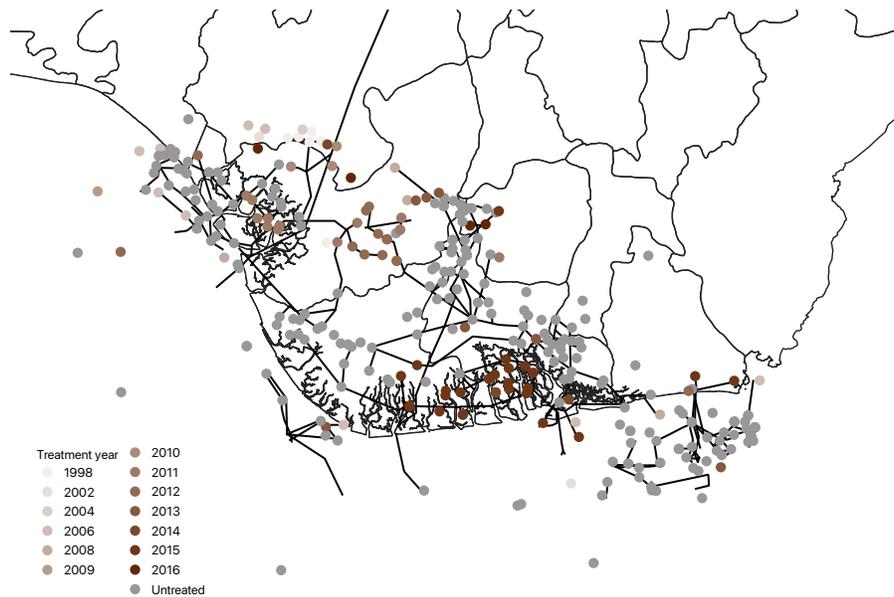
However, when combined, the longer panel of the NNPC data can fill in the always-treated fields, while the DI data can provide more precise timing and identify non-operating divestments from 2006-2016. A detailed breakdown of the treatment definition is provided in Appendix Table A1, which counts the treatment observations by data source. I take the earliest year of treatment across sources for a given field; the final treatment indicator $local_{it}$ is equal to one for all field-years after the treatment year, following the “staggered adoption” difference-in-differences setup. If there are reverse divestments from locals to multinationals, we may wrongly assign treated field-years and therefore create measurement error biasing the effect toward zero.²⁰ In total, there are 94 ever-treated fields.

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 treatment cohort. The 220 untreated fields are clustered in the tidal mangroves of Delta,

¹⁹In practice, the local subsidiaries of oil supermajors in Nigeria – Eni, Total, Shell, Chevron, and ExxonMobil, as well as the Chinese Sinopec subsidiary Addax Petroleum – are classified as MNC; all others are indigenous. The NNPC-owned Nigerian Petroleum Development Company (NPDC) is considered an indigenous firm.

²⁰These transactions are exceedingly rare in the DI data, affecting less than 1% of field-years from 2006-2016.

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. Grey markers are never-treated fields. Basemap is Nigerian states of the Niger Delta region, while lines indicate oil pipelines.

Bayelsa, and Rivers states – the heart of Nigeria’s oil sector – as well as in the shallow waters off Akwa Ibom state. The brown points indicate the 94 local fields and their takeover dates. Indigenized fields are clustered partly in the inland Niger Delta, with another cluster of recently-divested fields in coastal Rivers state and a handful of offshore assets.

3.2 Law enforcement data

Data on law enforcement comes from the text of news media reports from Nigerian and international publications on raids, seizures, arrests, and other oil theft-related activity by state forces. I first assemble a comprehensive sample of crime-related news articles, and then employ local research assistants to identify relevant articles and extract law enforcement events. From this procedure, we obtain 3261 unique geocoded law enforcement events from 2000-2020 related to the illegal oil sector. For each event, I observe all involved law enforcement agencies, the criminal activity,²¹ the items destroyed or seized,²² the number of arrests, and the number of fatalities. I collapse these to the field-year level by taking all enforcement actions within fifteen kilometers of the field centroid. The trend in enforcement over time is shown in Appendix Figure A7.

²¹Some prominent examples: transportation of stolen oil, piracy, pipeline vandalism.

²²Some examples are guns, illegal refineries, stolen oil, and boats.

3.3 Political connections data

To collect data on the political connections of firms in the Nigerian oil sector, I partnered with a market research firm called [Asoko Insight](#), specializing in research on corporate entities in sub-Saharan Africa. In order to obtain a measure of firm-level political connections over time, we first collated a comprehensive list of board members, shareholders, and second-level beneficial owners from corporate filings with the Nigerian Corporate Affairs Commission (CAC), as well as all officers' tenures and individual holdings, from each firm's incorporation date to the present. We then collected comprehensive biographical information on these individuals, detailing their career progression, political participation, and other relevant information. In total, we identify 706 Nigerian nationals across 49 firms; we obtain biographical information for 552 of these, or 78.2%.²³ Details on the data collection process and coverage rates by firm are in Appendix A.9. I use biographies to identify field-years in which the operating firm has any board member or shareholder that has ever served at any level of Nigerian government. I also identify connections to technocratic regulatory agencies, cabinet-level politicians, elected politicians, traditional leaders, and members of the military and police.

3.4 Summary statistics

I combine all the data into an unbalanced panel from 2006-2016, consisting of 3183 field-years. Summary statistics are presented in Table 1, comparing ever-treated and untreated fields. In the top panel, I consider time-invariant field-level characteristics. Fields are not significantly different in their distance from the coast, the state capital, or from militant camps. They are of a similar age in 2016,²⁴ on average 42 years old. 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 significantly greater latitude, since divested are more likely to be onshore and in the inland Niger Delta. Indeed, 85% of ever-treated fields are onshore, while only 65% of multinational ones are, a difference that is significant at 1%. Treated fields are also slightly smaller, with 6.7 fewer wells per field, significant at 5%. 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 in the bottom panel of Table 1. These comparisons use all field-years and therefore include before and after periods for the treated group. Treated fields experience less oil theft, but more conflict. There are no differences in shut-in (nonproducing) rates, but annual production is on average 1.39 million barrels (44%) higher among the never-treated, and annual oil spills are also higher. In order to determine which of these relationships are causal, and control for the differences across time-invariant covariates,

²³We additionally find biographies on 90 second-level beneficial owners.

²⁴Defined relative to the date of completion of the first oil well.

I move to a staggered-adoption differences-in-differences approach in Section 4.

Table 1: Summary statistics

	Untreated	Treated	Full Sample
<i>Covariates</i>			
Field latitude	4.90 (0.60)	5.23 (0.66)	5.00 (0.64)
Distance to coast (km)	34.08 (29.13)	29.39 (30.59)	32.67 (29.60)
Distance to Niger River (km)	90.20 (79.74)	63.07 (55.81)	82.08 (74.36)
Distance to state capital (km)	87.02 (49.72)	80.27 (53.25)	85.00 (50.81)
Distance to militant camp (km)	32.64 (25.40)	33.61 (30.00)	32.93 (26.81)
Number of wells	20.98 (34.09)	14.29 (18.62)	18.94 (30.35)
Field age (2016)	42.07 (12.00)	41.71 (12.61)	41.96 (12.17)
Onshore field	0.65 (0.48)	0.85 (0.36)	0.71 (0.46)
Max well depth (m)	2602.36 (870.25)	2821.66 (1008.06)	2669.16 (918.31)
<i>Outcomes</i>			
Oil theft events, 15 km	11.02 (21.72)	5.64 (11.39)	9.51 (19.55)
Total conflict deaths, 15km	1.80 (7.15)	2.17 (9.45)	1.90 (7.86)
Piracy attacks, 15 km	0.12 (0.67)	0.19 (0.96)	0.14 (0.76)
Shut-in field	0.22 (0.42)	0.26 (0.44)	0.23 (0.42)
Annual oil production (million barrels)	3.13 (8.06)	1.74 (4.28)	2.76 (7.28)
Oil spills, 15 km	7.93 (10.49)	3.73 (5.93)	6.75 (9.62)
Number of clusters	220	94	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 over the sample period.

4 Empirical strategy

To test whether local firms affect field-level outcomes, 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 local participation, and ψ is the parameter of interest. Fixed effects for year δ_t and field ξ_i complete the two-way fixed effects (TWFE) specification of the DD model, while X_{it} includes an additional vector of time-invariant covariates interacted with year dummies. I use a parsimonious set of controls that includes the field latitude, as well as distance to the state capital, the nearest river, and the coast. Standard errors are clustered at the field level. The primary outcomes are output, oil spills, and oil theft, though I also consider a range of additional field-level outcomes as well.

In a TWFE specification, variation in $local_{it}$ is driven by changes in participation within a field over time, holding common shocks and time-invariant characteristics fixed. This means that fixed differences in the age, size, or productivity of fields allocated to local firms are controlled for; only trends in outcomes correlated with ownership changes should contaminate the results. Local takeovers might occur when oil prices are low, or following a deterioration in output and theft trends on a given asset. Localization could also be spatially and temporally correlated with specific policy changes that independently influence outcomes. As a standard omnibus test of pre-trends, I estimate the “stacked” event-study specification for field i in treatment cohort c at calendar year t

$$y_{ict} = \alpha + \sum_{\tau=-T}^T \psi_{\tau} 1(t - c = \tau) * local_{ic} + \delta_{ct} + \gamma_{ic} + \epsilon_{ict}$$

The stacked model takes all fields treated in year c , and pairs them with control units that are never or not-yet treated as of year c , keeping field-years in an event-window around c . I use $\tau \in [-7, 10]$ for the main specification. I then stack each cohort together, and estimate the event-study on the stacked dataset. The model includes relative time τ dummies $1(t - c = \tau)$, interacted with an indicator that field i is treated in year c , $local_{ic}$, to identify dynamic effects ψ_{τ} . Finally, I include cohort-specific year and unit effects δ_{ct} and γ_{ic} to saturate the model. In staggered adoption settings, TWFE may produce biased estimates of the DD effect in the presence of heterogeneous or dynamic treatment effects. This phenomenon is primarily driven by the fact that early- and always-treated units are implicitly used as controls for later-treated units, even though their previous treatment status is likely to alter their subsequent trends. The stacked model has the benefit of excluding these “bad” control units explicitly. In Appendix D, I consider recent advances in the literature on DD estimation from de Chaisemartin and D’Haultfoeuille (2020), Goodman-Bacon (2021), Callaway and Sant’Anna (2021), and Baker et al. (2021), testing robustness of the main results and the event-studies to a wide array of diagnostic tests, new estimators, reweighting schemes, and specification choices.

Unconditional variation in local participation comes from three sources. First, fields are divested from multinational to local firms in asset sales or operator agreements. 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 are

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 transitions in 1), while 2) and 3) comprise always-treated units, and therefore serve as control groups in the TWFE model.²⁵

The critical question for identification, then, is whether multinational-to-local divestments can be considered exogenous, conditional on the fixed effects. Of course, the choice of divested assets is not random, as Table 1 shows. Differential trends across observable field-level characteristics can be flexibly controlled for in X_{it} . But the main concern is that differential trends on unobservables will bias the results. Intuitively, such trends should bias the results against local advantage, since incumbent multinationals possess inside information on trends in field quality, leading to adverse selection on trends. Furthermore, such confounders would likely show up as significant pre-trends in event-studies.

Still, flat pre-trends 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. As reported in Section 2.2, asset sales in Nigeria’s oil sector are governed by ministerial consent; in the recent wave of divestment, many transactions were stalled or terminated by the Ministry of Petroleum Resources. As such, the precise timing of local takeover is unlikely to be systematically correlated with unobserved field trends, since it is not directly manipulable by market participants. In Section 5.3.3, 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. In Appendix D, I also show that treatment effects relying *only* on timing variation are if anything larger than the main specification.

5 Main results

5.1 Main outcomes

The main results of the TWFE models are in Table 2. Columns (1)-(2) use output in millions of barrels as the outcome, (3)-(4) use oil spills, and (5)-(6) use oil theft. For each outcome, I estimate the TWFE model with and without controls. All results indicate the control group mean for reference. I find that local participation increases per-field output by roughly 0.89-0.96 million barrels on average annually, significant at 1% in both specifications, corresponding to 30.5-33.2% of the control group mean. The increase in output translates to 36-82 million dollars in additional revenue per year at world oil prices (see Appendix Table A14).

In Table 2 columns (3)-(4), I estimate the effect of local participation on equipment malfunctions that result in oil spillage. Local fields experience 0.79-1.38 more spills annually, or

²⁵Again, in OA D I consider these issues exhaustively.

Table 2: The effect of divestment on output, oil spills, and crime

Outcome	Output		Oil spills		Oil theft	
	(1)	(2)	(3)	(4)	(5)	(6)
Local firm	0.887*** (0.312)	0.963*** (0.308)	1.380* (0.767)	0.788 (0.828)	-4.805*** (0.804)	-5.703*** (1.023)
Control group mean	2.901		7.612		10.608	
Observations	2476	2476	3183	3183	3183	3183
R ²	0.861	0.878	0.590	0.649	0.712	0.756
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 parentheses 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. Oil spills 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. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

10.4-18.1% of the control group mean. These estimates are not significant at conventional levels. However, in Appendix Tables A22 and A23 I consider alternative estimation methods that correct for TWFE biases. These estimates suggest that TWFE *understates* the DD effect in our setting. Under these alternative estimators, the effects on oil spills are positive and significant. I interpret these effects as evidence that local firms are of lower quality, though the precise source of this quality gap is unobserved.²⁶ I also find that the effect on malfunctions is not mechanically driven by the effect of greater oil production.²⁷

What drives this gap in performance? In Table 2 columns (5)-(6), I find that local participation reduces oil theft. Localized fields experience 4.8-5.7 fewer theft incidents annually, or 45.3-53.8% of the control group mean, significant at 1%. In Appendix Tables A19 and A20, I use an instrumental variables approach to estimate that, quantitatively, the reduction in oil theft in (6) is large enough to explain roughly 63% of the increase in output in (2).

5.2 Other outcomes

In Table 3, I consider the impact of localization on other field-level outcomes. Columns (1)-(2) show that localization significantly reduces the overall level of conflict-related violence on

²⁶Operational oil spills could be driven by lower-quality physical capital, human capital, or management practices and standards.

²⁷In Appendix Table A17, 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 are only slightly smaller in magnitude.

oil assets by roughly 3 events per year on average.²⁸ In columns (3)-(4), the effect on maritime piracy is also large and negative but not significant at conventional levels.²⁹

Table 3: The effect of divestment, other outcomes

Outcome	Conflict deaths		Piracy		Shut-in		Gas flaring	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Local firm	-3.043*** (1.052)	-3.203*** (0.993)	-0.111* (0.063)	-0.097 (0.065)	-0.009 (0.056)	0.003 (0.057)	0.415** (0.187)	0.271 (0.195)
Control group mean	2.006		0.154		0.234		1.068	
Observations	3183	3183	3183	3183	2476	2476	1503	1503
R ²	0.232	0.317	0.227	0.311	0.657	0.671	0.895	0.899
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

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. Conflict deaths are the total number of conflict-related fatalities reported in news media within 15 km of the field. Piracy is the pirate attacks within 15 km of the field. Shut-in is an indicator for nonzero production in a field-year. Gas flaring is measured in million mscf. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Theft and violence may reduce output by *i*) generating technical losses on the intensive margin from spilled and stolen oil, and *ii*) affecting the incentive on the extensive margin to produce at all. Columns (5)-(6) test the extensive margin, estimating the impact of divestment on shut-ins. There is no evidence that divestment affects extensive margin production; local firms are not so much reviving shut-in fields as they are recovering lost production from active fields.³⁰ However, in Appendix Table A25, I find that takeovers in which local firms also assume operatorship lead to a large and significant reduction in the shut-in rate.

In addition to oil spills, gas flaring represents a major source of environmental pollution from oil production in the Niger Delta. In columns (7)-(8), I estimate the effect of localization on flaring in the panel of oilfields from 2012-2016, the years of our sample for which gas flaring data is available. Local participation increases gas flaring by 0.27-0.41 million mscf on average, 35-39% of the control group mean. This results in an additional 14.3-36.2 thousand tonnes of CO₂ emissions per field annually. However, these effects are only significant without control variables. The gas flaring data support the results of Table 2 in confirming that local firms are indeed more prolific polluters.

²⁸ Appendix Table A27 and Figure A21 consider the role of measurement error in conflict deaths.

²⁹ Piracy attacks primarily consist of hijackings of oil tankers and platforms. In additional tests, not reported, I find that the reduction in piracy is significant in the subsample of offshore assets where piracy is concentrated.

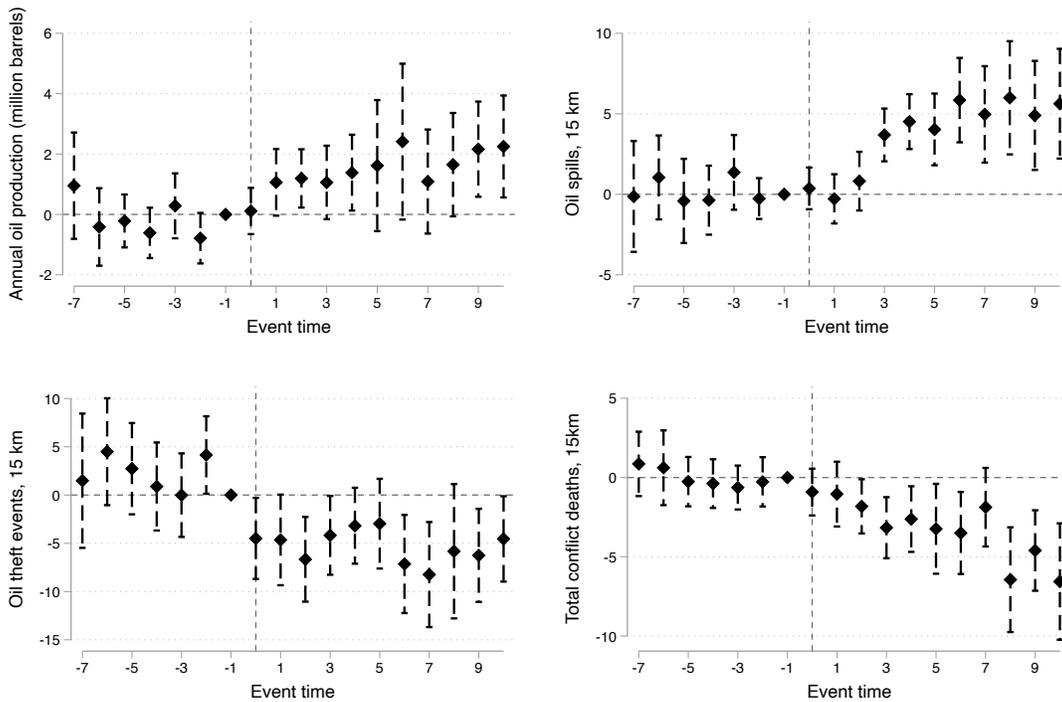
³⁰ This may be driven by the fact that long-standing shut-in fields may not be attractive targets for divestment. In fact, only 8 fields are shut-in at the time of their divestment. Furthermore, fields may shut-in during part of a year and still record positive production. Therefore, we cannot rule out within-year extensive margin effects.

5.3 Falsification and robustness tests

5.3.1 Parallel trends

I test for divergent pre-trends using a stacked event-study model, estimating the model for output, oil spills, oil theft, and violence. For each regression, I omit $\tau = -1$ as the pre-event reference year, and estimate the specification including interacted field-level controls for an event window of $\tau \in [-7, 10]$. Figure 3 presents the results. Overall, the pre-trends for all outcomes appear relatively parallel across treated and control fields, with no evidence of anticipation effects. For output (top left), the ψ_τ for $\tau < 0$ are all zero and insignificant. The event-study coefficients jump significantly in the year after divestment, remaining positive up to 10 years after the treatment. Oil spills, in the top-right panel, exhibit a similar trend.

Figure 3: Stacked-DD event-study: main outcomes



Note: Figure shows coefficients from stacked event-study regressions described in Section D.3 for oil production, oil theft, oil spills, and conflict. Standard errors are clustered at the field-level. Output is measured in millions of barrels of oil per year. Oil spills 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. Conflict deaths are the total number of conflict-related fatalities reported in news media within 15 km of the field. Specifications include interacted controls for latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital. Bottom-right panel shows the total number of treated fields in each event-time period.

For both theft and violence, the event-study coefficients are generally near zero in the

pre-period. The post-event coefficients for oil theft incidents (bottom left) are negative in all of the post-treatment periods, and significant at 5% in all but three. They display a large initial drop, followed by a consistent and sustained reduction in oil theft for 10 years, with minimal dynamic effects. In contrast, violence outcomes (bottom right) witness a small initial drop and take substantially longer to improve, with a sustained impact only emerging 2-3 years after divestment and growing thereafter. Appendix D.5 provides extensive robustness tests on the event-study models, considering alternative estimators, covariate specifications, control groups, and outcomes.

5.3.2 Asset type

Offshore assets have higher technology requirements and equipment costs but are less susceptible to direct theft and violence. Onshore assets, with their accessible, unprotected pipelines, are a soft target for oil theft, but also operationally low-cost for the firm. We should therefore expect reductions in criminality primarily on onshore assets. If oil theft drives the local output advantage, then we should further expect that output gains are also concentrated onshore. In contrast, given the greater technological requirements of offshore extraction, the multinational operational advantage should be concentrated in more complex offshore assets.

Table 4 replicates Table 2, splitting the sample into onshore (Panel A) and offshore (Panel B) fields. In columns (1)-(2), the output effect concentrated entirely onshore, while offshore it is small and insignificant. As predicted, patterns of heterogeneity in oil theft effects in columns (5)-(6) mirror those of output; the effect of localization on theft in is entirely concentrated in onshore fields, where the coefficient ranges from 6.5-6.8, significant at 1%. The political risks of onshore extraction give rise to a comparative advantage for local firms. Columns (2)-(3) confirm that local firms cause substantially more oil spills in offshore sites. Local takeover of an onshore field increases malfunctions by only 0.6-1.2, insignificant, while for offshore fields this number rises to 2.9-4.5, significant at 1%. The technological requirements of offshore extraction magnify the efficiency costs of local ownership. The bottom row shows that these differences are significant for oil spills and theft, though not output.

5.3.3 Transitions, and terminations

The DI data also allow for two useful falsification tests of our identification assumptions.³¹ First, the results may be driven by “transition effects,” whereby the observed effects are not driven by localization per se, but rather by *any* new owner. I use transactions that do not change assets’ local status to rule out these spurious effects. Second, for 36 fields, multinational-to-local divestments were planned but either delayed or terminated for bu-

³¹For a detailed description of the DI data, see Appendix A.2.

Table 4: The effect of divestment by asset type

Outcome	Output		Oil spills		Oil theft	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Onshore assets</i>						
Local firm	0.713*** (0.274)	0.769** (0.314)	1.185 (0.787)	0.643 (0.798)	-6.789*** (0.989)	-6.461*** (1.165)
Control group mean	1.475		7.580		14.977	
Observations	1736	1736	2296	2296	2296	2296
R ²	0.785	0.799	0.631	0.712	0.708	0.747
<i>Panel B: Offshore assets</i>						
Local firm	0.163 (0.715)	-0.093 (1.401)	2.918*** (0.843)	4.500*** (1.413)	0.048 (0.030)	0.071 (0.094)
Control group mean	5.967		7.689		0.048	
Observations	740	740	887	887	887	887
R ²	0.860	0.900	0.575	0.610	0.188	0.400
Differential effect	0.550 (0.762)	0.862 (1.410)	-1.732 (1.150)	-3.857** (1.600)	-6.836*** (0.990)	-6.532*** (1.172)
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 parentheses are clustered at the field level. Sample is the panel of oilfields from 2006-2016. Output is measured in millions of barrels of oil per year. Oil spills 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. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

reaucratic reasons typically exogenous to field characteristics or trends.³² 88% 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 \tilde{t} , the announcement date, but not actually divested until $\tilde{t} + k$, 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 oil output (columns (1)-(3)) and oil theft (columns (4)-(6)). Columns (1) and (4) test asset transfers between local firms (local-to-local transitions), while columns (2) and (5) test all non-localizing transitions. For both outcomes, the placebo coefficients are small and insignificant, while the coefficients on the “true” MNC-

³²For 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. See articles [here](#) and [here](#).

to-local divestment indicator remain large and significant.

Columns (3) and (6) 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 consummated, if at all. Figure A19 estimates the TWFE event-study using the announcement year for terminated divestments. There is no evidence of any systematic trends or treatment effects of the terminated transactions. One caveat is that the placebo transitions and terminated divestments are somewhat rare and may be underpowered.³³ 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.

Table 5: The effect of divestment on output and oil theft, placebo tests

Outcome	Output			Oil theft		
	(1)	(2)	(3)	(4)	(5)	(6)
Local firm	0.943*** (0.307)	0.926*** (0.307)		-5.703*** (1.034)	-5.568*** (1.029)	
Local-to-local sale	0.213 (0.870)			0.002 (1.021)		
Non-local transaction		0.421 (0.732)			-1.550 (1.354)	
Terminated divestment			0.182 (0.353)			0.523 (2.105)
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	2476	2476	2476	3183	3183	3183
R ²	0.878	0.878	0.877	0.756	0.756	0.753

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. “Local-to-local sale” is an indicator that equals one in all periods after a field is sold from a one Nigerian buyer to another, as derived from transactions data. Non-local transaction additionally includes MNC-to-MNC and local-to-MNC asset sales. 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5.3.4 Additional robustness tests

In Appendix C, I test the robustness of the main results to many possible sources of bias. C.1 considers violations of the identifying assumption driven by time-varying, locality-

³³There are 27 local-to-local transitions, three MNC-to-MNC, three local-to-MNC, and 36 terminated or delayed.

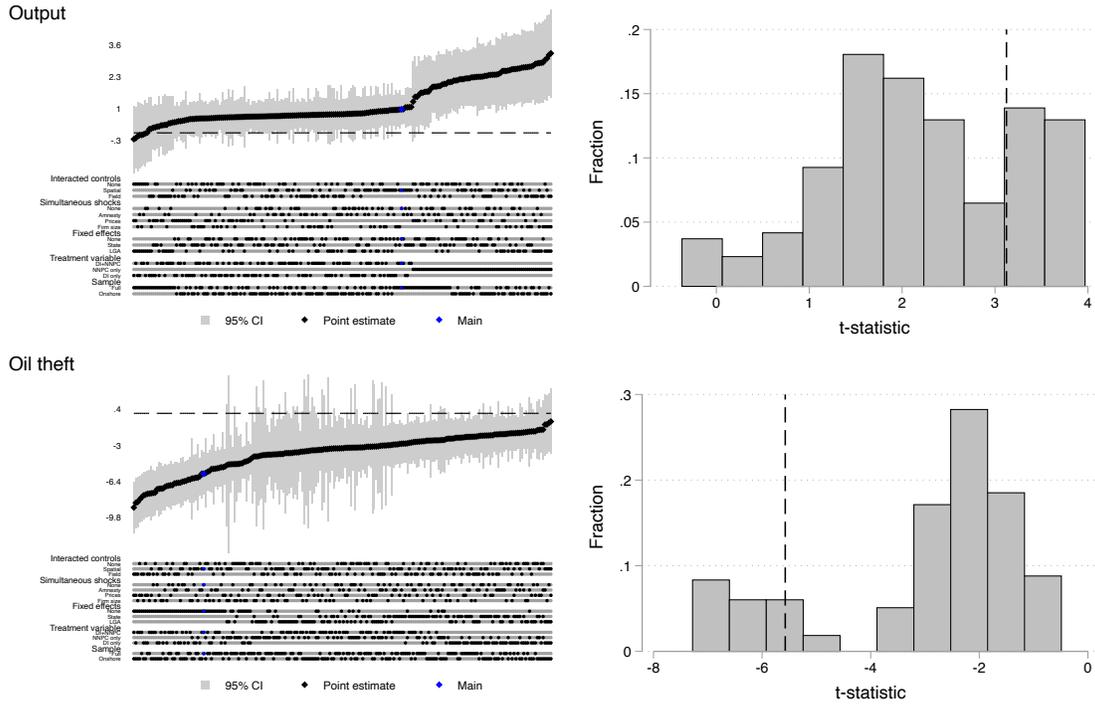
specific effects (Table A3), differential trends by field-level technical covariates (Table A4), differential oil price responses (Table A5), the impact of Federal amnesty for militant groups (Table A6), and the role of firm size (Table A7). In C.2 I estimate exponential models, finding similar results. C.3 considers inferential problems arising from correlation in treatment assignment, using randomization inference (Figure A9) and oil block-level clustering (Table A9). I also test robustness to sample restrictions (Table A10), nonlinear regression imputation of missing output (Table A11), Winsorizing outliers (Table A12), and distance radii for outcomes (Table A13). In C.9 I disaggregate treatment by state-owned vs. private local firms (Table A15) and NNPC data vs. DI data (Table A16). Table A18 investigates theft effects by asset type. In Appendix E, I further subject the results of Table 3 to additional tests for measurement error, treatment definition, and event-study specification.

Figure 4 summarizes the results of 216 different specifications for output and theft. The specification set features all possible combinations of 3 interacted control specifications, 4 simultaneous shocks, 3 interacted fixed effects, 3 different treatment measures, and 2 samples.³⁴ 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 not in the tail of the coefficient distribution, indicating the preferred specification is unlikely to be spurious. Nearly all estimates are of the right sign, while the majority are significant at the 5%.

Lastly, a growing literature in applied microeconometrics identifies potential bias in the TWFE estimator in staggered-adoption designs. The TWFE estimate is a weighted average of 2×2 DD comparisons across treatment cohorts and control groups. Unfortunately, some of these 2×2 comparisons use already-treated observations as controls, even though these units' treatment status should alter their subsequent trends. Appendix D considers a battery of diagnostic tests and alternative estimators. Table A21 decomposes the TWFE estimate into all 2×2 comparisons. While the results are driven largely by the treated vs. never-treated comparison, estimates that leverage *only* treatment timing – which was argued to be plausibly random – are similar, as are reweighted estimates that purge “forbidden” comparisons. Table A22 estimates a stacked specification (Baker et al. 2021), which restricts the control group only to never-treated and later-treated observations. The results are, if anything, stronger, implying that TWFE biases the main results downward. Table A23 uses the semiparametric estimator of Callaway and Sant’Anna (2021), yielding similar results. In Appendix D.5, I further consider many possible event-study specifications – unconditional (Figure A12), different control groups (Figure A13), TWFE (Figure A14), NNPC (Figure A15) vs. DI (Figure A16) treatments, as well as dynamic estimates following Callaway and Sant’Anna (2021) and de Chaisemartin and D’Haultfoeuille (2020) in Figures A17 and A18.

³⁴They are: *i*) with none, spatial, or field controls, *ii*) with none, price, or amnesty, or firm size controls, *iii*) with none, state-, or LGA-by-year fixed effects, *iv*) using both datasets, DI only, or NNPC only, and *v*) with full sample or onshore field only.

Figure 4: Robustness plots



Note: Figure displays estimated indigenization effects for robustness tests across 216 specifications, for oil production (top panel) and oil theft (bottom panel). Specification is indicated by points in the bottom of the figure. Righthand panel displays estimated t -statistics for these specifications, with dotted lines for the “main” specification of Table 2, columns (2) and (6). 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 shocks indicate inclusion of price interaction terms, as in Table A5. Amnesty shocks indicate inclusion of controls for the 2009 Niger Delta amnesty, as in Table A6. Firm size specifications include controls for the number of wells owned by the operating firm, as in Table A7. State and LGA indicate state- or local government area-by-year fixed effects, respectively. The main treatment variable is DI+NNPC, which uses both sources of data, as in Tables A16. Sample is either all or all offshore nonmissing observations for the outcome in question.

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 Appendix B, and here only describe the strategic environment and illustrate important predictions. In the model, firms bargain with state security agents over protection, 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 stage game is a simple two-stage, one-shot interaction between firms f , gangs g , and a state security agent. The firm produces a fixed level of oil output \bar{Q} , sold at the international oil price p . The firm and the gang simultaneously offer bribes b_f and b_g to the state. 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 $c - \epsilon_g > 0$, 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 enforce the law. The enforcement technology reduces the probability of theft (and gangsters' expected profits) to $\alpha < 1$ at cost $\eta > 0$.

If b_f is accepted, the firm pays an expected penalty $\lambda > 0$, capturing the legal or reputational costs of corrupt payments. In addition, firms only receive a share γ of \bar{Q} , to capture the important role of joint-ventures in Nigeria, as shown in Figure A4. Importantly, law enforcement may internalize the firm's profits by a factor $\mu \in [0, 1]$, measuring the strength of political connections. In the base case of the model, under uniform ϵ_g , the enforcement decision of the security forces is summarized in the following equation:

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

Government enforcement increases in μ , political connections, by lowering the threshold b_f required to sway the state. It also rises in γ , the ownership share, by increasing firm willingness-to-pay. An increase in q reduces enforcement, since it increases the value of criminal bribes more than it increases the collective costs of theft to government and firm, neither of which fully internalize the loss. Enforcement also declines in η and increases in government effectiveness $1 - \alpha$. A positive price shock increases enforcement only when $\frac{\kappa}{q} > \frac{(1 - \gamma - \mu)}{(\gamma + \mu)}$, so that the theft losses are sufficiently large relative to the increase in value of the criminal bribe. The central point is that inefficient theft may occur in equilibrium, driven by bargaining frictions. However, with political connections, government internalizes the losses from theft, lowering the cost of bribery and increasing the scope for efficient corruption.

The base case (Appendix B.2) optimistically assumes that all parties can perfectly commit. In reality, 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 B.3 that commitment problems substantially increase the importance of political connections. Now, government will only enforce the law when the firm is sufficiently connected:

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

Commitment problems imply that bribes alone will not suffice, since the state can always renege; enforcement is only provided when state security agents have a large enough stake in

the profits of the firm. Lastly, in Appendix B.4 I show that repeated interaction restores 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.

7 Mechanisms

7.1 Law enforcement corruption

If political connections provide local firms protection from black market predation, then local takeover should increase law enforcement actions on divested assets. In Table 6, I estimate the main TWFE model using data on anti-oil theft law enforcement actions from Nigerian news reports. The outcomes, defined in the Table header, are the counts of various anti-oil theft enforcement actions – any event (1-2), seizures of stolen oil (3-4), illegal refinery raids (5-6), and illegal export raids (7-8). I disaggregate these outcomes because enforcement patterns may plausibly differ along the black market value chain.³⁵ Local takeover leads to a large increase in enforcement provision. Anti-oil theft law enforcement actions increase by 1.2-1.7 events (roughly 83-117% greater than the control group mean). The effects hold across outcomes, and are always significant at 5 or 10%.³⁶

Table 6: The effect of local ownership on law enforcement activity

Enforcement outcome	All oil theft		Oil seizures		Illegal refineries		Illegal export	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Local firm	1.652** (0.770)	1.163* (0.688)	1.438** (0.620)	1.012* (0.545)	1.285*** (0.461)	0.882** (0.416)	0.411** (0.199)	0.412** (0.173)
Control group mean	1.418		0.844		0.694		0.270	
Observations	3183	3183	3183	3183	3183	3183	3183	3183
R ²	0.481	0.578	0.436	0.533	0.436	0.531	0.323	0.423
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

Standard errors, in parentheses, 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 the total count of anti-oil theft law enforcement events 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. ** $p < 0.01$, * $p < 0.05$, $p < 0.1$.

This result may simply reflect a generalized increase in law enforcement effort coinciding with local takeover. In Table A29, I consider several placebo outcomes – enforcement on non-oil crime, kidnapping, and gang activity. I find that that localization has no effect on these

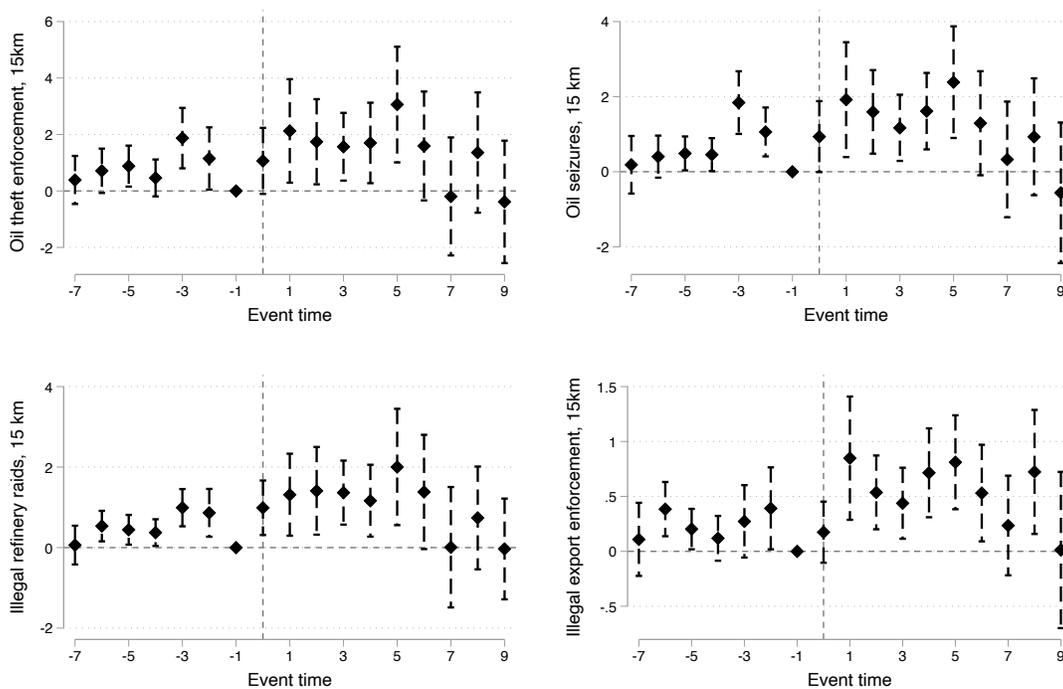
³⁵For example, export activities are controlled by powerful militant groups with connections to organized crime, while refineries are primarily the purview of smaller-scale entrepreneurs.

³⁶In Appendix Table A28, I estimate similar, slightly stronger results in the stacked model.

other enforcement outcomes that are unrelated to the protection of oil assets.

Figure 5 plots corresponding event-study coefficients using the stacked model. The top left panel estimates the model for all anti-oil theft enforcement. The dynamic coefficients indicate parallel trends in enforcement in the years leading up to local takeover, with the exception of $\tau = -2, -3$.³⁷ Enforcement spikes in the years after divestment and remains significantly elevated until $\tau = 6$, at which point it falls to zero. Combining this plot with the event-study results for oil theft in Figure 3, it appears that an immediate jump in enforcement drives a fast and durable reduction in theft, which in turn leads to an eventual tapering of enforcement activity over time. In the remaining panels of Figure 5, I consider the other subcategories of anti-oil theft enforcement from columns (3)-(8), which exhibit similar dynamics.

Figure 5: Stacked-DD event-study: enforcement outcomes



Note: Figure shows coefficients from stacked event-study regressions described in Section D.3 for oil theft enforcement outcomes. Standard errors are clustered at the field-level. All enforcement outcomes are the total count of enforcement events of each subtype within 15 kilometers of a field centroid. Specifications include interacted controls for latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital.

³⁷it's not entirely clear what drives this spike in treated fields' enforcement 2-3 years before divestment, though the most likely explanation is noise in the outcome variable.

7.2 Political connections

To obtain greater law enforcement protection, local firms may lean on their superior political connections. But are they better connected? Table 7 shows the share of connected firms by type for the years 1998-2016. Local and multinational firms have similar and high rates of political connectedness – 80% of firms have at least one shareholder or board member with any experience in government or traditional leadership over this period. This number falls to around 73% when we exclude connections only to traditional leaders, who do not hold official government posts. However, the composition of connections differs markedly between locals and multinationals. Local firms are more likely to pursue links with elected politicians and the security forces, while multinationals rely more heavily on connections to traditional leaders and high-level political appointees. I argue that the connections to security forces in particular provide local firms a key advantage in obtaining protection for their assets.

Table 7: Political connections by firm type

	Local	MNC	Total
Any politician	0.79	0.80	0.80
Any politician (excluding chief)	0.77	0.60	0.73
Elected	0.33	0.10	0.29
Technocrat	0.72	0.60	0.69
High-level	0.21	0.30	0.22
Security forces	0.28	0.00	0.22
Traditional leader	0.49	0.60	0.51
Number of observations	39	10	49

Table shows the share of connected firms by connection type for the multinational and local subsamples as well as the full sample. Data is the universe of operating firms in Nigeria’s oil sector; for a complete list of firms, see Table A2. Firms are coded as connected if they ever had a connection over the period of 1998-2016.

To provide additional support for this hypothesis, I estimate the value of connectedness in Table 8, which contains the results of the TWFE regression of theft on political connections measures.^{38,39} Column (1) shows that there is no association between political connections

³⁸The connections measures vary over time at the firm-level. This estimate is therefore identified both from field takeovers by politically connected firms *and* change in connectedness over time for a given field operator.

³⁹The 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 or traditional leadership on the board or among the shareholders and beneficial owners. “Technocrats” are those associated with ministerial posts or regulatory agencies, typically the NNPC and DPR. “Elected” indicates an elected official, such as a governor or senator. “Cabinet-level politician” indicates connection to a politician who at some point served in a ministerial post. “Chief” indicates connection to the holder of a traditional tribal title or chieftancy. Lastly, “security forces” are those linked to the military or police forces.

and oil theft, conditional on TWFE and interacted controls. Columns (2)-(5) further show that specific connections to technocrats, elected politicians, cabinet-level figures, and traditional leaders are not significantly associated with oil theft, although the point estimate on elected connections is negative and represents 17% of the outcome mean. Column (6), however, shows that connections to the Nigerian security forces are associated with a large reduction in theft, significant at 1% and equivalent to 42.7% of the outcome mean.⁴⁰ In Appendix F.1, I show that connections to the security forces are additionally associated with significantly greater output, oil spills, and enforcement (Table A31).⁴¹

Table 8: The effect of political connections on oil theft

Connection	Any	Tech.	Elected	Cabinet	Chief	Security
	(1)	(2)	(3)	(4)	(5)	(6)
Connected	0.168 (0.811)	0.070 (0.782)	-1.635 (1.033)	-0.761 (0.905)	-0.190 (0.954)	-4.068*** (1.492)
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	3183	3183	3183	3183	3183	3183
R^2	0.753	0.753	0.753	0.753	0.753	0.754

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

Political connections in general do not mitigate the activities of organized crime. However, connections to the security forces in particular – the primary political actors engaged in black market activity – are critically important. This suggests that connections to former high-ranking officials in the Nigerian armed forces give local 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 capacity of these agencies. Since the decision by security forces to enforce the law or collude with organized crime is a key determinant of the viability of theft, connected firms are able to reduce theft on their assets. Table A31 confirms that these effects are driven by increased enforcement, and result in greater output.

⁴⁰In additional results, not reported, I find that when field fixed effects are removed, all connection types are associated with significantly lower oil theft. However, the coefficient on security connections remains largest.

⁴¹Furthermore, event-study models in Figure A22 show parallel trends in oil theft before the introduction of a connected shareholder or board member, and large annual declines in theft thereafter.

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.

Ownership consolidation: Operators with larger ownership stakes γ internalize a greater share of the losses from theft. Because of indigenization policies and 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, 43 p.p. less likely to obtain sole-risk licenses, have government stakes roughly 85% higher than the average Nigerian independent operator. In Appendix F.2, I use oil block licensing data to show that divestment 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% (Table A32).

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. Given the relatively broad definitions of foreign officials contained in these laws, the prospect of legal liabilities could plausibly deter multinationals from bargaining with political power-brokers. In Appendix F.3, I test whether exposure to an international corruption law affects outcomes. I find that passage of a home-country corruption law is associated with a statistically significant annual increase in theft and violent conflict in the sample of multinational-operated fields (Table A33). Despite the small sample (there are only 5 distinct home countries), the evidence suggests that anti-corruption laws affect multinationals' ability to use corrupt payments to deter crime.

7.4 Alternative explanations

Spatial spillovers: In general equilibrium, localization could increase targeting of surrounding multinational fields if local fields are politically protected but their multinational neighbors are not. Such spillovers may bias the treatment effect by violating SUTVA (Rubin 2005), and would have important implications for the aggregate welfare effects of localization. In Appendix G.1, I estimate negligible spillovers in theft or output (Figure A23); accounting for these spillovers does not substantially alter the main treatment effects.

Differences in discount rates: The local advantage in production may be driven by different optimal extraction profiles given underlying time preferences. This is unlikely, as oil output is difficult to adjust along the intensive margin in the short-run for a given stock of fixed capital, and most of the increase in output from localization can be attributed to reducing the quantity lost to theft (see Appendix C.12). In Appendix G.2, I estimate nonparametric age-extraction profiles for multinational and local fields (Figure A24). Using linear splines, I do not find that the slopes of multinational extraction profiles are significantly different (Table A34).

Grievance toward multinationals: Criminal and militant activity may be driven by grievance rather than economic motives (Buhaug et al. 2014). Local sentiments toward Nigerian companies may be considerably better than toward multinationals. If so, we should expect to observe a reduction in community protest against oil companies, a relatively common occurrence in host communities, affecting 21% of fields during the sample period. I assess this in Appendix G.3, which re-estimates the TWFE specification using protests and riots as the outcome variables (Table A35). There is no evidence of a change in grievance after localization.

Local employment spillovers: Localization may generate positive economic spillovers to local labor markets, bidding up the opportunity cost of joining gangs, so that crime-reduction effects might be driven by labor costs in the criminal sector. To test this hypothesis, I use three rounds of a national panel survey on 16,211 working-age Nigerians from 2010-2016, linking each household to its nearest oilfield. In Appendix G.4, I find no evidence of employment or consumption spillovers from local ownership (Table A37 and Figures A25, A26).

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 than to negotiate protection directly. If local firms are better corporate citizens, this could account for localization benefits.⁴² In Appendix G.5, I use cross-sectional data on oil company CSR projects in 2016 to test whether local firms are better at targeting their investments. Though these cross-sectional correlations are somewhat speculative, I find that, if anything, multinational CSR investments are more responsive to recent levels of local militant violence (Figure A27).

8 Conclusion

Firms in the natural resource sectors of weak states are subject to violent extortion by armed groups and suffer losses to illicit markets. Despite the fact that this creates a willingness to pay for protection, the state typically fails to protect these firms. I study the roots of this bargaining failure in the context of the Nigerian oil sector. I argue that in Nigeria's oil sector – and others like it – where militant groups and organized crime are threats to firm operations and corruption buys law enforcement protection, local companies possess distinct advantages over their better-resourced multinational peers.

Using data on Nigerian oilfields from 2006-2016, I find that divestment to local firms significantly increases field output by nearly 1 million barrels per year, a 33% gain. However, consistent with an operational disadvantage, local fields experience more oil spills from equipment malfunctions and flare more natural gas. Both of these practices entail substantial

⁴²In 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 tiny fraction of the annual profits from oil theft, these projects are unlikely to meaningfully dissuade crime.

environmental costs. The key to the local output advantage lies in the multi-billion dollar black market for stolen oil. I find that local takeover reduces oil theft by 55%, and mitigates armed conflict as well. I further find that these gains are concentrated in the onshore fields most susceptible to crime and violence, whereas the technical oil spill losses 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 allow local firms to expand the scope for efficient corruption, purchase law enforcement protection, and outperform multinationals. 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 connections drive local advantage via law enforcement corruption, I first show that local firms indeed receive preferential state protection: localization leads to an 83% increase in law enforcement actions against oil theft. Using data on firms' boards and shareholders, I show that political connections are widespread in the sector across all firms. Local firms, however, are much more likely to possess the connections to security elites required to protect assets. Finally, I test for returns to political connections in reduced theft. Consistent with the model, connections to the security forces – the agencies at the center of black market corruption – are highly beneficial, while other connections do not affect theft.

The findings suggest that when institutions are weak, 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, the environmental costs reveal important tradeoffs in allowing politically connected local firms to dominate the marketplace. The findings also raise important questions around why multinationals fail to identify and build effective political influence networks, for example by using joint ventures. I leave these for future work.

Lastly, I argue that the Nigerian oil and gas sector is closer to a representative than a pathological case. In extractive industries 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. In these cases, we must reconsider the role of indigenous firms and the benefits of foreign investment.

References

- Acemoglu, D. (2003). Why not a political Coase theorem? Social conflict, commitment, and politics. *Journal of Comparative Economics*, 31(4), 620–652.
- Aigbovo, O., & Atsegbua, L. (2013). Nigerian anti-corruption statutes: an impact assessment. *Journal of Money Laundering Control*, 16(1), 62–78.
- Aitken, B., & Harrison, A. (1999). Do Domestic Firms Benefit from Direct Foreign Investment? Evidence from Venezuela. *American Economic Review*, 89(3), 605–618.
- Akcigit, U., Baslandze, S., & Lotti, F. (2018). Connecting to Power: Political Connections, Innovation, and Firm Dynamics. *NBER Working Paper*, 25136.
- Alfaro, L., & Chauvin, J. (2020). Foreign Direct Investment, Finance, and Economic Development. In F. L. Rivera-Batiz, M. Spatareanu, & C. Erbil (Eds.), *Encyclopedia of International Economics and Global Trade*. Singapore: World Scientific.
- Alfaro, L., & Chen, M. X. (2018). Selection and Market Reallocation: Productivity Gains from Multinational Production. *American Economic Journal: Economic Policy*, 10(2), 1–38.
- Arnold, J. M., & Javorcik, B. S. (2009). Gifted kids or pushy parents? Foreign direct investment and plant productivity in Indonesia. *Journal of International Economics*, 79(1), 42–53.
- Asuni, J. B. (2009). Understanding the Armed Groups of the Niger Delta. *Council on Foreign Relations Working Paper*.
- Autor, D. H., Palmer, C. J., & Pathak, P. A. (2014). Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts. *Journal of Political Economy*, 122(3), 661–717.
- Baker, A. C., Larcker, D. F., & Wang, C. C. (2021). How Much Should We Trust Staggered Difference-InDifferences Estimates? *Working Paper*.
- Balsvik, R. (2011). Is Labor Mobility a Channel for Spillovers from Multinationals? Evidence from Norwegian Manufacturing. *The Review of Economics and Statistics*, 93(1), 285–297.
- Baranek, B., & Titl, V. (2020). The cost of favoritism in public procurement. *FEB Research Report Department of Economics, KU Leuven*.
- Bardhan, P. (1997). Corruption and Development: A Review of Issues. *Journal of Economic Literature*, 35(3), 1320–1346.
- Beckert, J., & Dewey, M. (2017). *The architecture of illegal markets : towards an economic sociology of illegality in the economy*. Oxford University Press.
- Berman, N., Couttenier, M., Rohner, D., & Thoenig, M. (2017). This Mine is Mine! How Minerals Fuel Conflicts in Africa. *American Economic Review*, 107(6), 1564–1610.
- Blair, G., Christensen, D., & Wirtschafter, V. (2022). How Does Violence Shape Investment? Evidence from Mining. *The Journal of Politics*, 84(1).
- Bloom, N., Sadun, R., & van Reenen, J. (2012). Americans Do IT Better: US Multinationals and the Productivity Miracle. *American Economic Review*, 102(1), 167–201.

- Bloom, N., & van Reenen, J. (2010). Why Do Management Practices Differ across Firms and Countries? *Journal of Economic Perspectives*, 24(1), 203–224.
- Bruederle, A., & Hodler, R. (2019). Effect of oil spills on infant mortality in Nigeria. *PNAS*, 116(12), 5467–5471.
- Buhaug, H., Cederman, L.-E., & Gleditsch, K. S. (2014). Square Pegs in Round Holes: Inequalities, Grievances, and Civil War. *International Studies Quarterly*, 58(2), 418–431.
- Buonanno, P., Durante, R., Prarolo, G., & Vanin, P. (2015). Poor Institutions, Rich Mines: Resource Curse in the Origins of the Sicilian Mafia. *Economic Journal*, 125(August), 175–202.
- Burger, M., Ianchovichina, E., & Rijkers, B. (2015). Risky Business: Political Instability and Sectoral Greenfield Foreign Direct Investment in the Arab World. *World Bank Economic Review*, 30(2), 306–331.
- Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225(2), 200–230.
- Chapman, T. L., Jensen, N. M., Malesky, E. J., & Wolford, S. (2021). "Leakage" in International Regulatory Regimes: Did the OECD Anti-bribery Convention Increase Bribery? *Quarterly Journal of Political Science*, 16(4), 387–427.
- Chimeli, A. B., & Soares, R. R. (2017). The Use of Violence in Illegal Markets: Evidence from Mahogany Trade in the Brazilian Amazon. *American Economic Journal: Applied Economics*, 9(4), 30–57.
- Christensen, D. (2019). Concession Stands: How Mining Investments Incite Protest in Africa. *International Organization*, 73(1), 65–101.
- Correia, S., Guimarães, P., & Zylkin, T. (2020). Fast Poisson estimation with high-dimensional fixed effects. *The Stata Journal*, 20(1), 95–115.
- Couttenier, M., Grosjean, P., & Sangnier, M. (2017). The Wild West is Wild: The Homicide Resource Curse. *Journal of the European Economic Association*, 15(3), 558–585.
- Criscuolo, C., & Martin, R. (2009). Multinationals and U.S. Productivity Leadership: Evidence from Great Britain. *The Review of Economics and Statistics*, 91(2), 263–281.
- de la Sierra, R. S. (2020). On the Origin of States: Stationary Bandits and Taxation in Eastern Congo. *Journal of Political Economy*, Forthcoming.
- de Chaisemartin, C., & D'Haultfoeuille, X. (2020). Two-way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, 110(9), 2964–2996.
- Deshpande, M., & Li, Y. (2019). Who Is Screened Out? Application Costs and the Targeting of Disability Programs. *American Economic Journal: Policy*, Forthcoming.
- Diamond, R., & McQuade, T. (2019). Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development. *Journal of Political Economy*, 127(3), 1063–1117.

- Dube, O., & Vargas, J. F. (2013). Commodity Price Shocks and Civil Conflict: Evidence from Colombia. *Review of Economic Studies*, 80, 1384–1421.
- Faccio, M. (2006). Politically Connected Firms. *American Economic Review*, 96(1), 369–386.
- Fetzer, T., & Kyburz, S. (2018). Cohesive Institutions and Political Violence. *Centre for Competitive Advantage in the Global Economy, Working Paper*.
- Fisman, R. (2001). Estimating the Value of Political Connections. *American Economic Review*, 91(4), 1095–1102.
- Garcia-Jimeno, C. (2016). The Political Economy of Moral Conflict: An Empirical Study of Learning and Law Enforcement Under Prohibition. *Econometrica*, 84(2), 511–570.
- Gehring, K., Langlotz, S., & Kienberger, S. (2020). Stimulant or Depressant? Resource-Related Income Shocks and Conflict. *Households in Conflict Network (HiCN) Working Paper No. 286*.
- Goodman-Bacon, A. (2021). Difference-in-Differences with Variation in Treatment Timing. *Journal of Econometrics*, 225(2), 254–277.
- Gormley, T. A., & Matsa, D. A. (2011). Growing Out of Trouble? Corporate Responses to Liability Risk. *Review of Financial Studies*, 24(8), 2781–2821.
- Guadalupe, M., Kuzmina, O., & Thomas, C. (2012). Innovation and Foreign Ownership. *American Economic Review*, 102(7), 3594–3627.
- Guidolin, M., & Ferrara, E. L. (2007). Diamonds Are Forever, Wars Are Not: Is Conflict Bad for Private Firms? *American Economic Review*, 97(5), 1978–1993.
- Harrison, A., & Rodriguez-Clare, A. (2010). Trade, Foreign Investment, and Industrial Policy for Developing Countries. In D. Rodrick & M. R. Rosenzweig (Eds.), *Handbook of Development Economics, Volume 5* (pp. 4040–4214). Elsevier B.V.
- Hirschman, A. O. (1969). *How to divest in Latin America, and why*. International Finance Section, Princeton University.
- 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. *PhD Thesis, Department of Geography, Birkbeck University*.
- Huntington, S. (1968). *Political Order in Changing Societies*. Yale University Press.
- Javorcik, B. S., & Wei, S.-J. (2009). Corruption and cross-border investment in emerging markets: Firm-level evidence. *Journal of International Money and Finance*, 28(4), 605–624.
- Kaufmann, D., Kraay, A., & Zoido-Lobaton, P. (1999). *Governance matters* (Policy Research Working Paper Series No. 2196). The World Bank.
- Khwaja, A. I., & Mian, A. (2005). Do Lenders Favor Politically Connected Firms? Rent Provision in an Emerging Financial Market. *The Quarterly Journal of Economics*, 120(4), 1371–1411.

- Li, H., Meng, L., Wang, Q., & Zhou, L.-A. (2008). Political Connections, financing, and firm performance: Evidence from Chinese private firms. *Journal of Development Economics*, 87, 283–299.
- Lipsey, R. G., & Lancaster, K. (1956). The General Theory of Second Best. *The Review of Economic Studies*, 24(1), 11–32.
- Méon, P.-G., & Weill, L. (2010). Is Corruption an Efficient Grease? *World Development*, 38(3), 244–259.
- NBS. (2017). *Demographic Statistics Bulletin* (tech. rep.). Nigeria National Bureau of Statistics. Abuja, Nigeria.
- NEITI. (2016). *2016 Oil and Gas Industry Audit Report* (tech. rep.). Nigeria Extractive Industries Transparency Initiative. Abuja, Nigeria.
- Network, S. D. (2021). *Divesting from the Delta: Implications for the Niger Delta as international oil companies exit onshore production* (tech. rep.). Stakeholder Democracy Network.
- Nwokolo, A. (2018). Oil Price Shocks and Civil Conflict: Evidence from Nigeria. *HiCN Working Paper*, 274.
- Obi, C., & Rustad, S. A. (2011). *Oil and Insurgency in the Niger Delta: Managing the Complex Politics of Petroviolence*. Zed Books.
- Ologunorisa, T. E. (2009). A review of the effects of gas flaring on the Niger Delta environment. *International Journal of Sustainable Development and World Ecology*, 8(3), 249–255.
- Palepu, K. G., & Khanna, T. (1998). Institutional Voids and Policy Challenges in Emerging Markets. *The Brown Journal of World Affairs*, 5(1), 71–78.
- Poole, J. P. (2013). Knowledge Transfers from Multinational to Domestic Firms: Evidence from Worker Mobility. *The Review of Economics and Statistics*, 95(2), 393–406.
- Ralby, I. (2018). *Downstream Oil Theft: Global Modalities, Trends, and Remedies* (tech. rep.). The Atlantic Council.
- Rexer, J., & Hvinden, E. (2022). Violence and Black Markets: Evidence from the Niger Delta Conflict. *Working Paper*.
- Roth, J. (2019). Pre-test with Caution: Event-study Estimates After Testing for Parallel Trends. *Working Paper*.
- Rubin, D. B. (2005). Causal Inference Using Potential Outcomes: Design, Modeling, Decisions. *Journal of the American Statistical Association*, 100(469), 322–331.
- SDN. (2019a). *Communities not Criminals: Illegal Oil Refining in the Niger Delta* (tech. rep.). Stakeholder Democracy Network. Port Harcourt, Nigeria.
- SDN. (2019b). *Economic Dynamics of the Artisanal Oil Industry in the Niger Delta over Five Years* (tech. rep.). Stakeholder Democracy Network. Port Harcourt, Nigeria.
- SDN. (2019c). *Pipeline surveillance contracts in the Niger Delta* (tech. rep.). Stakeholder Democracy Network. Port Harcourt, Nigeria.

- Sexton, R. (2019). Unpacking the Local Resource Curse: How Externalities and Governance Shape Social Conflict. *Journal of Conflict Resolution*, *Forthcoming*.
- Shleifer, A., & Vishny, R. W. (1993). Corruption. *The Quarterly Journal of Economics*, *108*(3), 599–617.
- Teece, D. (1977). Technology Transfer by Multinational Firms: The Resource Cost of Transferring Technological Know-how. *Economic Journal*, *87*(346), 242–261.
- Watts, M. (2007). Petro-insurgency or Criminal Syndicate: Conflict and Violence in the Niger Delta. *Review of African Political Economy*, *34*(114), 637–660.
- Wooldridge, J. M. (1999). Distribution-free estimation of some nonlinear panel data models. *Journal of Econometrics*, *90*(1), 77–97.
- Zabyelina, Y., & van Uhm, D. (2020). The New Eldorado: Organized Crime, Informal Mining, and the Global Scarcity of Metals and Minerals. In Y. Zabyelina & D. van Uhm (Eds.), *Illegal Mining: Organized Crime, Corruption, and Ecocide in a Resource-Scarce World* (pp. 3–30). Springer International Publishing.

ONLINE APPENDIX

— For Online Publication Only —

A Data appendix

A.1 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)⁴³ 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 1998-2016.⁴⁴ 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 are 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 A1. 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 A2 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.

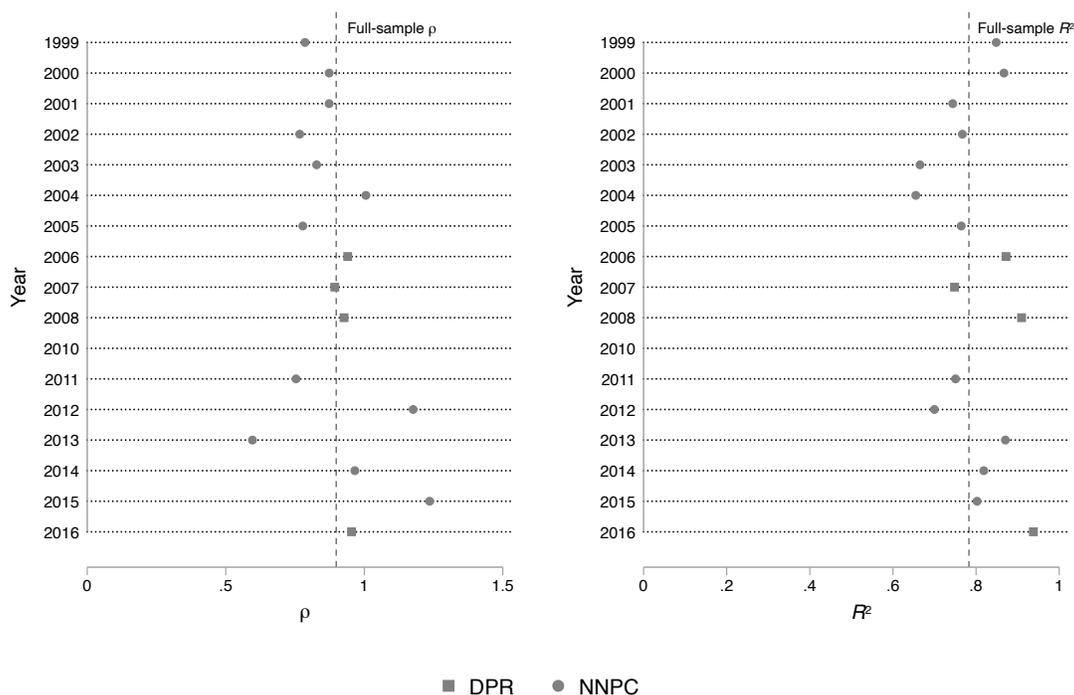
From DPR I also get field-level time-invariant covariates: the number of wells (field size), date of completion of the first well (field age), and the depth of the deepest well. Finally, I use infrastructure maps to obtain centroid locations for the fields in the DPR-NNPC data, which are then used to 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.

Measurement error in output may be correlated with localization. For example, firms may underreport output in order to evade royalty taxation. Since multinational firms tend to re-

⁴³The DPR is Nigeria’s primary petroleum sector regulatory body.

⁴⁴Unfortunately, disaggregated data are unavailable for 2009.

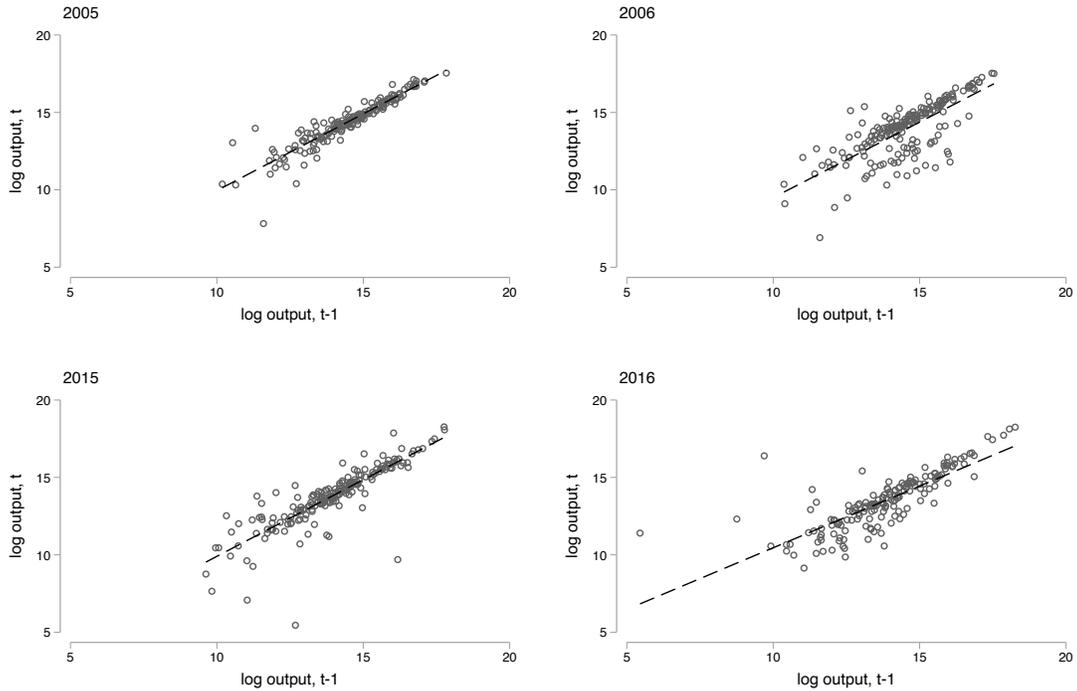
Figure A1: Year-to-year correlations in oil output



Note: Figure shows coefficient estimates (left panel) and R^2 (right panel) from separate AR(1) regressions of oil output for each consecutive year pair in the data, indicated on the vertical axis by the second year of the pair. Vertical dashed lines indicates the coefficient or R^2 from an AR(1) regression on the pooled full sample. Marker symbols indicate data source by year. Sample is an unbalanced panel of 314 oilfields from 1998-2016. Oil production data are missing for 2009, so estimates for 2008-2009 and 2009-2010 are excluded.

ceive less preferential treatment, the incentive to misreport may be greater. In this case, any local output advantage that materializes after divestment may simply be driven by reporting biases. I attempt to corroborate the reliability of Nigerian administrative data on output using independent data on export values from UN COMTRADE. This is a reasonable comparison, since the vast majority of Nigerian oil (85-90%) is exported. I aggregate all administrative NNPC data from 1998-2015 to the year level and value this production at annual world crude oil prices. I then correlate these against reported COMTRADE crude oil export values, in Figure A3. The two series are highly correlated over time. Furthermore, if selective reporting were driving the results, then the correlation between the two series should strengthen over time (in particular, after 2010) as the local market share grows and under-reporting falls. There is no evidence that post-2010 observations are systematically more correlated; in fact, observations in both periods are tightly clustered around the regression line. However, small sample caveats apply, since the analysis relies on aggregate time series data.

Figure A2: Oil output, selected consecutive years

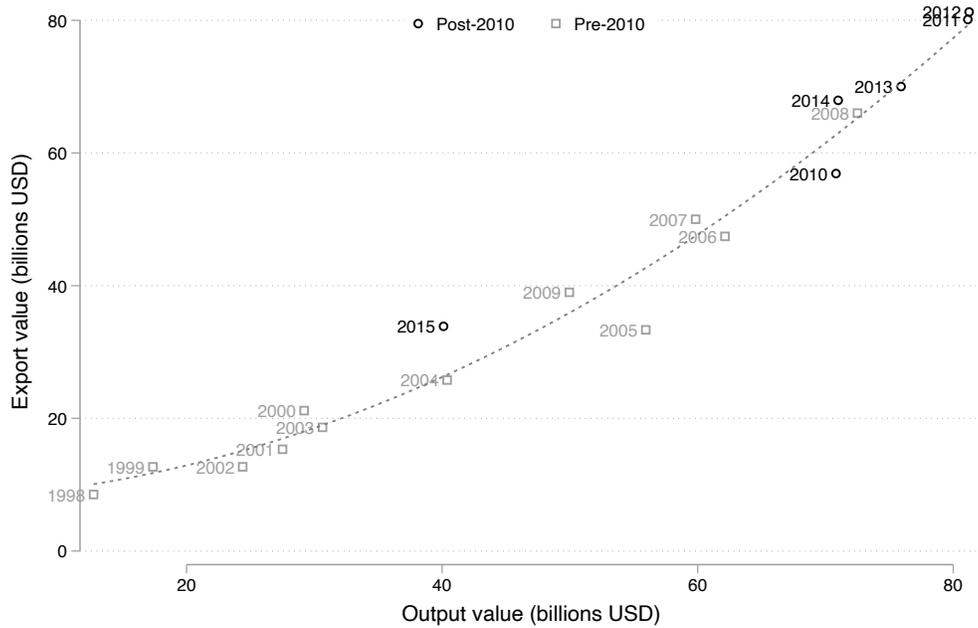


Note: Figure shows field-level scatterplots and linear fits of the log of output in t against $t - 1$ for selected consecutive year-pairs in which the data source for oil production changes. Sample is an unbalanced panel of 314 oilfields from 1998-2016. Oil production data are missing for 2009.

A.2 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 155 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 type of transaction (corporate M&A, new discoveries not yet developed, exploration blocks previously awarded, fields under development, producing fields, and new awards). I then drop new exploration awards, which cover license awards from the Nigerian government to private firms, since these licenses contain unexplored fields that do not enter into the data. I also drop corporate M&A transactions, which typically do not refer to specific assets but rather reflect changes in the ownership structure of entire firms. 117 transactions remain after these sample restrictions. For each asset transaction, I retain the nationality of buyers, sellers, the transaction opening and closing dates, and whether it was successful.

Figure A3: Output measurement validation: export value



Note: Figure shows time-series correlation between aggregate oil production value as measured by Nigerian administrative sources and Nigerian oil export values from UN Comtrade data, both measured in billions of USD. Pre- and post-2010 observations are indicated in figure.

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. Of the 117 DI transactions, 74 contain specific fields, covering 104 unique fields. The remaining 43 transactions mention only oil blocks, and cover 44 distinct blocks. I then merge these data to the main field-level dataset. In total, 43 out of 104 field-level transactions are matched, and 15 out of 44 blocks are matched. Since only 27% of the 117 transactions cover assets that are actively producing at the time of the transaction, these match rates are not unreasonable.

A.3 Oil block concessions data

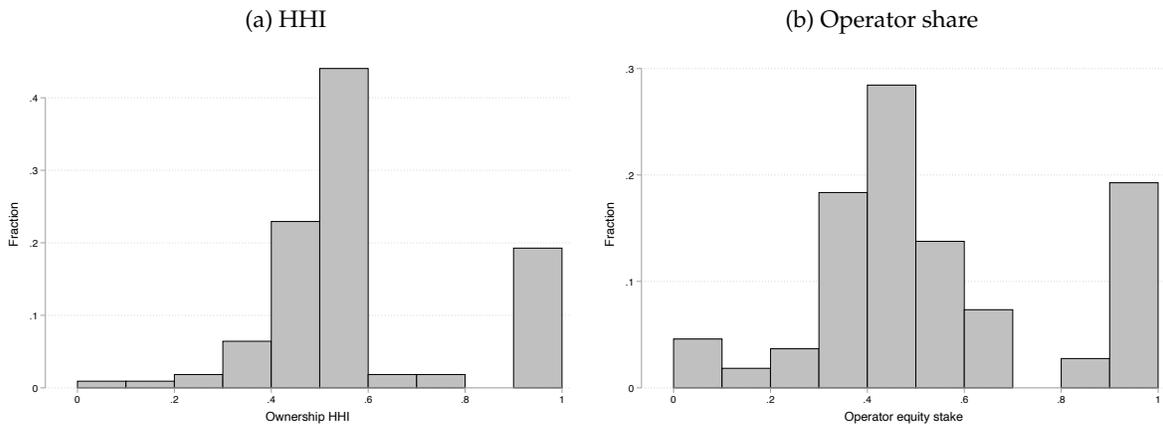
Concessions – large blocks of territory, typically containing several oilfields – are the primary unit of ownership in the Nigerian oil market. The exceptions to this rule are 30 “marginal” fields, which are independently-owned fields awarded to local operators that do not belong to larger concession 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 this data I obtain the ownership shares of all partners for all active oil mining leases, as well as the operating firm. Unfortunately, this detailed ownership data is only available from 2012-2018. I therefore exclude it from the main analysis and use it only to test mechanisms.

Figure A4 shows the block-level distribution of ownership shares in concessions as of 2016. Panel A shows the Herfindahl index of ownership, while Panel B gives the share owned by the primary operating firm. There is substantial variation in ownership structure and concentration across assets in the Niger Delta.

Figure A4: Ownership mixes in 2016



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

A.4 Treatment definition

Both main sources of data contain information on the identity of firms in control of oil fields. Still, both datasets contain substantial gaps and drawbacks; therefore, in my preferred specifications, I combine information from both sources to measure the localization “treatment” at the field-year level.

The DPR-NNPC dataset includes information on the firm operating each field in each year. However, as mentioned, this data has substantial gaps – many fields are missing information for years after they first enter the data. In addition, operatorship information is likely to lag

divestments, perhaps substantially given government data collection challenges. Therefore, it is difficult from this data alone to determine the exact year in which a given treatment occurs. Furthermore, using the operatorship measure exclusively overlooks cases in which local firms are non-operating shareholders, which may also be important.

In contrast, the DI data provides detailed information on a substantially wider set of transactions, covering all cases in which any ownership stake in a given field is transferred from a multinational to a local firm. It also contains precise information on the date of a transaction. However, it does not have information before 2006, so we cannot identify which fields are always-treated (that is, divested as of 2006) in our difference-in-differences setup using only the DI data. Furthermore, while it provides a wealth of transaction-level detail, it does not include comprehensive information on operatorship per se, simply changes in ownership.⁴⁵

However, when combined, the longer panel of the NNPC data can fill in the always-treated firms, while the DI data can provide more precise timing and information on non-operating divestments from 2006-2016. Furthermore, there will invariably be divestments missing in one dataset that are more likely to show up in the other. Using both data sources to generate the treatment indicator therefore provides the most detailed picture of local participation, defined as any amount of operator and/or ownership by an indigenous Nigerian firm. An indigenous Nigerian firm, in turn, is any firm that is headquartered in Nigeria and not majority-owned by a firm headquartered outside of Nigeria.⁴⁶

A detailed breakdown of the treatment definition is provided in Table A1. Firstly, I define a MNC-to-local “divestment” indicator from the DI data d_i^{DI} , which equals one for all fields that observed a transaction from 2006-2016 in which any buyer was Nigerian and any seller was multinational. Next, I create a dummy d_i^{NNPC} for all fields that are operated by local firm according to NNPC from 1998-2016. The first row of Table A1 corresponds to the 32 fields that experience a divestment from 2006-2016 according to both NNPC and DI records, so that $d_i^{DI} = 1$ and $d_i^{NNPC} = 1$. If DI and NNPC disagree on timing, these fields take as their treatment year the *earliest* year that local participation is reported by either source. These observations likely correspond to locally owned and operated fields. The next row indicates a further 24 fields that have any local participation from 2006-2016 according to DI data, but which are listed as solely multinationally operated over this period by NNPC, so $d_i^{DI} = 1$ and $d_i^{NNPC} = 0$. These are likely local firms that take stakes in multinationally operated fields without assuming operatorship. They take as their treatment year the closing date of the DI transaction, where available; otherwise the opening date.

The third and fourth rows, together, contain fields where $d_i^{DI} = 0$ and $d_i^{NNPC} = 1$, so the

⁴⁵In the deal description, there is mention of operatorship in only 68 of the 117 transactions.

⁴⁶In practice, this means that the local subsidiaries of oil supermajors in Nigeria – Eni, Total, Shell, Chevron, and ExxonMobil, as well as the Chinese Sinopec subsidiary Addax Petroleum – are classified as MNC; all others are indigenous. Note that the national oil company, the Nigerian Petroleum Development Company (NPDC) is considered an indigenous firm.

Table A1: Field counts by treatment type

	Number of fields	Share
Both	32	34.04
DrillingInfo only	24	25.53
NNPC only (always treated)	33	35.11
NNPC only (transition)	5	5.32
Total	94	

Table displays the number of fields that are marked as treatment by different data sources. Both refers to fields that are divested in both DI and NNPC data. DI only contains fields only divested in the DI data. NNPC only (always treated) contains fields that are treated in NNPC data before the DI data begins in 2006. NNPC only (transition) are fields that are divested between 2006-2016 in the NNPC but not DI data.

treatment year is taken from the NNPC data as the first year that the field has a local operator. These fall into two categories. First are the 33 always-treated fields, or those that have local participation from 1998-2005 according to NNPC, but where $d_i^{DI} = 0$ by construction. These are the fields that we know from NNPC must be local prior to the start of the DI data in 2006. However, since the NNPC data only contains operators, this may well omit some always-treated fields. That is, some of our eventually-treated (or even never-treated) fields may have actually been treated before 2006 through ownership and not operatorship. However, since 2010 marks the key inflection point in local participation (see Figure 1), this is likely to be at most a minor issue. Lastly, the fourth row contains the final 5 treated fields, which become treated between 2006-2016 according to NNPC but not DI. Reassuringly, this is a small number, which we should expect since DI is a more expansive dataset capturing both owner and operator transactions. In total, there are 94 ever-treated fields. The final treatment indicator $local_{it}$ is then equal to one for all field-years after the treatment year, following the “staggered adoption” difference-in-differences setup. If there is reverse divestments from locals to multi-nationals, we may wrongly assign treated field-years and therefore create measurement error biasing the effect toward zero. These transactions are exceedingly rate in the DI data, affecting less than 1% of field-years from 2006-2016.

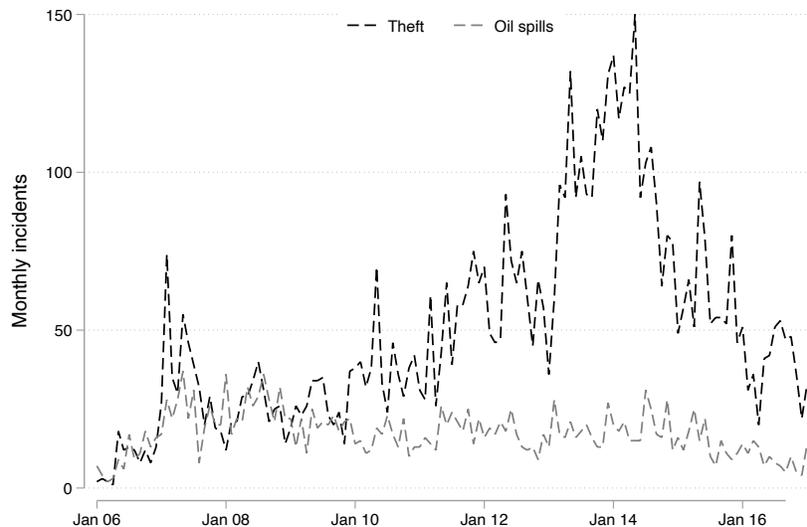
For the placebo tests, I define additional variables from the DI data, indicating a field’s exposure to local-to-local (27) and all other transactions (6). Finally, I construct an indicator of delayed divestments which equals one for all years after the announcement of an MNC-to-local divestment but before it’s consummation. In some cases, these are terminated/nullified transactions, while in others, this reflects a delay between the opening and closing dates. This indicator equals zero if and when the field eventually becomes “treated” according to the DI data. 36 fields are exposed to a delayed or terminated divestment in the sample period.

A.5 Oil spills 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.⁴⁷ 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.

Figure A5: 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 2006-2016. 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.

⁴⁷Rexer and Hvinden (2022) for a discussion about measuring oil theft.

Figure A5 charts the evolution of the black market by plotting the monthly incidents of pipeline sabotage and operational oil spills from 2006 through 2016. Oil spills due to theft rise dramatically from 2010-2014, and then fall thereafter. Oil spills due to operational failure, in contrast, decline over the whole period.

A.6 Conflict data

To measure violent conflict, 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 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. For each field, I aggregate the sum of annual attacks and fatalities of different types within 15 kilometers of the field centroid. ACLED event data are derived from news media reports (see for a discussion of the methodology), and report the journalistic source for each conflict event. In some cases, I subset to only conflict events reported by local news media sources in order to test for different sources of measurement error; see Appendix E.2 for a discussion of measurement error correlated with the divestment treatment.

A.7 Gas flaring data

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,⁴⁸ and contributes to climate change with CO₂ emissions. The practice has been subject to regulation since 1969, but meagre fines of 10 Naira per mscf flared (roughly 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.⁴⁹ Still, enforcement of this penalty is uneven, not least because companies under-report flaring volumes.⁵⁰

Data on gas flaring volumes comes from the Nigeria Gas Flare Tracker,⁵¹ a joint project by NOSDRA and the NGO Stakeholder Democracy Network. I download monthly panel data

⁴⁸See Ogunorisa (2009) for a review of studies on the negative impacts of Niger Delta flaring.

⁴⁹Current 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>

⁵⁰In 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.

⁵¹<https://nosdra.gasflaretracker.ng/>

on total gas flaring volume from March 2012 to May 2020, measured in thousands of cubic feet (mscf), for 210 flare sites. These location-specific volume estimates can then be converted to CO₂ emissions, since according to U.S. Energy Information Administration, flared natural gas emits 54.75 kg of CO₂ per mscf.⁵² 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 contained in the DPR/NNPC output data. These matched fields account for 93.4% of the flared gas volume over the period.

A.8 Law enforcement data

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 keywords⁵³, iii) mention at least one of a set of exact oil crime-related phrases.⁵⁴ Some examples of relevant articles are shown in Figure A6.

This procedure yields 17146 total articles potentially related to oil theft enforcement.⁵⁵ We then hired Nigerian research assistants to first identify all articles that are relevant to law enforcement activity in Nigeria, yielding a total of 3932.⁵⁶ From this set of relevant articles we then manually extract all *law enforcement events*, where an event is defined as a unique

⁵²https://www.eia.gov/environment/emissions/co2_vol_mass.php

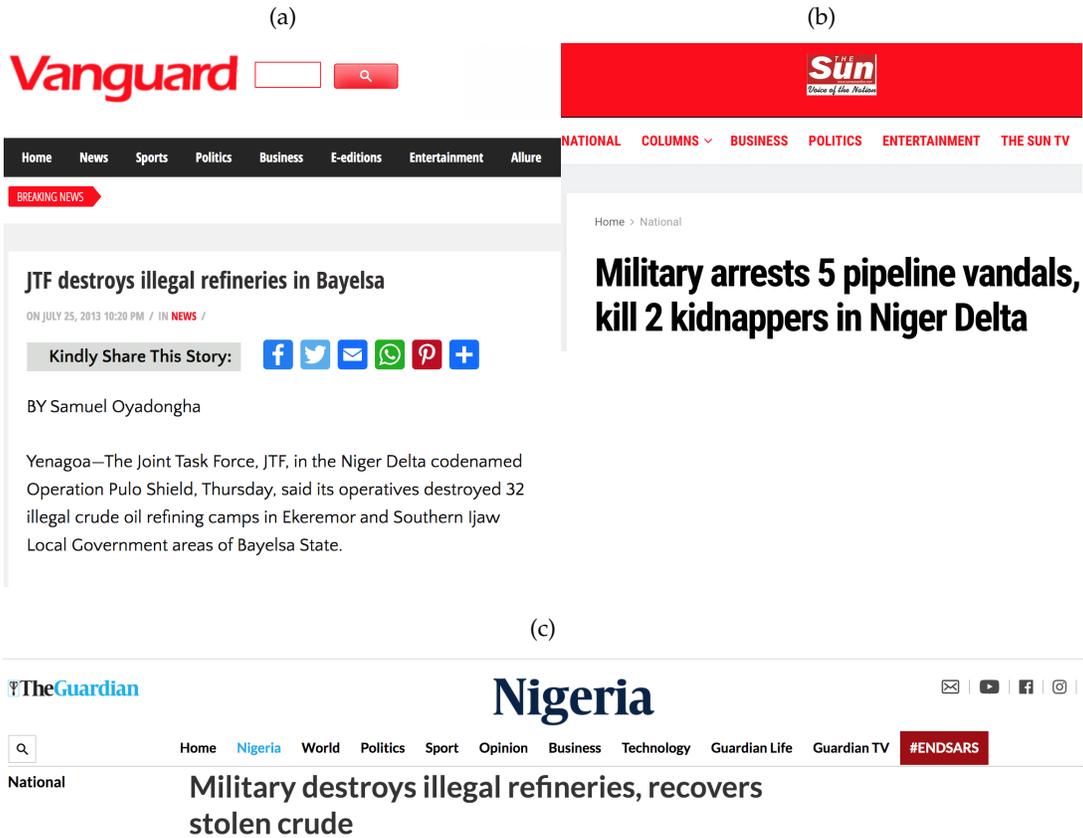
⁵³These are: "raid, raids, raided, seize, seized, seizes, seizure, seizures, destroy, destroys, destroyed, operation, capture, captures, captured, arrest, arrested, arrests, kill, killed, kills, apprehend, apprehended, apprehends, burn, burns, burned, invade, invaded, invades, search, searches, search"

⁵⁴These 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, pipeline vandalism, pipeline vandals, pirate, pirates, stolen crude, stolen, diesel, stolen oil, swamp buggy.

⁵⁵Of 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.

⁵⁶We 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

Figure A6: Sample relevant articles from local news media sources



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

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,⁵⁷ *iv*) the items seized or destroyed in the law enforcement action, selected from a pre-coded list,⁵⁸ *v*) the total number of arrests, and *vi*) the total number of fatalities. Extensive manual quality checks were conducted on weekly researcher submissions.

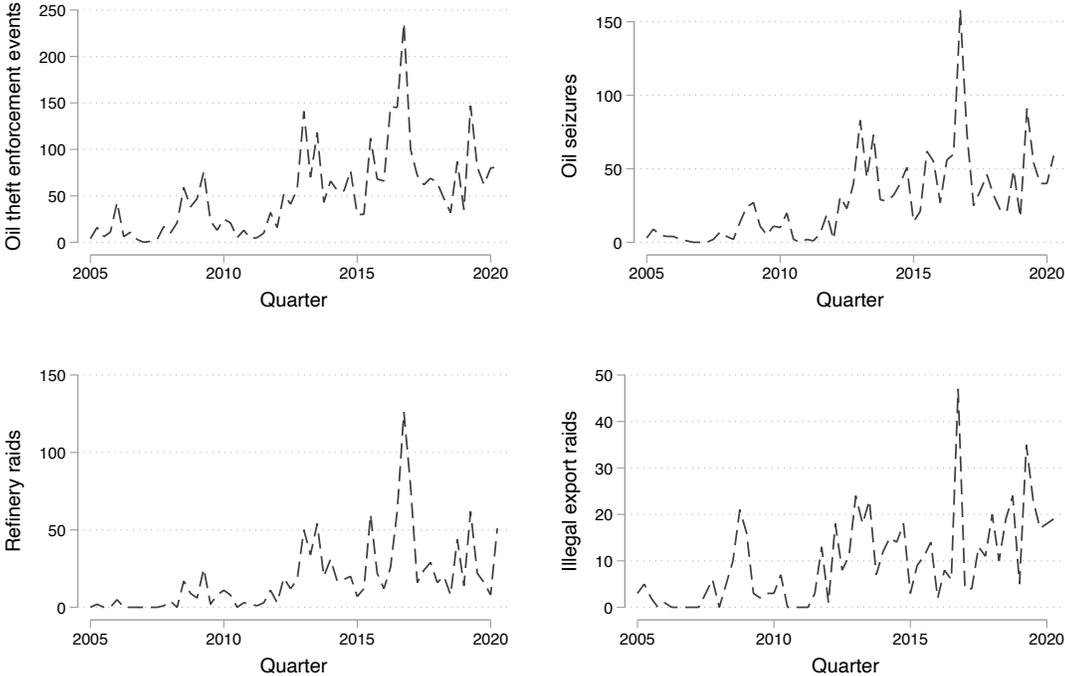
We consider any two articles as duplicates if they are published in the same calendar week and their headlines exceed a string similarity score threshold as defined by Levenshtein edit distance. After grouping duplicates into unique articles, we take the union of all events

⁵⁷These are: oil theft, piracy, illegal refining, pipeline vandalism, transportation of stolen oil, kidnapping, cultism/gang activity, militancy, and other illegal activities.

⁵⁸These are: no item, boats, stolen oil, arms, illegal refineries, trucks, oil theft equipment, and other items.

identified by the researchers to allow that duplicate articles may contain both repeated events as well as independent information.⁵⁹ 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 A7 plots quarterly total law enforcement actions for several different criminal activity categories.

Figure A7: Anti-oil theft enforcement



Note: This figure shows total quarterly oil theft-related law enforcement actions for the following categories of enforcement: all oil theft (top left), seizures of stolen petroleum products (top right), raids on illicit refineries (bottom left), and raids on vessels engaged in illegal export of stolen oil (bottom right).

⁵⁹For example, if there are two articles about the same raid, one may mention a second event, while the other does not.

A.9 Political connections data

To collect data on the political connections of firms in the Nigerian oil sector, I partner with a market research firm called Asoko Insight,⁶⁰ headquartered in London and Nairobi, and specializing in research on corporate entities in sub-Saharan Africa. In order to obtain a measure of firm-level political connections over time, we collected biographical information on the universe of corporate officers and shareholders in the sector. Below, I provide detailed steps that Asoko Insight uses to collect and assemble the biographical data that forms the basis of the firm-level political connections dataset.

1. Document retrieval: As a first step in the procedure, Asoko requests corporate filings from the Nigerian Corporate Affairs Commission (CAC) for a given firm. These documents contain all information that the firm has ever provided to CAC, from the date of incorporation to the present. For each firm, there is at minimum an initial filing at incorporation, detailing all of the shareholders and directors of the firm. In nearly all cases, firms will file subsequent additional reports of current shareholders and directors when there is a change in the company ownership and/or personnel.
2. Search report: The documents for a given firm are then compiled by Asoko Insight into a “search report.” This is a .pdf file that contains basic firm-level data as well as shareholder and director history, assembled from the corporate filings. Basic data include full name, date of incorporation in Nigeria, registered address, and issued share capital. Individual-level data includes name, address, status, and share allotment (if shareholder).
3. Transcription: Asoko then transcribes the search reports into two Excel files: one containing company-level data, and another containing person-company-filing-level data on directors and shareholders.
4. Cross-referencing: The names in the search report are then cross-referenced by Asoko against a list provided by the researcher. This list contains biographies of oil sector personnel that were able to be identified before engaging Asoko. These are taken as given and removed from the list of personnel for whom Asoko must obtain biographies. In addition, entity shareholders (e.g., holding companies) and foreign personnel are removed from the directors/shareholders list.
5. Ultimate beneficial owners (UBOs): We then take those shareholders that are themselves corporate entities and find additional information on their officers using ng-check.com – a public Nigerian corporate registry containing less detail than the CAC filings. We add all names associated with these firms to the list of personnel in 4), as “second-level

⁶⁰<https://www.asokoinsight.com/>

UBOs.” The result of steps 4 and 5 is a final list of Nigerian individuals associated with the firm for whom Asoko must obtain biographies.

6. Biographical research: The last step in the procedure is to obtain biographies for these individuals. Asoko’s Nigerian field researchers use a variety of methods, including desk research of all publicly available information in the local and international business press. In addition, researchers fill in gaps by employing key informant interviews that leverage their substantial corporate network in Nigeria. The result is an individual-level dataset containing the name and biographic details for all the unique individuals identified in steps 4 and 5.

In total, we obtain a list of 706 Nigerian nationals associated with 49 distinct corporate entities, covering all the firms listed in the DI and NNPC data. We are able to find biographic data for 552 of these individuals, implying an overall match rate of 78.2%. However, this rate varies by firm; Table A2 lists all of the firms included in the data collection procedure and their match rates. We further obtain data on 90 second-level UBOs who are not listed on corporate filings. I then code these biographies into the following indicator variables:

- Any politics: if the individual has any previous political activities in Nigeria.
- Elected politician: if the individual has ever held any elected office in Nigeria.
- Technocrat: if the individual has ever worked for the DPR, the NPDC, the NNPC, the Ministry of Petroleum Resources, or any other oil-related regulatory agency in the Nigerian Federal Government.
- High-level: if the individual has ever held a cabinet-level position in the Nigerian Federal Government.
- Security forces: if the individual has ever worked for the Nigerian Federal Police or in any military branch.
- Chief: if the individual holds any non-governmental, inherited, traditional title, e.g., the Oba of Benin or the Emir of Kano.

Using these dummy variables and the dates of the corporate filings as start and end dates of individual tenures, I transform the data into a firm-year panel, defining time-varying dummy variables for each connection types at the firm-level. Finally, I match this panel to the field-level data using the operators indicated in the NNPC data. Note that the political connections variables vary over-time at the field level because of both turnover in firm-level personnel and changes in field-level ownership.

Table A2: Coverage rates by firm for biographical data

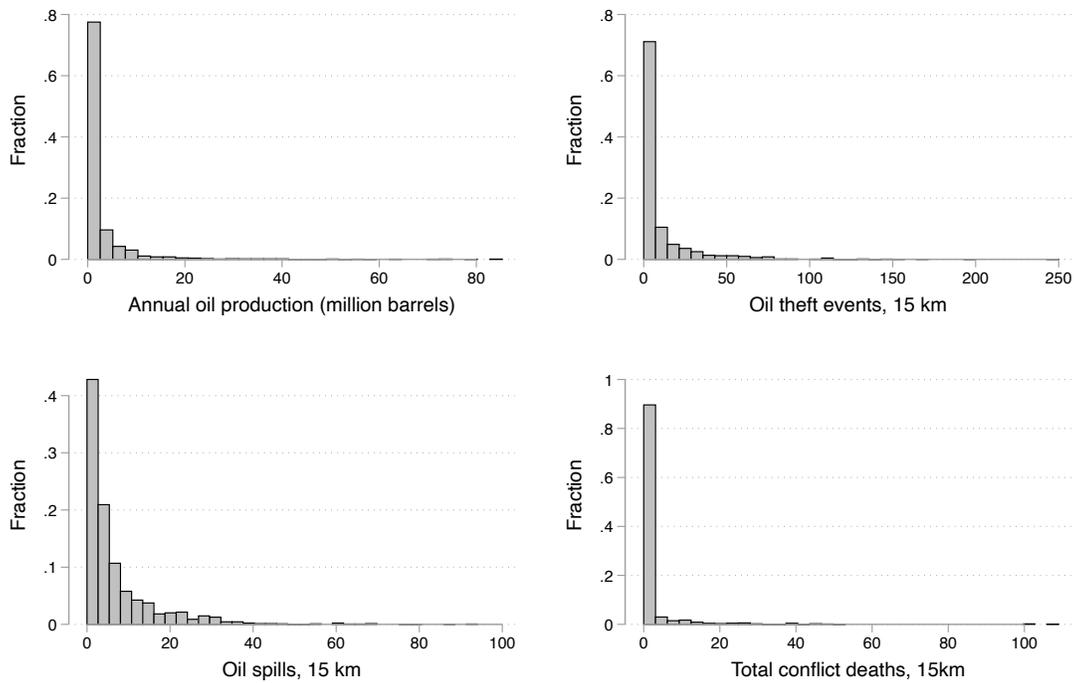
Firm	Officers	Foreign	Nigerian	Coverage
ORIENTAL ENERGY	10	0	10	1.00
TEXACO OVERSEAS (NIGERIA) PETROLEUM CO LTD	29	24	5	1.00
MIDWESTERN OIL & GAS CO LTD	15	1	14	1.00
NECONDE ENERGY LTD	19	5	14	1.00
FIRST HYDROCARBON NIGERIAN LTD	14	1	13	1.00
SHORELINE NATURAL RESOURCES LTD	7	4	3	1.00
ND WESTERN	11	3	8	1.00
NETWORK EXPLORATION AND PRODUCTION NIGERIA LTD	9	0	9	1.00
SOUTH ATLANTIC PETROLEUM LTD	17	7	10	1.00
ELCREST EXPLORATION AND PRODUCTION NIGERIA LTD	13	3	10	1.00
EROTON EXPLORATION AND PRODUCTION CO LTD	14	2	12	1.00
ESSO EXPLORATION AND PRODUCTION NIGERIA LTD	42	36	6	1.00
YINKA FOLAWIYO ENERGY LTD	4	0	4	1.00
ATLAS PETROLEUM INTERNATIONAL LTD	8	1	7	1.00
NEWCROSS PETROLEUM LTD	9	0	9	1.00
NIGERIAN AGIP OIL CO LTD	14	11	3	1.00
AITEO EASTERN E&P CO LTD	3	0	3	1.00
CAVENDISH PETROLEUM NIGERIA LTD	4	0	4	1.00
MOBIL PRODUCING NIGERIA UNLTD	55	30	25	0.96
ENERGIA LTD	20	1	19	0.95
PRIME EXPLORATION AND PRODUCTION LTD	16	0	16	0.94
TOTAL NIGERIA PLC	94	72	22	0.91
PLATFORM PETROLEUM LTD	20	0	20	0.90
SHELL PETROLEUM DEVELOPMENT CO OF NIGERIA LTD	90	58	32	0.88
STAR DEEP WATER PETROLEUM LTD	35	19	16	0.88
NEWCROSS EXPLORATION AND PRODUCTION LTD	7	0	7	0.86
CHEVRON NIGERIA LTD	59	38	21	0.86
SHELL NIGERIA EXPLORATION AND PRODUCTION CO LTD	50	26	24	0.83
NIGER DELTA PETROLEUM RESOURCES LTD	6	0	6	0.83
PAN OCEAN OIL CORPORATION (NIGERIA) LTD	20	14	6	0.83
SEPLAT ENERGY PLC	25	10	15	0.80
FRONTIER OIL LTD	14	0	14	0.79
ADDAX PETROLEUM	37	28	9	0.78
TOTAL UPSTREAM NIGERIA LTD	58	50	8	0.75
AMNI INTERNATIONAL PETROLEUM CO LTD	17	5	12	0.75
EXPRESS PETROLEUM & GAS CO. LTD	7	0	7	0.71
BRITANIA-U NIGERIA LTD	21	4	17	0.71
NIGERIAN PETROLEUM DEVELOPMENT CO LTD	53	0	53	0.68
CONTINENTAL OIL AND GAS LTD	9	0	9	0.67
UNIVERSAL ENERGY RESOURCES LTD	34	3	31	0.65
WALTER SMITH PETROLEUM OIL LTD	24	5	19	0.63
CONOIL PLC	99	48	51	0.61
PILLAR OIL LTD	24	0	24	0.58
ALLIED ENERGY PLC	30	6	24	0.54
AITEO EXPLORATION AND PRODUCTION CO LTD	6	0	6	0.50
MONI PULO LTD	20	0	20	0.50
CAMAC NIGERIA LTD	17	4	13	0.38
DUBRI OIL CO LTD	16	0	16	0.31
STERLING OIL EXPLORATION & ENERGY PRODUCTION CO LTD	5	5	0	
ALL	1230	524	706	0.78

Table shows counts of total unique identified officers (boardmembers and shareholders), foreign officers, Nigerian officers, and the coverage rate, by firm. The coverage rate is the share of Nigerian officers for whom biographical information can be found.

A.10 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 violent conflict, piracy, oil theft, oil spills, 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. Therefore, while the estimation sample for all non-production outcomes is 3,183 field-years, the sample for regressions in which production is the outcome falls to only 2476 field-years.⁶¹

Figure A8: Histogram of main outcomes



Note: Figure shows histogram of main outcomes for the primary estimation sample of field-years from 2006-2016. Oil theft events are measured as all sabotage-related oil spills within 15 km of the oilfield. Oil spills are measured as all malfunction-related oil spills within 15 km of the oilfield. Conflict deaths are all conflict deaths within 15 km of the field.

Figure A8 provides histograms for the main outcomes used throughout the paper – oil output, oil spills, oil theft, and conflict deaths – in our estimation sample. All of the outcome variables are non-negative and right-skewed with long tails and a mass point at zero. This suggests a role for exponential models, explored in Appendix C.2.

⁶¹I choose not to restrict the sample for all estimation in order to make full use of available data for non-production outcomes. However, I test robustness of results to the restricted sample in Appendix C.4, and consider an imputation procedure for output in Appendix C.5.

B Theoretical appendix

B.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 with probability λ if the behavior is discovered. For simplicity, normalize $\lambda = 1$. 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 of Q , to capture the important role of joint-ventures in Nigeria, as shown in Figure A4. Importantly, law enforcement may internalize firm f 's output based on the parameter μ . This measures how the strength of political connections determines enforcement behavior. If a firm is unconnected, then $\mu = 0$. Note that $\gamma + \mu \leq 1$

B.2 Base case

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

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

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 $S(0) = p(\bar{Q} - k) - c$.⁶² 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

$$\bar{b}_g = (1 - \alpha)(pq - c + \epsilon_g) \quad \bar{b}_f = (1 - \alpha)\gamma p(q + k) - \lambda$$

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 $b_g = \bar{b}_g$.⁶³ 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

⁶²Note 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.

⁶³Note 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

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

Proposition 1. Comparative statics: enforcement and theft. Given $\gamma + \mu \leq 1$ 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:

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

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

B.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.⁶⁴ 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 $U_f^0 - \lambda - b_f < U_f^0$. 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

⁶⁴I 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 renegeing 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 $U_f^0 - \lambda - b_f < U_f^0$. 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 $\bar{\mu} > 0$. \square

B.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:

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

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

$$r_f \leq \frac{\gamma p\bar{Q} - \lambda - b_f}{1 - \delta}$$

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

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

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:

$$\gamma p\bar{Q} + \delta r_f \leq \frac{\gamma p\bar{Q} - \lambda - b_f}{1 - \delta}$$

Yielding a maximal willingness to pay:

$$\tilde{b}_f = \delta \gamma p(q + \kappa) - \lambda$$

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

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

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

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 $\delta > 0$. \square

C Robustness tests: main outcomes

C.1 Robustness to confounders

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, though the event-studies in Figure 3 and OA D.5 do not suggest this to be the case. To further control for differential location-specific time trends, I include location-by-year interacted fixed effects, using both states and municipalities as larger and smaller geographic areas. The resulting specification essentially compares fields within a given location, with the final estimate a weighted average across localities of these within-locality comparisons. Note that this also controls for electoral cycles and changes in political leadership at the local level. I also estimate specifications that include field-specific linear time trends. Table A3 presents the results for each of the main outcomes. The results are all qualitatively similar to Table 2, with the exception of column (4); the effect of divestment on oil spills is no longer positive after including field-specific time trends.

Table A3: The effect of divestment, robustness to fixed effects

Outcome	Output			Oil spills			Oil theft		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Local firm	0.848* (0.447)	1.054*** (0.390)	0.684 (0.517)	-0.570 (0.695)	1.480 (0.995)	2.595 (1.735)	-6.362** (2.531)	-2.932*** (0.940)	-3.845** (1.730)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Field FE \times time trend	Yes	No	No	Yes	No	No	Yes	No	No
State \times Year FE	No	Yes	No	No	Yes	No	No	Yes	No
Locality \times Year FE	No	No	Yes	No	No	Yes	No	No	Yes
Controls \times Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2476	2476	2476	3183	3183	3183	3183	3183	3183
R^2	0.915	0.885	0.911	0.744	0.692	0.790	0.813	0.786	0.861

Standard errors in parentheses 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 spills are the total number of non-sabotage spills within 15 km of the field. Oil theft is the total number of sabotage spills within 15 km of the field. Localities are local government areas, the lowest level administrative unit in Nigeria. Time trend is a linear trend. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Another threat to identification is that there may be selection into field takeover based on field characteristics. Table 1 demonstrates that localized fields are younger, smaller, and more likely to 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 A4, 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 2,392 field-years for output and 3,038 for other outcomes because 17

fields have missing characteristics. Despite this, the results are unchanged.

Table A4: The effect of divestment, robustness to field-level covariates

Outcome	Output		Oil spills		Oil theft	
	(1)	(2)	(3)	(4)	(5)	(6)
Local firm	0.707** (0.291)	0.882** (0.356)	1.735** (0.732)	1.524* (0.785)	-6.831*** (1.038)	-6.579*** (1.101)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Main controls × Year FE	No	Yes	No	Yes	No	Yes
Field controls × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2392	2392	3038	3038	3038	3038
R ²	0.870	0.884	0.624	0.669	0.737	0.763

Standard errors in parentheses 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 spills are the total number of non-sabotage spills within 15 km of the field. Oil theft is the total number of sabotage spills within 15 km of the field. Main 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A5: The effect of divestment on output and criminality, robustness to prices

Outcome	Output		Oil spills		Oil theft	
	(1)	(2)	(3)	(4)	(5)	(6)
Local firm	0.702* (0.360)	0.598* (0.333)	1.047 (0.824)	0.287 (0.890)	-5.526*** (0.891)	-6.793*** (1.173)
Treated × Oil price (USD/barrel)	-0.015 (0.010)	-0.028*** (0.011)	-0.031** (0.013)	-0.044*** (0.013)	-0.068** (0.029)	-0.097*** (0.032)
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	2476	2476	3183	3183	3183	3183
R ²	0.861	0.878	0.590	0.651	0.713	0.757

Standard errors in parentheses 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 spills are the total number of non-sabotage spills within 15 km of the field. Oil theft is the total number of sabotage spills within 15 km of the field. Main 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 measured as the annual average world crude oil price in dollars per barrel. Field controls are number of wells, initial year, onshore dummy, and maximum well depth. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A central threat to the assumption of parallel trends is the existence simultaneous shocks that have differential effects on local and multinational firms. The most obvious of these is oil price shocks, which may have outsized impacts on the production of more capital-constrained

local firms. I test robustness of the main results to differential oil price effects in Table A5 by including 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 shocks among localized fields are driving the results.

Rexer and Hvinden (2022) show that the 2009 amnesty for Niger Delta militants reduced violence and increased oil theft differentially in conflict-affected 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 A6 by including the interaction between a post-amnesty indicator and distance to the nearest militant camp in the main TWFE model. The results are unaffected.

Table A6: The effect of divestment on output and criminality, robustness to amnesty policy

Outcome	Output		Oil spills		Oil theft	
	(1)	(2)	(3)	(4)	(5)	(6)
Local firm	0.775** (0.324)	0.969***	1.281* (0.760)	0.938 (0.804)	-5.234*** (0.845)	-5.196*** (1.026)
Distance to militant camp (km) \times Post-amnesty	-0.027*** (0.010)	-0.002 (0.012)	-0.041 (0.026)	-0.073** (0.032)	-0.179*** (0.023)	-0.246*** (0.044)
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	2476	2476	3183	3183	3183	3183
R^2	0.863	0.878	0.592	0.651	0.721	0.761

Standard errors in parentheses 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 spills are the total number of non-sabotage spills within 15 km of the field. Oil theft is the total number of sabotage spills within 15 km of the field. Main controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. Amnesty date is 2009. Field controls are number of wells, initial year, onshore dummy, and maximum well depth. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Another source of endogeneity is that firm rather than field-level characteristics may be highly correlated with localness, and thus drive the results. The clearest example here is firm size – oil theft gangs may target the assets of deep-pocketed larger firms because sabotage threats are more likely to generate direct payments, or because larger firms can finance operations under difficult conditions for longer. At the same time, we know that multinationals are much larger, on average, than local firms. The local advantage may therefore be a “small firm effect” rather than a local effect. This is inherently difficult to test, since there is a high degree of correlation between overall firm size and indigeneity. As such, in Table A7 I control for firm size using measures that capture the size of multinational *subsidiaries*, rather than parent firms. In particular, I calculate the (log of) total fields or wells operated by the operating firm of field i at time t . This is a time-varying, field-specific characteristic and so is not absorbed by fixed effects.

Panel A estimates the impacts on output while Panel B looks at oil theft. Columns (1)-(3)

Table A7: The effect of divestment on output and criminality, robustness to firm size

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Output</i>						
Local firm		0.945*** (0.345)	0.902*** (0.338)		0.859** (0.355)	0.841** (0.351)
Log number of wells	-0.045 (0.193)	0.114 (0.209)	-0.114 (0.148)			
Log number of fields				-0.238 (0.270)	-0.054 (0.294)	-0.246 (0.194)
Observations	2476	2476	2476	2476	2476	2476
R ²	0.861	0.861	0.878	0.861	0.861	0.878
<i>Panel B: Oil theft</i>						
Local firm		-5.160*** (0.863)	-6.256*** (1.086)		-5.303*** (0.877)	-6.182*** (1.083)
Log number of wells	-0.190 (0.229)	-0.869*** (0.286)	-1.275*** (0.412)			
Log number of fields				-0.246 (0.310)	-1.252*** (0.419)	-1.221*** (0.416)
Observations	3183	3183	3183	3183	3183	3183
R ²	0.710	0.713	0.756	0.710	0.713	0.756
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	No	Yes	No	No	Yes

Standard errors in parentheses 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 spills are the total number of non-sabotage spills within 15 km of the field. Oil theft is the total number of sabotage spills within 15 km of the field. Main controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital. Firm size at the field-level is measured as the log of the total number of fields or wells owned by the operating firm. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

use the log number of wells, while columns (4)-(6) use the log number of fields, which is a more aggregate measure. The main results all hold. Furthermore, columns (2)-(3) and (5)-(6) of Panel B suggest that, after conditioning on local status, the assets of larger firms actually appear to be targeted relatively *less*, an association which is significant at 1%. This may be because, all else equal, larger firms have more capital to invest in security for their assets, either directly or through protection rackets.

C.2 Maximum likelihood estimation

All of the outcome variables are bounded below by zero, and both oil spills and oil theft take the form of count data. As such, it may be reasonable to consider that a linear model estimated using TWFE is misspecified, and that a non-linear exponential model may improve specification. I estimate an exponential conditional mean model without additional distributional assumptions using the quasi-maximum likelihood Poisson regression methods developed in Wooldridge (1999). Given our identification strategy, this requires conditional likelihood optimization with high-dimensional, multiway fixed effects, a non-trivial computational problem. I follow the implementation of Correia et al. (2020) for high-dimensional quasi-MLE with exponential Poisson models.

Table A8: The effect of divestment: Poisson model

Outcome	Output		Oil spills		Oil theft	
	(1)	(2)	(3)	(4)	(5)	(6)
Local firm	0.311* (0.184)	0.471*** (0.172)	0.069 (0.189)	0.326** (0.160)	-0.347*** (0.092)	-0.444*** (0.105)
Average marginal effect	0.859* (0.508)	1.301*** (0.475)	0.465 (1.277)	2.203** (1.079)	-3.302*** (0.875)	-4.222*** (1.001)
Observations	2476	2476	3183	3183	3183	3183
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 parentheses are clustered at the field level. Sample is the panel of 314 oil-fields from 2006-2016. Output is measured in millions of barrels of oil per year. Oil spills are the total number of non-sabotage spills within 15 km of the field. Oil 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. All models are estimated with Poisson pseudo-maximum likelihood (PPML) regression. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The results are in Table A8, which contains both the estimated coefficients from the Poisson regression and the average marginal effects for all three outcomes. The estimates are all correct directionally, mostly all significant. More importantly, the magnitudes of the average marginal effects are very similar to the linear effects estimated in Table 2.

C.3 Inference

Many of the divestments occur not necessarily at the field but rather the oil block-level. An oil block is a geographic ownership unit (concession) in the Nigerian oil sector that typically contains some number of fields. However, assets within an oil block can be individually

divested. Furthermore, many locally-owned assets take the form of marginal fields, which while geographically part of an oil block are not necessarily owned by firm with the mineral rights to that block. As such, the level of treatment is not exactly obvious – treatment status is likely to be correlated across fields within blocks, though not perfectly so. To account for the impact of this on inference, Table A9 reproduces the main results clustering at the oil block level rather than the field level.⁶⁵ The results for output and theft remain significant at the 1% level. However, the results for oil spills are no longer significant.

Table A9: The effect of divestment on output and criminality, block-level clustering

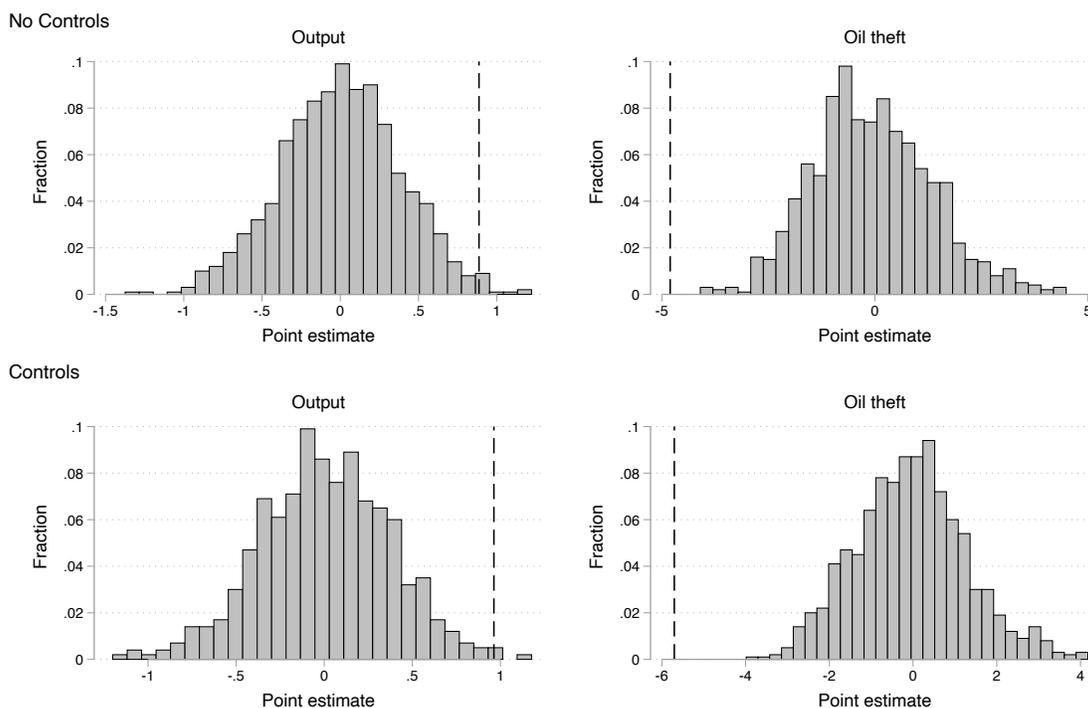
Outcome	Output		Oil spills		Oil theft	
	(1)	(2)	(3)	(4)	(5)	(6)
Local firm	0.887*** (0.313)	0.963*** (0.286)	1.380 (1.197)	0.788 (1.205)	-4.805*** (1.353)	-5.703*** (1.901)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	No	Yes	No	Yes
Control group mean	2476	2476	3183	3183	3183	3183
Observations	0.861	0.878	0.590	0.649	0.712	0.756

Standard errors in parentheses are clustered at the oilblock level. Sample is the panel of 314 oilfields from 2006-2016. Output is measured in millions of barrels of oil per year. Oil spills are the total number of non-sabotage spills within 15 km of the field. Oil theft is the total number of sabotage spills within 15 km of the field. Main controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

I also consider the robustness of the significance levels of the effects on main outcomes (output and theft) to randomization inference. I generate 1000 independent permutations of treatment status and treatment dates (conditional on status) across all field-years in the data. Each of these alternate treatment timing vectors preserves the overall proportion of treated and untreated field-years, as well as the proportion of fields treated in any individual event-year. I then re-estimate the main TWFE models with and without controls for each of these placebo treatment assignments. Figure A9 plots the empirical approximation to the sampling distribution of treatment effects generated by this procedure. Vertical lines indicate the effect from the “true” assignment. The randomization inference procedure produces well-behaved, approximately-normal sampling distributions both without controls (top panel) and with them (bottom panel), for both outcomes. The location of the true estimates in these distributions implies p -values that attain 1% significance, consistent with the standard inference results of Table 2.

⁶⁵This is a substantially higher level of clustering – there are 314 fields in our data, and only 80 oil blocks.

Figure A9: Randomization inference



Note: Figure shows histograms of coefficient estimates for 1000 independent draws of a randomization inference routine. Randomization inference permutes the date of treatment assignment across all observations in the data, holding the share of each event-year and the untreated fixed at their empirical levels. Outcome variable is oil output in millions of barrels (left panel) or oil theft, measured as total number of sabotage spills within 15 km of the field (right panel). Top panel estimates models without controls, while bottom panel controls for latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital interacted with year fixed effects. Vertical line indicates the estimate for the observed data.

C.4 Estimation sample

The estimation sample for columns (1)-(2) of Table 2 consists only of field-years with non-missing oil production. However, the sample of columns (3)-(6) consists of all field-years after a field first enters the data. To harmonize the sample, Table A10 tests whether the results for oil theft and oil spills hold in the subsample of non-missing observations (the “oil sample”). The results in columns (1)-(2) suggest that the effect on oil theft is in fact even larger when we restrict to the oil sample. However, the effect on oil spills in columns (4)-(5) is smaller and insignificant.

Columns (3) and (6) replicate the results of Table 4 by interacting the treatment variable with the onshore indicator. The results are the same as the main sample – the local oil theft advantage is concentrated in risky onshore assets, while the increase in spills is concentrated

in complex offshore assets. These differential effects remain statistically significant at the 1% and 10% level, respectively.

Table A10: The effect of divestment on theft and oil spills, oil output sample

Outcome	Oil theft			Oil spills		
	(1)	(2)	(3)	(4)	(5)	(6)
Local firm	-5.279*** (1.000)	-6.491*** (1.217)	0.094 (0.141)	0.840 (0.820)	0.490 (0.908)	5.538** (2.233)
Local firm × Onshore field			-7.403*** (1.445)			-4.605* (2.376)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	Yes	No	Yes	Yes
Observations	2464	2464	2464	2464	2464	2464
R ²	0.723	0.769	0.787	0.598	0.660	0.684

Standard errors in parentheses are clustered at the field level. Sample is all field-year observations for which oil output is nonmissing. Oil spills are the total number of non-sabotage spills within 15 km of the field. Oil theft is the total number of sabotage spills within 15 km of the field. Main controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C.5 Output imputation

An alternative method of harmonizing the sample is to fill in missing output values. Given gaps in government administrative data reporting, there are 707 field-years with missing output data. These missing values are unlikely to be completely random, and may be correlated with field-level covariates. If that is the case, then the main results on output may be affected by sample selection bias, leading to incorrect inferences. I address this concern by using regression imputation to estimate the expected value of output in missing field years conditional on characteristics. I model the conditional expectation for output as a function of the interaction between field-specific covariates and time trends

$$E[y_{it}|X_i, t] = g \left(X_i' \beta + \sum_{\tau} 1(t = \tau) \times (\delta_{\tau} + X_i' \beta_{\tau}) \right)$$

That is, missing values are predicted by the time trend, which is allowed to vary with field characteristics, X . This covariate vector includes the main controls – latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital – as well as distance to a militant camp and onshore. Unfortunately, I must omit field-level variables used in Table A4 because they are also missing for 17 fields and so cannot generate complete

predictions. After estimating the conditional expectation, I use these predictions to “fill in” missing values in output. All non-missing values remain as in the original data. Table A11 then re-estimates the main output results using imputed values. Column (1) contains the main result on the nonmissing subsample for comparison.

Table A11: The effect of divestment on output: regression imputation

Imputation	None	Linear	Log-linear	Log-link
	(1)	(2)	(3)	(4)
Local firm	0.887*** (0.312)	1.012*** (0.275)	0.951*** (0.290)	0.730*** (0.268)
Field FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	2476	3183	3183	3183
R^2	0.861	0.799	0.793	0.798

Standard errors in parentheses 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. Predictor variables for each model are latitude of the field centroid, distance to coast, distance to Niger River, distance to the capital, distance to militant camps, and an onshore indicator, interacted with year fixed effects. Imputed values are predictions from the model (indicated in table header) for observations for which output is missing. Linear model is an OLS regression. Log-linear is an OLS regression with log output as the outcome variable, where predicted values are $\exp(E[\log(y)|X] + 0.5\sigma^2)$. Log-link estimates a generalized linear model with a log link function and a gamma distribution for y .
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Of course, a primary issue is the choice of functional form for g . In column (2), I use a simple linear form estimated with OLS for the imputation. This produces estimates slightly larger than in (1). However, a linear functional form has the drawback of predicting negative levels of output for a non-negative variable. Given that the predicted outcome must be weakly positive, an exponential functional form for g is more appropriate. One approach is to run a log-linear regression, and then calculate the expected values using smearing transformation, that is, $\hat{y} = \exp\left(E[y_{it}|X_i, t] + \frac{\sigma^2}{2}\right)$. This approach assumes homoskedastic errors and a lognormal distribution of y , so that the regression errors are normally distributed. The results are in column (3), and are very similar to the main specification in (1). However, this approach drops observations with zero production in estimating the parameters of the conditional expectation function. Another approach, used in (4), is to estimate a generalized linear model with a logarithmic link function (which generates an exponential conditional mean), with a gamma distribution, since y is a continuous variable. The results are slightly smaller than in (1), yet still strongly positive and significant.

C.6 Outliers

Figure A8 shows that all of the main outcome variables are long-tailed and right-skewed. This suggests that outliers could be driving the main results. To address this possibility, I Winsorize (top-code) the outcome at the 95th or 99th percentile of the empirical distribution in the full sample. The 95th percentile top-coding somewhat reduces the size of the output and theft effects, though it increases the size of the oil spills effect. All of the estimates remain significant at the 10% level or lower.

Table A12: The effect of divestment: robustness to outliers

Winsorized at	95		99	
	(1)	(2)	(3)	(4)
<i>Panel A: Output</i>				
Local firm	0.543** (0.251)	0.493* (0.259)	0.869*** (0.300)	0.817*** (0.291)
Observations	2476	2476	2476	2476
R ²	0.852	0.857	0.873	0.880
<i>Panel B: Oil spills</i>				
Local firm	1.366** (0.537)	1.435** (0.605)	1.309* (0.685)	1.040 (0.768)
Observations	3183	3183	3183	3183
R ²	0.687	0.729	0.656	0.705
<i>Panel C: Oil theft</i>				
Local firm	-2.997*** (0.562)	-2.996*** (0.648)	-4.463*** (0.751)	-5.106*** (0.915)
Observations	3183	3183	3183	3183
R ²	0.777	0.816	0.756	0.796
Field FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	No	Yes

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. Outcomes are winsorized at the 95th or 99th percentile of the distribution, as indicated in table header. Output is measured in millions of barrels of oil per year. Oil spills are the total number of non-sabotage spills within 15 km of the field. Oil 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C.7 Main estimates by distance

The main theft outcome variable is defined as the number of theft incidents within 15km of the field centroid. This relatively wide radius is to account for the fact that oilfields can be large, and even beyond the boundaries of the field there may be significant pipeline infrastructure vulnerable to theft. However, using a wide radius also introduces estimation challenges. First, this may induce mechanical spatial correlation between nearby fields, biasing inference. Second, we may be capturing theft on infrastructure not owned/operated by the firm. If this measurement error is systematically related to localization, it risks biasing the results.

Table A13: The effect of divestment: robustness to distance radius

Outcome	Oil theft					
	(1)	(2)	(3)	(4)	(5)	(6)
Local firm	-0.735*** (0.281)	-2.488*** (0.599)	-5.703*** (1.023)	-10.843*** (1.578)	-14.236*** (2.067)	-14.807*** (2.429)
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
Distance (km)	5	10	15	20	25	30
Control group mean	1.768	5.334	10.608	18.233	26.378	34.720
Observations	3183	3183	3183	3183	3183	3183
R ²	0.629	0.730	0.756	0.775	0.781	0.806

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. Oil theft is the total number of sabotage spills within a given distance radius of the field centroid. Distance radius is indicated in table footer. 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 A13 shows the robustness of the results to different distance radii, from 5-30 km in intervals of 5km. The control group mean is given in each specification so that the results can be rescaled. The results are negative, large, and significant at 1% in all specifications. It does not appear that is spurious correlations or measurement error induced the definition of oil theft is driving our results.

C.8 Revenue effects

Table A14 estimates the impact of divestment on field-level revenue by multiplying output by the global oil price. Columns (1) and (2) show a 36-82 million USD per year increase in per-field revenue, equivalent to a 15-35% increase on the control group pre-divestment mean.

Table A14: The effect of divestment on revenue

Outcome	Revenue (millions of USD)		log(Revenue)	
	(1)	(2)	(3)	(4)
Local firm	81.739*** (27.072)	35.934 (29.252)	0.474** (0.185)	0.551*** (0.192)
Field FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Controls \times Year FE	No	Yes	No	Yes
Control group mean	229.354			
Observations	2476	2476	1899	1899
R^2	0.834	0.851	0.764	0.778

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016 for which output information is available in (1)-(2) and nonzero in (3)-(4). 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C.9 Treatment definition

The main results in Table 2 include all non-multinational firms in “local.” In Table A15 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 positive effect on output is driven almost exclusively by private firms. In contrast, the efficiency costs of localness in terms of greater malfunctions essentially vanishes when we disaggregate the treatment, with a small insignificant point estimate, while the effect size rises to 3.9 for state-run fields. At the same time, the reductions in oil theft is 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 minimal, resulting in a smaller output effect.

As mentioned in Section 3, the data on local firm participation at the oilfield level comes from two sources. The first is the administrative records of the NNPC, which records the operating firm of each oilfield-year. The second source is the DrillingInfo data on corporate transactions. This data provides more detail on local firms’ stakes in oilfields, and helps fill in the substantial gaps in the NNPC data. However, it does not distinguish between divestments of operatorship or ownership. Throughout the paper I use a conservative approach that leverages all of the information in both of these datasets, defining treatment as all field-years with any local participation in either dataset.

Table A16 investigates the implications of different treatment definitions for the results.

Table A15: The effect of divestment on output and criminality, public and private

Outcome	Output		Oil spills		Oil theft	
	(1)	(2)	(3)	(4)	(5)	(6)
Private local operator	0.910*** (0.312)		0.390 (0.809)		-5.619*** (1.055)	
Government operator		0.216 (0.489)		3.939*** (0.665)		0.867 (1.969)
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	2476	2476	3183	3183	3183	3183
R^2	0.878	0.877	0.649	0.650	0.756	0.753

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. 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 NNPC/NNPC in a given field-year. Output is measured in millions of barrels of oil per year. Oil spills are the total number of non-sabotage spills within 15 km of the field. Oil 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Column (1) reprints the main results for reference. Column (2) uses a similar treatment definition as the main specification, leveraging information from both sources. However, when the two sources disagree on the event-year of a given divestment, it takes the DrillingInfo date, since companies may be updated in the NNPC data with a lag whereas DI contains precise dates. This changes treatment status for only a handful of field-years and so does not materially affect any of the results. Column (3) ignores all NNPC data and uses only divestments mentioned in DI from 2006-2016, which yields 56 ever-treated fields. The results remain similar, though the local output advantage falls slightly and the quality disadvantage rises. Finally, column (4) uses only changes in operatorship that are identified in the NNPC administrative data, which gives 70 ever-treated. The results here are still qualitatively similar. However, the output advantage is over twice as large as in (1), while oil theft reduction falls slightly. Both remain significant at 1%.

C.10 Oil spills output adjustment

The results of Table 2, column (3)-(4) may not be driven by lower operating quality among local firms. Instead, the increase in per-field output may mechanically be driving up equipment failure-related oil spills. As a robustness test, I consider whether the increase in output is large enough to explain the effect on oil spills. Let γ denote the marginal effect of an ad-

Table A16: The effect of divestment, treatment definition

Treatment	Main	DI first	DI only	NNPC only
	(1)	(2)	(3)	(4)
<i>Panel A: Output</i>				
Local firm	0.887*** (0.312)	0.867*** (0.309)	0.852*** (0.327)	2.181*** (0.609)
Observations	2476	2476	2476	2476
R ²	0.861	0.861	0.861	0.862
<i>Panel B: Oil spills</i>				
Local firm	1.380* (0.767)	1.395* (0.758)	1.600** (0.791)	1.431 (0.927)
Observations	3183	3183	3183	3183
R ²	0.590	0.590	0.590	0.589
<i>Panel C: Oil theft</i>				
Local firm	-4.805*** (0.804)	-4.726*** (0.791)	-4.370*** (0.728)	-2.972** (1.170)
Observations	3183	3183	3183	3183
R ²	0.712	0.712	0.712	0.711
Field FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. Treatment definition is given in table header. Main is the primary treatment indicator used throughout the paper, which defines the localization event year as the first year of treatment in either dataset. DI first takes both DI and NNPC treatments, but uses the DI event year if a field is treated in both datasets. DI only uses only localizations that occur in the DI data. NNPC only uses only localizations that occur in the NNPC data. Output is measured in millions of barrels of oil per year. Oil spills are the total number of non-sabotage spills within 15 km of the field. Oil theft is the total number of sabotage spills within 15 km of the field. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

ditional (million) barrels of oil annually on recorded oil spills. In column (1) of Table A17, I estimate this quantity using a simple fixed effects regression of oil spills on output, controlling for field and year fixed effects.

Now recall the treatment effect of localization, ψ and index it by y for the output outcome and s for the oil spills outcome. Columns (2) and (3) display these effects, taken from Table 2. Finally, column (4) subtracts from ψ_s the implied increase in spills that would result only from the output gains, $\psi_y\gamma$. The three-equation system is estimated jointly to enable nonlinear hypothesis testing. Column (4) shows that once we account for this effect, localization still

Table A17: The effect of divestment on oil spills: adjustment for increased oil production

	γ	ψ_y	ψ_s	$\psi_s - \psi_y\gamma$
Estimate	0.131***	0.848***	1.380*	1.269*
	(0.042)	(0.297)	(0.765)	(0.759)

Standard errors in parentheses are clustered at the field level. Sample is the same as in Table 2. Parameters are in the table header, with y and s indexing the output and oil spills outcomes, respectively. Estimation of the three-equation system is conducted jointly with seemingly unrelated estimation for nonlinear hypothesis testing across equations. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

increases oil spills – the impact falls only slightly and remains significant at the 10% level.

C.11 Infrastructure type

In the main results of Table 2, 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; unfortunately, we lack detailed data on the size of thefts. 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. 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 A18, re-estimating the main TwFE 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.

C.12 Theft-output elasticity

The primary interpretation of the results in this paper is that improvements in the security situation and consequent reductions in oil theft lead to output advantage for local firms. This is supported by both the time-path of effects in the main event-study and the analysis of

Table A18: 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 firm	-2.864*** (0.570)	-2.222*** (0.576)	0.324 (0.282)	0.260 (0.371)	-1.647*** (0.457)	-2.690*** (0.687)	-0.624*** (0.107)	-0.853*** (0.151)
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	3183	3183	3183	3183	3183	3183	3183	3183
R^2	0.476	0.535	0.438	0.487	0.764	0.808	0.415	0.534

Standard errors in parentheses 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 infrastructure type. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

heterogeneous effects by asset type. Still, it is possible that different features of local firms' operations cause both increased output and reduced theft simultaneously. One important falsification test for the causal mechanism is to determine whether the increase in output quantitatively consistent with the increase oil theft. That is, given the elasticity of oil output to theft, how much of the increase in output can be "explained" by the treatment effect on oil theft.

First, define the following three parameters of interest: $\eta_{y,b}$ is the impact of black-market theft b on output y , ψ_y is the ATE of local ownership on output, and ψ_b is the ATE of ownership on oil theft. Then the "residual" increase in output that cannot be quantitatively explained by the reduction in theft is $\psi_y - \eta_{y,b}\psi_b$, and the total share of ψ_y explained by oil theft is $\frac{\eta_{y,b}\psi_b}{\psi_y}$. Of course, we have estimates for ψ_y and ψ_b from Table 2. However, we do not have reliable estimates for $\eta_{y,b}$, the causal effect of theft on oil output.

I estimate $\eta_{y,b}$ in Table A19 using an instrumental variables approach. Oil theft and output are endogenous equilibrium outcomes, likely exhibiting both reverse causality and omitted variables bias. Identification of $\eta_{y,b}$ requires an exogenous shock that satisfies the exclusion restriction – it must alter incentives in the oil black market but not directly affect oil production decisions except by changing theft risk. One such shock is the national energy market. Despite producing 2-2.5 million barrels of crude oil per day, Nigeria meets the vast majority of domestic fuel demand through imports. As Rexer and Hvinden (2022) show, the period between 2006-2016 was one of steadily worsening domestic fuel shortages, a result of shrinking domestic refining capacity, mismanagement and corruption in import market, and increasingly unsustainable fuel subsidies. This fuel crisis has coincided with the aggregate growth in oil theft *and* a shift toward supplying the domestic market SDN (2019a). Still, since Nigeria exports nearly 90% of its oil output, these domestic market conditions should not materially affect incentives for production except by increasing black market oil theft.

Table A19: Impact of theft on output: Instrumental variables estimation

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: First stage</i>						
Log gasoline imports \times Distance to coast (km)	-0.217*** (0.075)	-0.324*** (0.111)	-0.235** (0.103)	-0.348*** (0.081)	-0.521*** (0.117)	-0.603*** (0.162)
Observations	2476	2476	2476	2476	2476	2476
R ²	0.723	0.728	0.761	0.749	0.761	0.773
<i>Panel B: Reduced form</i>						
Log gasoline imports \times Distance to coast (km)	0.048** (0.024)	0.057** (0.027)	0.069*** (0.025)	0.035** (0.014)	0.037** (0.016)	0.037** (0.014)
Observations	2476	2476	2476	2476	2476	2476
R ²	0.861	0.862	0.873	0.869	0.869	0.879
<i>Panel C: 2SLS</i>						
Oil theft events, 15 km	-0.223 (0.154)	-0.175 (0.115)	-0.296* (0.178)	-0.101** (0.046)	-0.071** (0.033)	-0.061** (0.028)
F-statistic	8.242	8.530	5.242	18.569	19.998	13.938
Observations	2476	2476	2476	2476	2476	2476
Field FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Amnesty controls	No	Yes	Yes	No	Yes	Yes
Main controls \times Year FE	No	No	Yes	No	No	Yes
Militant controls \times Year FE	No	No	No	Yes	Yes	Yes

Standard errors in parentheses are clustered at the field level. Sample is the panel of 2476 field-years from 2006-2016 for which oil output information is available. Outcome in Panel A is oil theft, measured as the total number of sabotage spills within 15 km of the field. Outcome in Panels B and C is oil output, measured in millions of barrels of oil per year. Main controls are latitude of the field centroid, distance to Niger River, and distance to the capital. Amnesty controls is the interaction between a post-2009 indicator and distance to the coast. Militant controls includes distance to the nearest militant camp. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

I measure fuel shortages by the log of aggregate refined gasoline imports. However, this quantity varies only at the national level over time, while the instrument must have field-specific variation. To generate cross-sectional variation, I interact the national trend in log gasoline imports with distance to the coast. Rexer and Hvinden (2022) show that distance to the coast is a reasonable proxy for black market costs, since coastal locations are proximate to export destinations and coastal waterways used for illegal transport. As imports rise and shortages of legal fuel are alleviated, black market margins shrink and oil theft becomes unprofitable first in higher-cost inland locations, while remaining profitable in low-cost coastal regions. As such, we should expect a negative first stage coefficient on the interaction between national import trends and distance to the coast (i.e., alleviation of fuel crises reduces oil theft more in further inland locations). This is exactly the result of Table A19, Panel A, which shows first stage estimates with different combinations of controls, all of which include two-way fixed effects. All estimates are significant at 1% suggesting a relevant instrument. The instrument is strongest in columns (4) and (5), achieving an F -statistic of 19-20. These

specifications control for militant presence (4) and amnesty policy (5), both of which may be correlated with costs and trends in oil importation.

The reduced form shows consistent results in Panel B. Alleviating gasoline shortages is associated with differentially large increases in output in the low-coast inland locations where theft falls. The reduced form effects are significant at 5 or 1% in all specifications. Finally, Panel C estimates a 2SLS model using the interaction between log national gas imports and distance to coast as an instrument for theft, conditioning on year and field fixed effects. The results indicate a robustly negative impact of theft on oil output, ranging 0.06-0.22 million fewer barrels annually per additional theft incident. These estimates, however, are only significant at conventional levels for specifications in (4)-(6) with high first-stage F -statistics.

Table A20: Theft-output elasticity

	$\eta_{y,b}$	ψ_b	ψ_y	$\psi_y - \eta_{y,b}\psi_b$	$\frac{\eta_{y,b}\psi_b}{\psi_y}$
Estimate	-0.101**	-5.932***	0.945***	0.346	0.634*
	(0.045)	(1.110)	(0.323)	(0.414)	(0.355)

Standard errors in parentheses are clustered at the field level. Sample is the panel of 2476 field-years from 2006-2016 for which oil output information is available. $\eta_{y,c}$ is the effect of oil theft on oil output, from the 2SLS specification in column (2), Panel C in Table A19. $\eta_{c,l}$ is the effect of local ownership on oil theft. $\eta_{y,l}$ is the effect of local ownership on oil output. All equations include fixed effects for field and year and control only for militant presence interacted with year dummies. Estimation of the three-equation system is conducted jointly with GMM for nonlinear hypothesis testing. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

I choose column (4) of Table A19 as the preferred estimate of $\eta_{y,b}$ primarily because it sits in the middle of the range of magnitudes and obtains a strong first stage. However, one can easily repeat the exercise of Table A20 with different estimates in order to bound the share of ψ_y explained by theft. For Table A20, I jointly all of the required parameters using a GMM system with three moment conditions and the instrument for theft, on the sample of field-years with nonmissing oil output. These estimates are in columns (1)-(3).⁶⁶ I then conduct nonlinear hypothesis tests to generate standard errors for columns (4) and (5). The main output effect is 0.945, significant at 1%, while column (4) shows the residual effect after removing the role of theft is only 0.346, and not significantly different from zero. This implies that 63.4% of the output effect can be explained by theft (column 5), an estimate significant at the 10% level.

⁶⁶Each equation controls for all of the variables indicated in Table A20 column (4).

D Robustness tests: two-way fixed effects

The main estimates of divestment effects use TWFE estimation methods. A growing literature has identified the fundamental challenges in interpreting TWFE estimates as average treatment effects (ATEs) in settings with staggered treatment timing and treatment effect heterogeneity across cohorts or over time. The central issue is that the TWFE estimate is a weighted average of 2x2 difference-in-difference comparisons across different treatment and control groups (Goodman-Bacon 2021). Unfortunately, some of these 2x2 comparisons are quite bad in the presence of dynamic effects; for example, always-treated observations act as controls even though their previous treatment status should alter their trends. Similar logic applies to comparing units treated later to those earlier. These bad control have been alternatively expressed as negative weights on some unit-time-specific heterogeneous ATEs (de Chaisemartin and D’Haultfoeuille 2020). Baker et al. (2021) show that these issues have real empirical implications, and can lead TWFE estimates to be quantitatively misleading and even wrong-signed relative to underlying ATEs. Fortunately, several diagnostic decomposition tools and alternative estimators can help address this problem (Goodman-Bacon (2021), Callaway and Sant’Anna (2021), de Chaisemartin and D’Haultfoeuille (2020)).

D.1 TWFE decomposition

As an initial diagnostic tool, Goodman-Bacon (2021) shows that the staggered-adoption TWFE estimate can be decomposed into a weighted average of all 2×2 difference-in-difference comparisons. These weights depend on the size of the groups and the variance of the treatment in each 2×2 comparison; TWFE will tend to place lower weight on 2×2 estimates for units treated early or late in the panel. The key insight is that these weights identify which comparisons are driving the overall TWFE results. Table A21 presents weights and average treatment effect estimates for each 2×2 DD comparison type.⁶⁷ Panel A presents the results for output, while Panel B decomposes the oil theft effect.

Because of the large sample of never-treated clusters, the TWFE estimate heavily weights the “treated vs. never treated” 2×2 comparison, which accounts for 85% of the treatment effect. In Panel B, every 2×2 comparison estimate is negative, ranging from -1.1 to -11.2. In Panel A, every 2×2 group estimate is positive and of comparable magnitude, except the “Later treated (T) vs. Earlier treated (C)” comparison, which uses earlier treated units as controls for fields that switch into treatment towards the end of the panel. Note, however, that this is a so-called “forbidden” comparison, since it uses already-treated units as controls, despite the fact that their treatment status has already affected their subsequent trend. To adjust for this, I report the “purged estimate”, which removes the “Later treated (T) vs. Earlier treated (C)” and “Treated (T) vs. Already treated (C)” comparisons, both of which rely

⁶⁷The estimation is run on a subsample of 275 fields for which a balanced panel is available.

Table A21: Goodman-Bacon (2021) TWFE weights

Comparison	Weight	Estimate
<i>Panel A: Output</i>		
Earlier treated (T) vs. Later treated (C)	0.051	0.811
Later treated (T) vs. Earlier treated (C)	0.022	-0.858
Treated (T) vs. Never treated (C)	0.848	0.856
Treated (T) vs. Already treated (C)	0.079	0.567
TWFE estimate		0.792
Purged estimate		0.853
<i>Panel B: Oil theft</i>		
Earlier treated (T) vs. Later treated (C)	0.051	-11.202
Later treated (T) vs. Earlier treated (C)	0.022	-8.658
Treated (T) vs. Never treated (C)	0.848	-4.684
Treated (T) vs. Already treated (C)	0.079	-1.147
TWFE estimate		-4.824
Purged estimate		-5.054

Sample is the subset of 275 fields for which a balanced panel is available ($N = 3025$). Outcome variable in panel header. Output is measured in millions of barrels of oil per year, using the GLM method of Table A11 (4) to impute missing values. Oil spills are the total number of non-sabotage spills within 15 km of the field. Oil theft is the total number of sabotage spills within 15 km of the field. All models estimate two-way fixed effects weights and ATEs for different 2x2 comparison groups using the method explained in Goodman-Bacon (2021). T and C in parentheses indicates which observations are used as treatment and which as control, respectively, for a given comparison. Purged estimate refers to the weighted ATE which removes Treated (T) vs. Already treated (C) and Later (T) vs. Earlier treated (C) comparisons. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

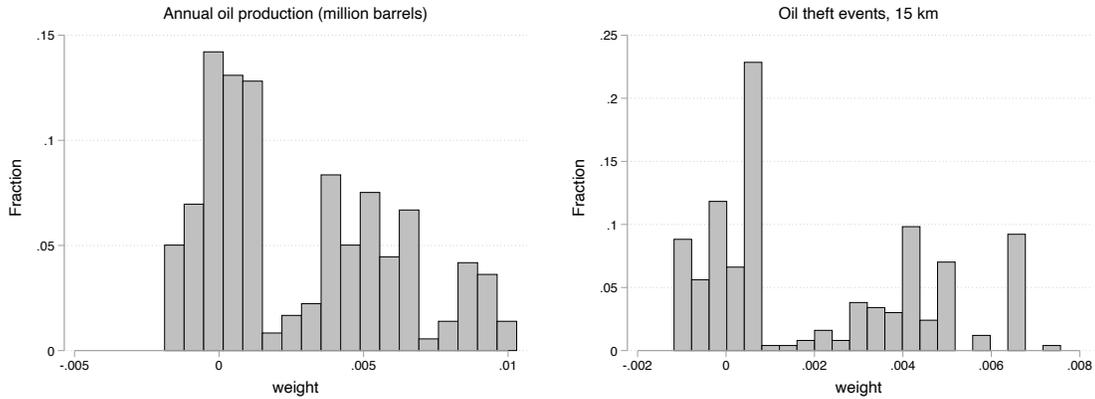
on already-treated fields to serve as controls, and reweights the estimate accordingly. Both purged estimates are even larger than the main estimate, suggesting that TWFE biases our estimates toward zero.

Table A21 also decomposes treatment effects by identifying variation, allowing us to probe the identification assumption. In Section 4, I argue that treatment timing may be more plausibly exogenous than treatment assignment. The comparison that leverages only timing variation among ever-treated units is “Earlier treated (T) vs. Later treated (C),” while the comparison that relies on variation between treatment and never-treated fields is “Treated (T) vs. Never treated (C).” Despite the fact that the latter drives much of our observed effect, the signs and magnitudes of the former are similar. Relying only on exogenous treatment timing does not weaken the results.

D.2 Unit weights

In a similar exercise, de Chaisemartin and D’Haultfoeuille (2020) derive a general formula that expresses the problem of “forbidden comparisons” as possibly negative weights on unit-by-time-specific average treatment effects (ATEs). If there are sufficiently many ATEs with large enough negative weights, it is plausible that the TWFE estimate will have a different sign as the individual ATEs. Figure A10 displays histograms of estimated weights for the two main outcomes. In all cases, the distributions are centered to the right of zero, and 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.

Figure A10: Histogram of de Chaisemartin and D’Haultfoeuille (2020) TWFE weights



Note: Figure shows implied TWFE weights for unit-and-time-specific ATEs, as derived in de Chaisemartin and D’Haultfoeuille (2020). Sample is the panel of 314 oilfields from 2006-2016. Output (left panel) is measured in millions of barrels of oil per year. Oil theft (right panel) is the total number of sabotage spills within 15 km of the field.

D.3 Stacked DD

Moving beyond diagnostics, several recent papers propose estimators for addressing the issues in TWFE. In general, these methods amount to different ways of removing already-treated controls from the estimation. One alternative estimation method is the stacked DD,⁶⁸ as suggested by Goodman-Bacon (2021). 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 event-time takes the place of calendar year. This ensures that already-treated observations are never used as controls. We then estimate the following equation, for unit i in cohort-stack c

⁶⁸See Gormley and Matsa 2011, Deshpande and Y. Li 2019, and Baker et al. 2021 for examples.

for event-time t

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

Standard errors are clustered at the field level. The parameter β is a variance weighted average of cohort-specific treatment effects, where each cohort-specific comparison is only between newly treated and not-yet-treated groups. An additional robustness test is to further restrict the sample either to ever-treated or never-treated fields, in order to isolate the role of treatment timing, as in the Goodman-Bacon (2021) decomposition.

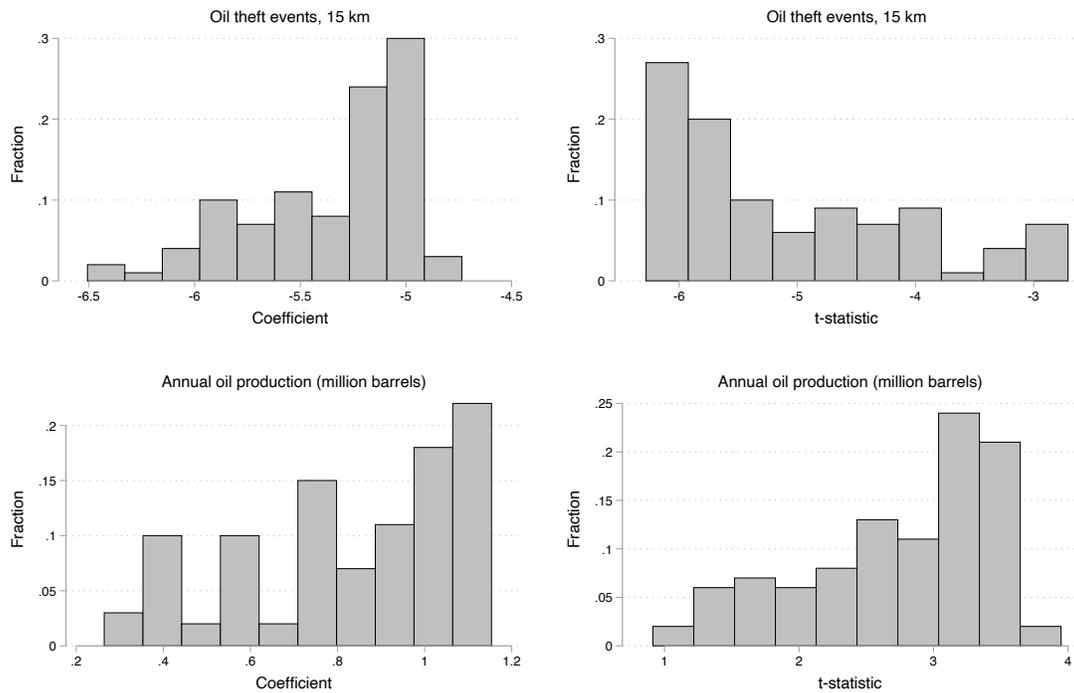
Table A22: The effect of divestment: stacked-DD estimation

	(1)	(2)	(3)
<i>Panel A: Output</i>			
Local firm	1.024*** (0.327)	1.020*** (0.336)	1.164** (0.536)
Observations	20637	18695	2506
R^2	0.872	0.873	0.750
<i>Panel B: Oil spills</i>			
Local firm	1.632** (0.778)	1.619** (0.816)	2.859** (1.182)
Observations	26229	23675	3338
R^2	0.584	0.584	0.622
<i>Panel C: Oil theft</i>			
Local firm	-5.001*** (0.827)	-4.718*** (0.784)	-10.885*** (2.514)
Observations	26229	23675	3338
R^2	0.715	0.716	0.729
Field-by-Cohort FE	Yes	Yes	Yes
Year-by-Cohort FE	Yes	Yes	Yes
Control group	All	Untreated	Treated

Standard errors in parentheses are clustered at the field level. Outcome variable in panel header. Output is measured in millions of barrels of oil per year. Oil spills are the total number of non-sabotage spills within 15 km of the field. Oil theft is the total number of sabotage spills within 15 km of the field. All models use the stacked difference-in-differences estimation method explained in Baker et al. (2021). All models use a symmetric event window of +/- 10 years. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The results of this analysis are given in Table A22 for the three main outcomes. Column (1) uses all possible control units, while column (2) uses only never-treated and column (3) uses only ever-treated. I find that full-sample stacked-DD estimates (columns 1-3) are robustly negative and significant for oil theft, and positive and significant for output and oil spills. The magnitude of effects is in fact somewhat larger than the TWFE estimates in Table 2. 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 A21. If anything, TWFE biases our main results toward zero.

Figure A11: Stacked-DD histogram over event-windows



Note: Figure shows histograms of coefficients (left panel) and t -statistics (right panel) from the stacked-DD specification from Table A22 for oil theft (top panel) and oil production (bottom panel) outcomes. Output is measured in millions of barrels of oil per year. Oil theft is the total number of sabotage spills within 15 km of the field. Treatment effects are estimated for all possible combinations of event windows up to 10 years before and 10 years after the event. Standard errors are clustered at the field level.

This estimation restricts to a ± 10 year event window. I also test robustness to estimating the stacked DD regression over all possible event-windows (including asymmetric windows) for output and theft, the two main outcomes. The resultant β coefficients and t -statistics are plotted in Figure A11. As desired, they are clustered around large negative and positive values, respectively.

D.4 CSDID

Callaway and Sant’Anna (2021) propose a semi-parametric DD estimator to address the “negative weights” problem, which also corrects for the down-weighting of 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 across various dimensions (cohort, time, or both). 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.

Table A23: The effect of divestment: Callaway and Sant’Anna (2021) estimation

Outcome	Output		Oil spills		Oil theft	
	(1)	(2)	(3)	(4)	(5)	(6)
Local firm	0.792*	0.778*	1.689***	1.623***	-4.169***	-4.366***
	(0.461)	(0.459)	(0.468)	(0.466)	(1.119)	(1.195)
Observations	1874	1875	2459	2459	2459	2459

Standard errors in parentheses are clustered at the field level. Output is measured in millions of barrels of oil per year. Oil spills are the total number of non-sabotage spills within 15 km of the field. Oil theft is the total number of sabotage spills within 15 km of the field. All models use the difference-in-differences estimation method for staggered adoption settings detailed in Callaway and Sant’Anna (2021). Columns (1), (3), and (5) use only never-treated observations as controls. Columns (2), (4), and (6) use both never and not-yet treated as controls. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A23 provides results using doubly-robust inverse-probability weighting (sz2020) for the three main outcomes. Columns (1), (3), and (5) use only never-treated observations as controls. Columns (2), (4), and (6) use both never and not-yet treated as controls. None of the specifications include control variables. All results are directionally robust and statistically significant at the 10% level or lower. The output effects are slightly smaller than in Table 2 and only significant at the 10% level, while the oil spill effects are larger and more significant. Note, however, that the estimation routine drops fields that are not “pair-balanced”, that is, observed in both $t = 0$ and $t = 1$ of event-time. This smaller sample may explain slightly different results and loss of significance.

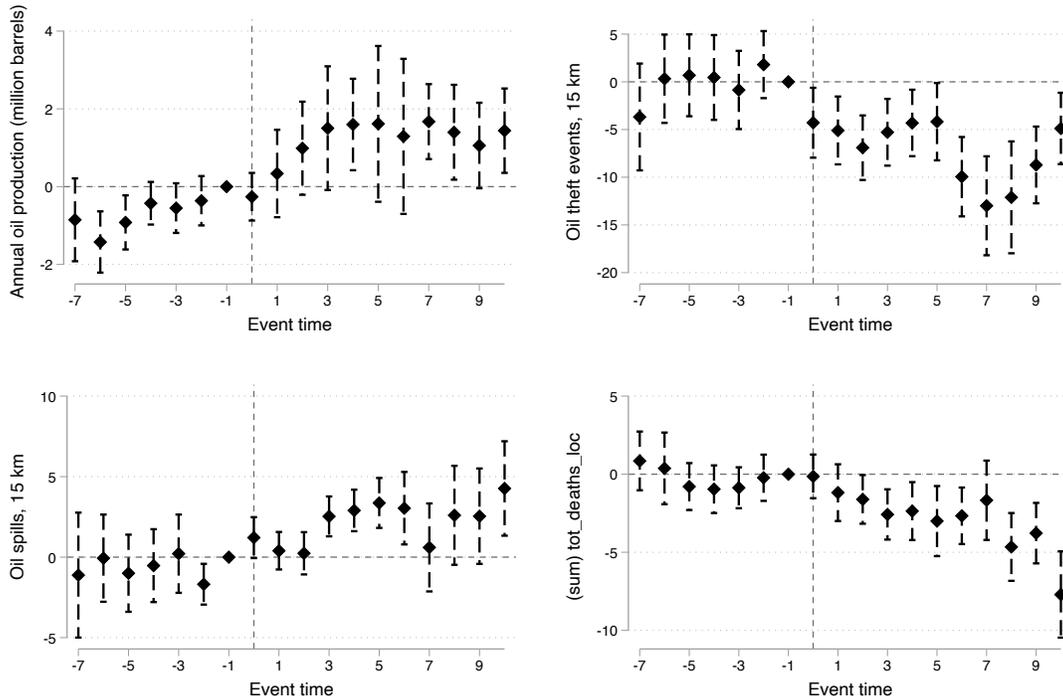
D.5 Event-study plots: main effects

The main event-study plot in Figure 3 employs the stacked DD configuration for the three main outcomes, using all yet-untreated fields as controls, and controlling for interacted co-variates. In this section, I consider event-studies using different estimators, samples, control

variables, and outcome variables, in order to verify that the parallel trends obtained in Figure 3 are robust.

Figure A12 replicates the stacked-DD event-study specification for four key outcomes *without* any control variables.

Figure A12: Stacked-DD event-study: main outcomes without controls

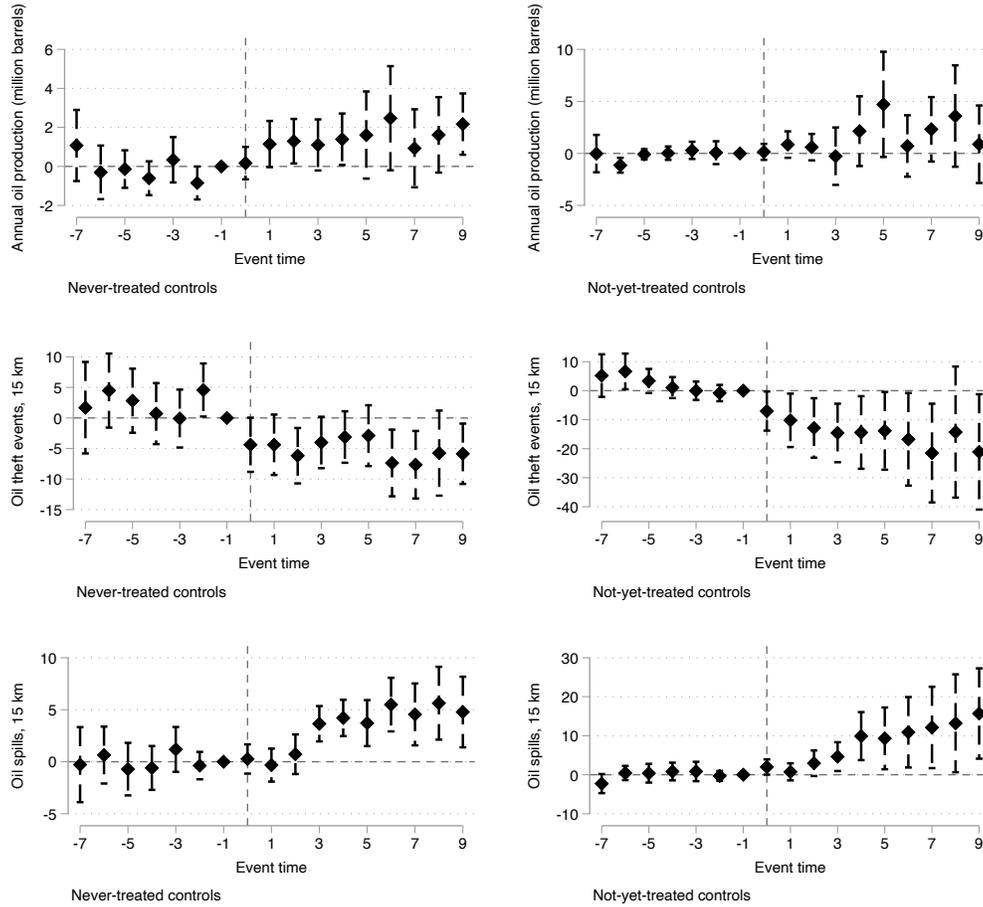


Note: Figure shows coefficients from stacked event-study regressions for oil production, oil theft, and oil spills. Standard errors are clustered at the field-level. Output is measured in millions of barrels of oil per year. Oil spills are the total number of non-sabotage spills within 15 km of the field. Theft is the total number of sabotage spills within 15 km of the field. Conflict is the total number of conflict deaths within 15 km of the field reported by local news media.

Figure A13 decomposes the stacked DD event-study by control group following the method in Table A22 for the three main outcomes of output, theft, and oil spills. The left panel uses only never-treated controls, while the right panel uses only ever-treated controls. There is evidence for parallel trends in both comparisons.

Figure A14 uses TWFE for the event-study specification, for the three main outcomes as well as conflict deaths. All regressions include the main set of spatial controls used throughout the paper. The results are visually very similar to the main event-study, with the exception of a noisy zero coefficient estimated for year 6 after treatment for the output outcome. Figure A15 then re-estimates the TWFE specification using only the NNPC data to define the treatment variable. Perhaps predictably, the results differ somewhat from the main event-

Figure A13: Stacked-DD event-study: main outcomes by control group

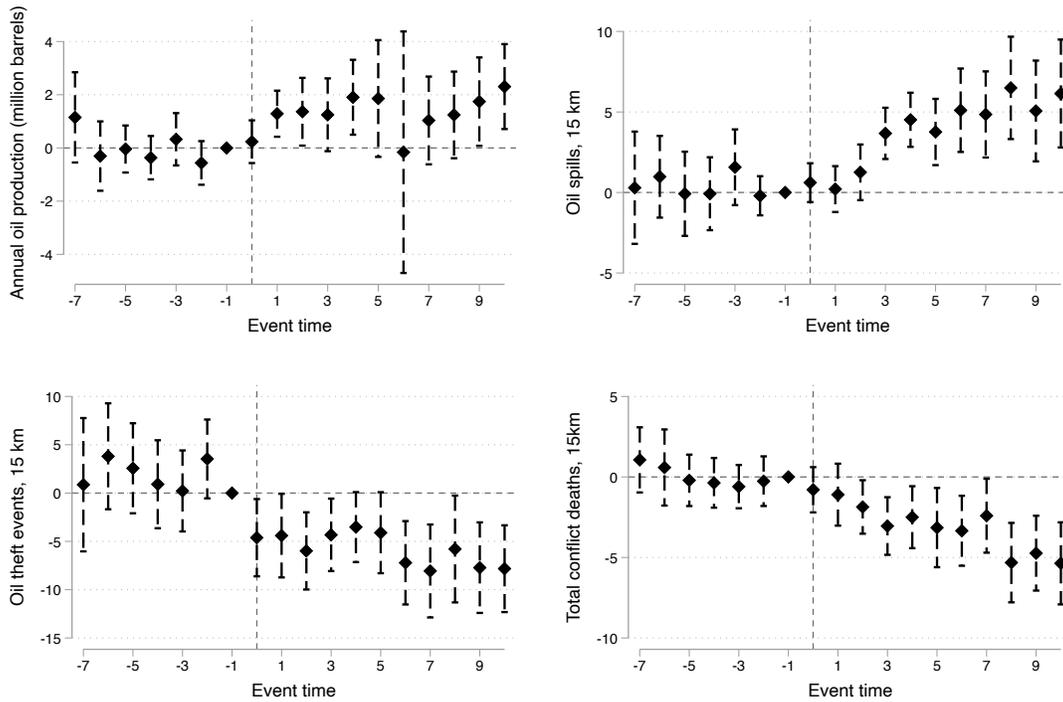


Note: Figure shows coefficients from stacked event-study regressions for oil production (top panel), oil spills (middle panel), and oil theft (bottom panel). Standard errors are clustered at the field-level. Left panel uses only never-treated controls, while right panel uses only ever-treated controls. Output is measured in millions of barrels of oil per year. Oil spills 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. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km.

study. However, the qualitative trends are similar and pre-trends remain flat. Finally, Figure A16 estimates the event-study using the DI data only to generate the event-year. The effects are generally smaller and shorter-lived, though again, pre-trends remain parallel. The TWFE event-study is robust across all different treatment indicators.

Figure A17 uses the de Chaisemartin and D’Haultfoeuille (2020) estimates aggregated by event-year to generate event-study plots. Given the smaller sample, the results appear somewhat noisier than stacked and TWFE specifications, but the patterns are remarkably consistent. The only major difference is that the post-event coefficients for oil theft converge

Figure A14: TWFE event-study: main outcomes

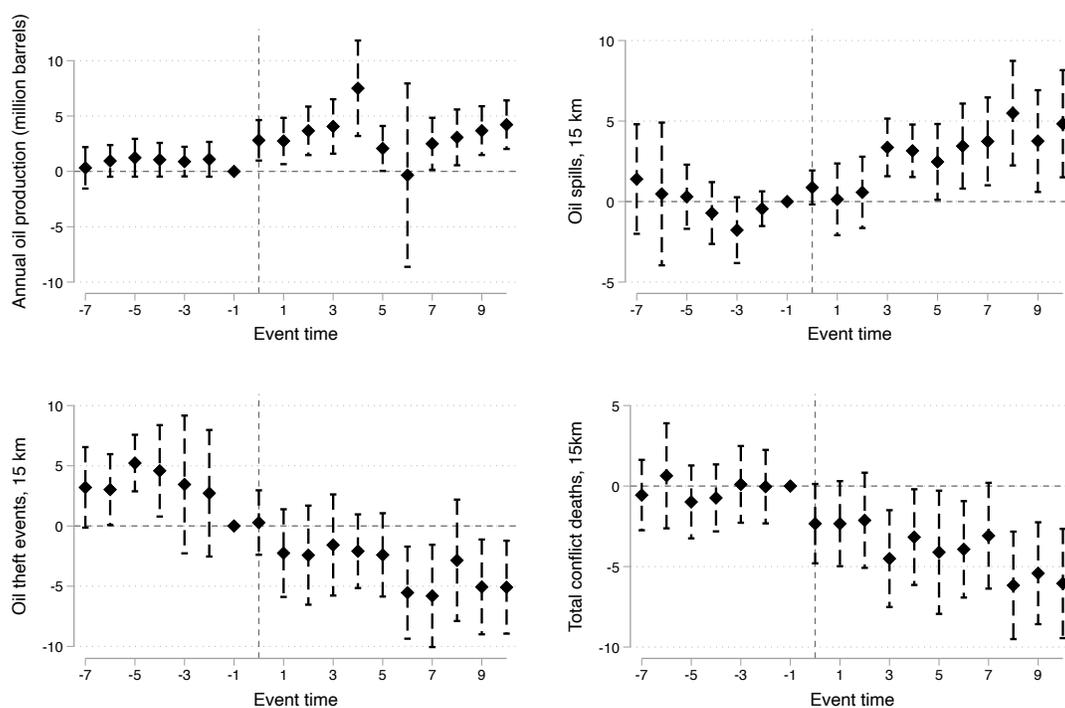


Note: Figure shows coefficients from TWFE event-study regressions for oil production, oil theft, and oil spills. Standard errors are clustered at the field-level. Output is measured in millions of barrels of oil per year. Oil spills are the total number of non-sabotage spills within 15 km of the field. Theft is the total number of sabotage spills within 15 km of the field. Conflict is the total number of conflict deaths within 15 km of the field reported by local news media.

back to zero in years past five, suggesting smaller long-run effects on theft than the other specifications. Figure A18 uses the switching estimator proposed by de Chaisemartin and D’Haultfoeuille (2020) to estimate event-study coefficients. Interestingly, the results of this estimator are more similar to the de Chaisemartin and D’Haultfoeuille (2020) specification than the stacked or TWFE plots.

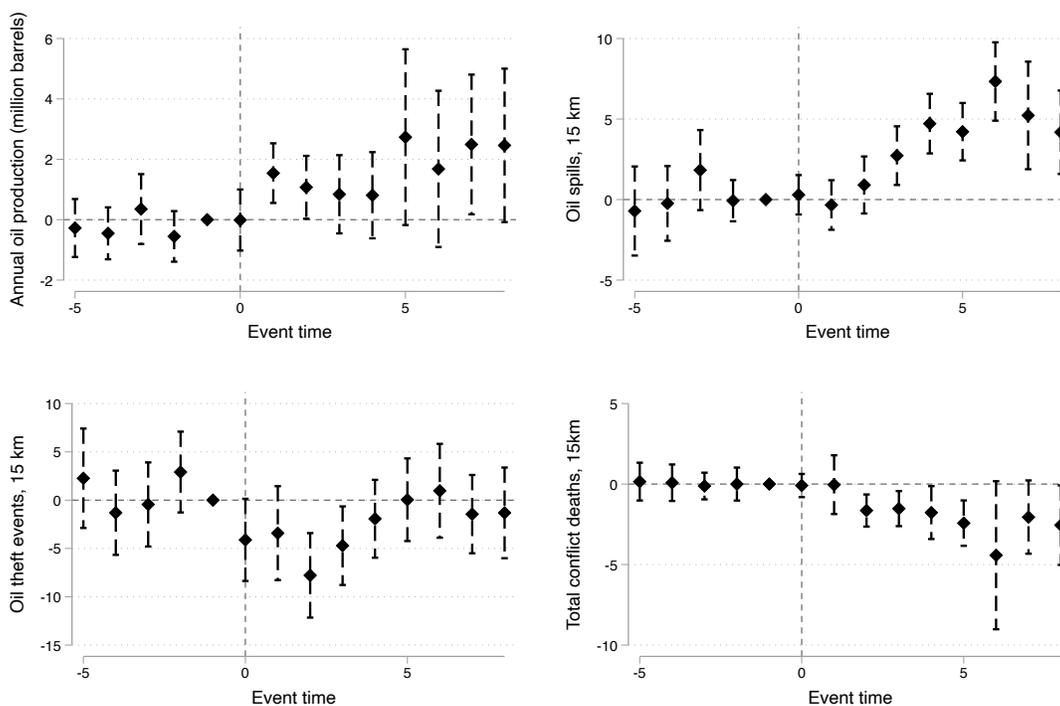
Lastly, Figure A19 estimates the TWFE event study corresponding to Table 5, columns (3) and (6), which look at terminated divestments. In this case, the date in which a terminated divestment was initiated is taken as the event-date, yielding 36 treated fields. As expected, the plots indicate no measurable pre-trends and no post-treatment effects.

Figure A15: TWFE event-study with NNPC only: main outcomes



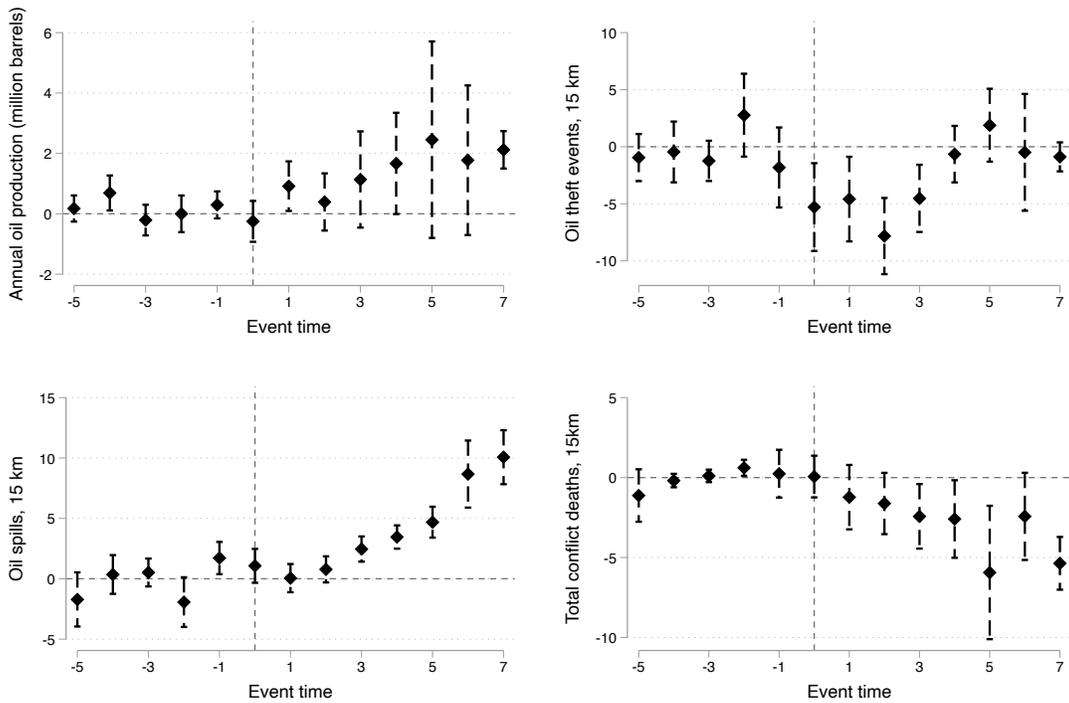
Note: Figure shows coefficients from TWFE event-study regressions for oil production, oil theft, and oil spills. Treatment timing defined using NNPC data only. Standard errors are clustered at the field-level. Output is measured in millions of barrels of oil per year. Oil spills are the total number of non-sabotage spills within 15 km of the field. Theft is the total number of sabotage spills within 15 km of the field. Conflict is the total number of conflict deaths within 15 km of the field reported by local news media.

Figure A16: TWFE event-study with DI only: main outcomes



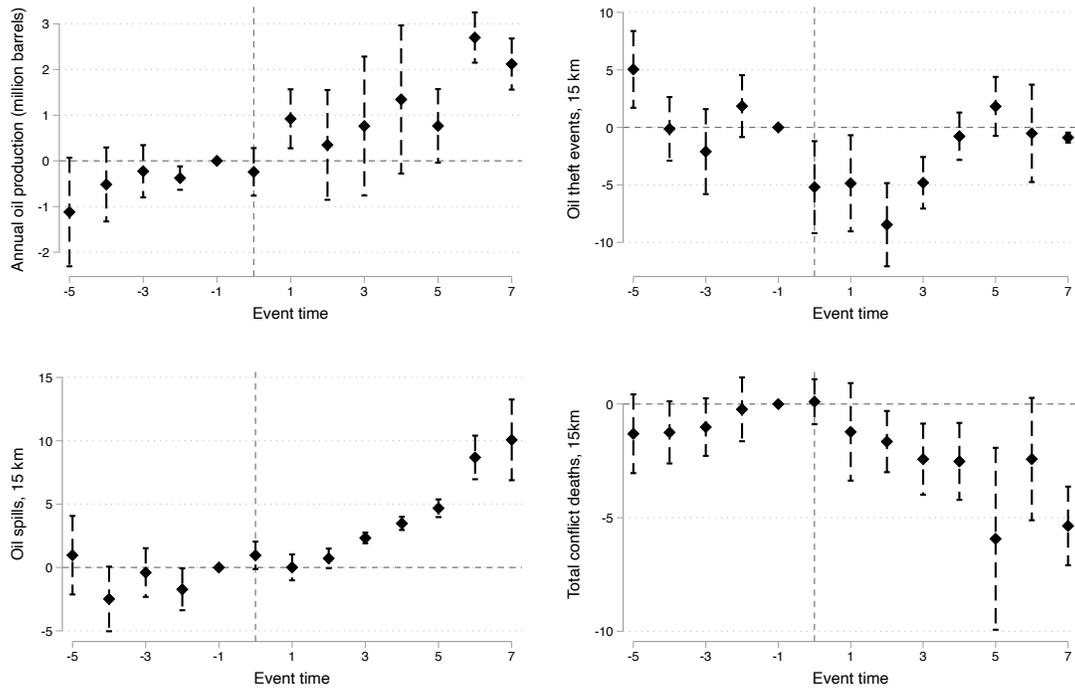
Note: Figure shows coefficients from TWFE event-study regressions for oil production, oil theft, and oil spills. Treatment timing defined using DrillingInfo data only. Standard errors are clustered at the field-level. Output is measured in millions of barrels of oil per year. Oil spills are the total 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. Conflict is the total number of conflict deaths within 15 km of the field reported by local news media.

Figure A17: Callaway and Sant'Anna (2021) event-study: main outcomes



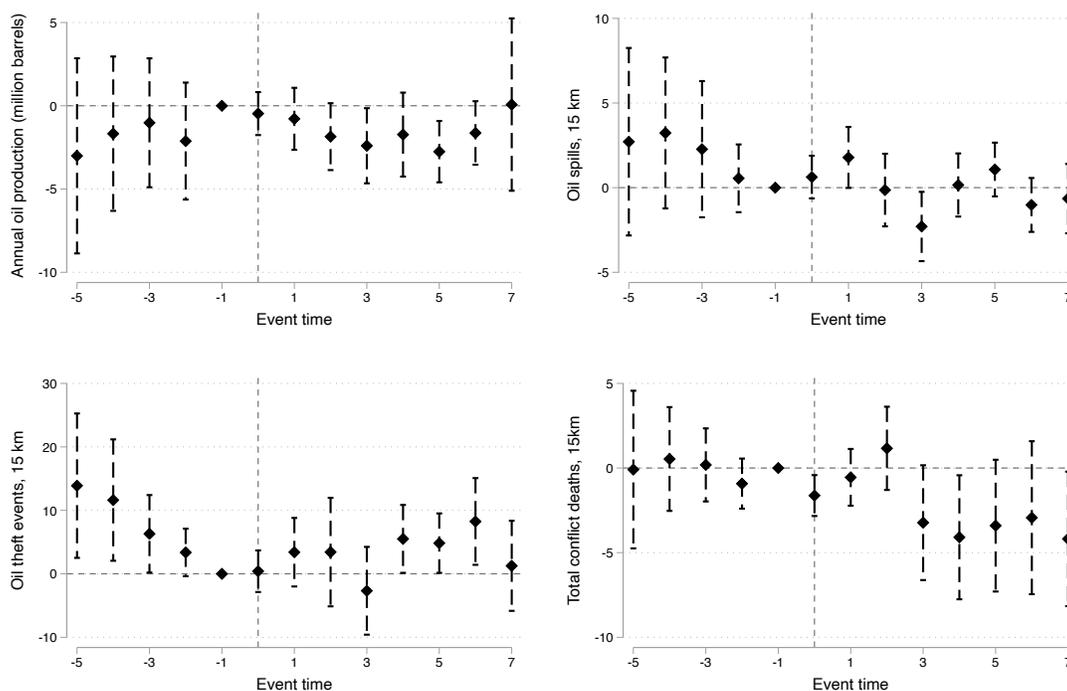
Note: Figure shows coefficients from event-study estimation following Callaway and Sant'Anna (2021) for oil production, oil theft, and oil spills. Standard errors are clustered at the field-level. Output is measured in millions of barrels of oil per year. Oil spills are the total number of non-sabotage spills within 15 km of the field. Theft is the total number of sabotage spills within 15 km of the field. Conflict is the total number of conflict deaths within 15 km of the field reported by local news media.

Figure A18: de Chaisemartin and D’Haultfoeuille (2020) event-study: main outcomes



Note: Figure shows coefficients from event-study estimation following de Chaisemartin and D’Haultfoeuille (2020) for oil production, oil theft, and oil spills. Standard errors are clustered at the field-level. Output is measured in millions of barrels of oil per year. Oil spills 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. Conflict is the total number of conflict deaths within 15 km of the field reported by local news media.

Figure A19: Placebo event-study: main outcomes



Note: Figure shows coefficients from TWFE event-study estimation for oil production, oil theft, and oil spills. Event time is defined relative to the first year of a terminated divestment. Standard errors are clustered at the field-level. Output is measured in millions of barrels of oil per year. Oil spills 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. Conflict is the total number of conflict deaths within 15 km of the field reported by local news media.

E Robustness tests: Other outcomes

E.1 Other outcomes

I consider the impact of divestment on other outcomes in Table A24. Columns (1)-(2) show that localization significantly reduces the overall level of conflict-related violence on oil assets, by roughly 3 events per year on average. In columns (3)-(4), the effect on maritime piracy – mostly tanker hijackings – is also large and negative but not significant at conventional levels.⁶⁹ Columns (5)-(6) test the extensive margin of oil production, estimating the impact of divestment on “shut-ins,” defined as zero-production field-years. There is no evidence that divestment affects extensive margin production. This null result, precisely estimated at zero, helps illustrate what is driving the main output result. Theft and violence may reduce output in two ways: *i*) by technical losses on the intensive margin from spilled and stolen oil, and *ii*) by affecting the incentive on the extensive margin to produce at all. In our context, it appears that local firms are not so much reviving shut-in fields as they are recovering lost production from producing fields.⁷⁰

Table A24: The effect of divestment, other outcomes

Outcome	Conflict deaths		Piracy		Shut-in		Gas flaring	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Local firm	-3.043*** (1.052)	-3.203*** (0.993)	-0.111* (0.063)	-0.097 (0.065)	-0.009 (0.056)	0.003 (0.057)	0.415** (0.187)	0.271 (0.195)
Control group mean	2.006		0.154		0.234		1.068	
Observations	3183	3183	3183	3183	2476	2476	1503	1503
R ²	0.232	0.317	0.227	0.311	0.657	0.671	0.895	0.899
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

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. Conflict deaths are the total number of conflict-related fatalities reported in news media within 15 km of the field. Piracy is the pirate attacks within 15 km of the field. Shut-ins is an indicator for nonzero production in a field-year. Gas flaring is 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Columns (7)-(8) estimate the impact on gas flaring. The sample only covers 2012-2016, and so is substantially smaller. Local ownership increases gas flaring by 0.27-0.41 million

⁶⁹In additional tests, I find that the reduction in piracy is significant in the subsample of offshore assets where piracy is concentrated. Results available on request

⁷⁰This may be driven by the fact that long-standing shut-in fields may not be attractive targets for divestment. In fact, only 8 fields are shut-in at the time of their divestment. Furthermore, fields may shut-in during part of a year and still record positive production. Therefore, using aggregated annual data we cannot definitely rule out within-year extensive margin effects.

mcsf on average, 35-39% of the control group mean. This results in an additional 14.3-36.2 thousand tonnes of CO₂ emissions per field annually. However, these effects are only significant without control variables. 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 A25: The effect of operatorship, other outcomes

Outcome	Conflict deaths		Piracy		Shut-in		Gas flaring	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Local operator	-1.623** (0.814)	-1.786** (0.893)	-0.015 (0.088)	-0.008 (0.087)	-0.180** (0.072)	-0.178** (0.070)	0.386 (0.296)	0.307 (0.283)
Control group mean	2.006		0.154		0.234		1.068	
Observations	3183	3183	3183	3183	2476	2476	1503	1503
R ²	0.228	0.313	0.226	0.310	0.660	0.673	0.895	0.899
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

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. Conflict deaths are the total number of conflict-related fatalities reported in news media within 15 km of the field. Piracy is the pirate attacks within 15 km of the field. Shut-ins is an indicator for nonzero production in a field-year. Gas flaring is measured in million mcsf. 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 A26 re-estimates the models of Table 3 using the stacked DD estimator and excluding the shut-in outcome for which effects were zero. The results are generally consistent across specifications (1) and (2), which use never-treated controls. However, the results for column (3) – which uses only eventually-treated fields as controls – are generally not significant.

In Figure A20, 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 sometimes significant. The dynamic path of the coefficients provides some evidence that the treatment effect increases over time.

E.2 Robustness: conflict measurement

The results of Table 3 columns (1)-(2) may be biased by measurement error. Since ACLED's measurement of conflict activity is derived from news media reports, it is possible that multi-

Table A26: The effect of divestment on other outcomes: stacked-DD estimation

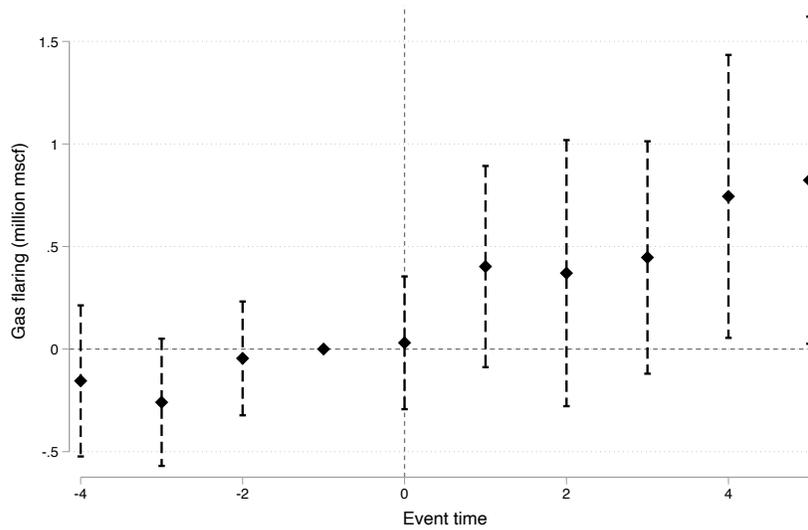
	(1)	(2)	(3)
<i>Panel A: Conflict deaths</i>			
Local firm	-3.332*** (1.055)	-3.397*** (1.058)	-1.620 (1.338)
Observations	26229	23675	3338
R ²	0.268	0.256	0.336
<i>Panel B: Piracy</i>			
Local firm	-0.096 (0.063)	-0.118 (0.072)	0.276 (0.199)
Observations	26229	23675	3338
R ²	0.242	0.216	0.337
<i>Panel C: Gas flaring</i>			
Local firm	0.428** (0.192)	0.430** (0.192)	0.242 (0.290)
Observations	11634	11006	1013
R ²	0.895	0.893	0.923
Field-by-Cohort FE	Yes	Yes	Yes
Year-by-Cohort FE	Yes	Yes	Yes
Control group	All	Untreated	Treated

Standard errors in parentheses are clustered at the field level. Outcome variable in panel header. Conflict deaths are the total number of conflict-related fatalities reported in news media within 15 km of the field. Piracy is the pirate attacks within 15 km of the field. Shut-ins is an indicator for nonzero production in a field-year. Gas flaring is measured in million mscf. All models use the stacked difference-in-differences estimation method explained in Baker et al. (2021). All models use a symmetric event window of +/- 10 years. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

national firms are considered more newsworthy. As such, a local takeover may reduce the media attention to a given field, rather than the underlying level of violent conflict.

Table A27 subsets conflict events by news media source. Column (1) re-prints the main results, while column (2) includes only conflict events reported by local Nigerian news media. Local reports are less susceptible to media bias because indigenous Nigerian firms are likely to be newsworthy to a local audience. Column (3) shows the results for only internationally-reported events. Since all events in ACLED are assigned a single source, the coefficients in (2) and (3) add up to (1). While the international effect is slightly larger, they all remain

Figure A20: Stacked-DD event-study: gas flaring



Note: Figure shows coefficients from stacked event-study regressions for gas flaring, measured in millions of mscf. Standard errors are clustered at the field-level. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km.

Table A27: The effect of divestment on conflict: robustness to measurement error

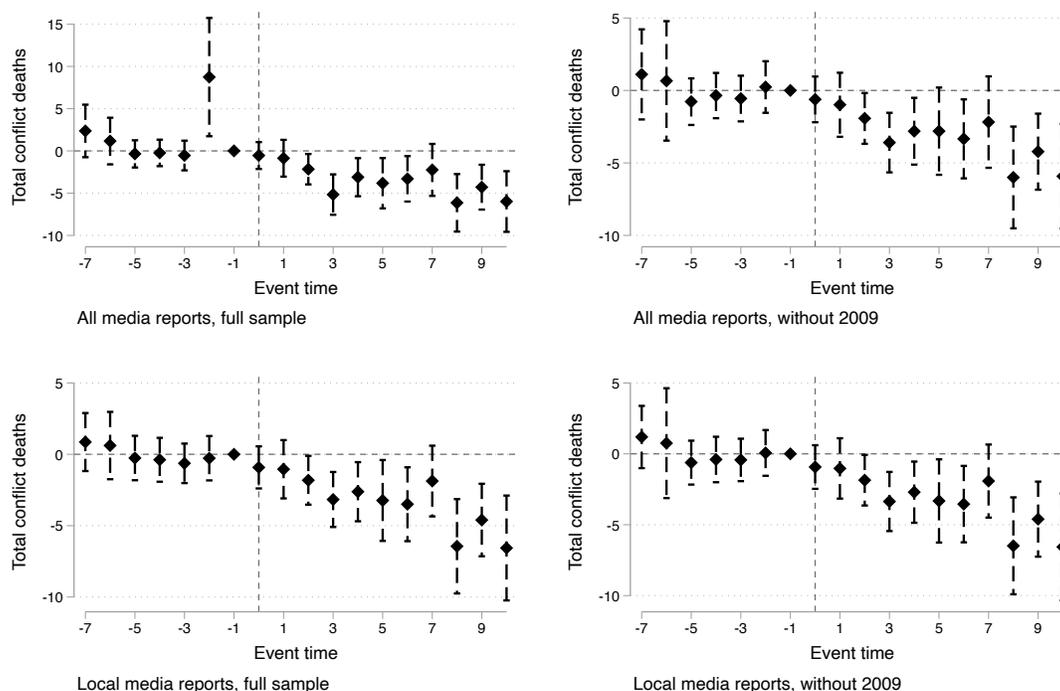
Outcome	All	Local	Int'l	Militant	Oil mil
	(1)	(2)	(3)	(4)	(5)
Local firm	-3.203*** (0.993)	-1.231*** (0.461)	-1.972** (0.902)	-3.138*** (0.953)	-1.522* (0.897)
Field FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Controls \times Year FE	Yes	Yes	Yes	Yes	Yes
Control group mean	2.006	1.099	0.907	1.720	0.621
Observations	3183	3183	3183	3183	3183
R^2	0.317	0.348	0.261	0.292	0.221

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. Conflict deaths in (1) are the total number of conflict-related fatalities reported in news media within 15 km of the field. Columns (2) and (3) subset conflict events to only those reported by international and local news media, respectively. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

statistically significant at the 1% level. It is therefore unlikely that the conflict results are driven primarily by measurement error. Column (4) subsets conflict events to only those that include organized militant groups, summing militant conflict deaths within 15 kilometers of

the field, while column (5) further restricts to militant events targeting the oil sector. The reduction in conflict is robust to restricting to explicitly oil-related organized violence.

Figure A21: Stacked-DD event-study: conflict deaths by measurement approach



Note: Figure shows coefficients from stacked event-study regressions for total conflict deaths within 15 km of the field. Top panel uses all news media reports in ACLED data. Bottom panel drops conflict events in ACLED reported by international news media sources. Left panel uses the full sample, while right panel drops 2009, the year of the Niger Delta amnesty. Standard errors are clustered at the field-level. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km.

Figure A21 assesses the role conflict measurement error in the stacked event-study framework. The top-left panel includes all conflict events and all years of data. While the trends are broadly parallel, there is a spike in conflict at event-time $t = -2$. This outlier coefficient is likely to be driven by fields that were divested in 2011, such that $t = -2$ corresponds to 2009, the year of the culmination of the Niger Delta conflict, which witnessed an unprecedented spike in violent conflict (see Rexer and Hvinden (2022)). To account for this, I drop the year 2009 from the estimation in the top-right panel. Pre-trends become flat and insignificant.

The Niger Delta conflict primarily targeted multinational firms, and was highly publicized in international media. It is therefore possible that the 2009 spike in conflict is an artefact of the data, driven by over-reporting of the conflict among international news sources. The bottom panel of the figure uses only local news media reports, as in Table A27 column (2). The pre-

divestment spike in conflict disappears, regardless of whether 2009 is included (left panel) or dropped (right panel). The results suggest that conflict measurement using international news media is indeed sensitive to outliers, and the restriction to local news media sources may be more appropriate. I therefore maintain this measurement assumption and use only local media reports in all of the event-study figures of Appendix D.5.

E.3 Robustness: enforcement outcomes

In this section, I consider robustness of the enforcement results to specification and outcome definition. Table A28 re-estimates the main results of Table 6 using the stacked DD specification. As in Appendix D.3, I find that once problematic controls in the TWFE are removed, the results are stronger and more significant. Estimates for all types of enforcement remain positive and significant at the 5% level or lower. Effects are also robust to using all controls (columns 1, 3 5, and 7), or using only never-treated controls (columns 2, 4, 6, and 8).

Table A28: The effect of local ownership on law enforcement activity: Stacked DD

Enforcement outcome	All oil theft		Oil seizures		Illegal refineries		Illegal export	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Local firm	1.598** (0.768)	1.717** (0.784)	1.405** (0.616)	1.480** (0.630)	1.272*** (0.464)	1.319*** (0.475)	0.399** (0.198)	0.421** (0.201)
Field-by-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-by-Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	26289	23730	26289	23730	26289	23730	26289	23730
R ²	0.477	0.482	0.424	0.430	0.419	0.420	0.327	0.327
Control group	All	NT	All	NT	All	NT	All	NT

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. Outcome variable is the total count of anti-oil theft law enforcement events within 15 km of the field. All models use the stacked difference-in-differences estimation method explained in Baker et al. (2021). All models use a symmetric event window of +/- 10 years. NT stands for never-treated. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

I argue that local law enforcement agents offer preferential protection to the assets of Nigerian firms. This protection is specific to the black market for stolen oil, the primary production risk faced by firms. However, a plausible alternative mechanism is that localization simply coincides with a generalized increase in law enforcement activity. I consider this hypothesis in Table A29, which estimates the impact of divestment on law enforcement actions against non-oil crime. Columns (1)-(2) aggregate all non-oil related criminal activities, while columns (3)-(6) disaggregate this category into two important crimes – kidnapping and gang activity. The results for all non-oil crime are quantitatively small and insignificant. Kidnapping produces somewhat larger positive point estimates, but remains noisy and insignificant. Columns (5)-(6) show, if anything, a reduction in law enforcement actions against gangs. Overall, there is no evidence of a generalized increase in law enforcement activity following divestment.

Table A29: The effect of divestment, non-oil law enforcement activity

Outcome	All non-oil		Kidnapping		Gangs	
	(1)	(2)	(3)	(4)	(5)	(6)
Local firm	0.091 (0.606)	0.007 (0.595)	0.359 (0.243)	0.381 (0.254)	-0.013 (0.035)	-0.066* (0.039)
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
Control group mean	1.739		0.568		0.077	
Observations	3183	3183	3183	3183	3183	3183
R ²	0.390	0.501	0.456	0.513	0.208	0.284

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. Enforcement category is given in the table header, defined as the total number of enforcement actions reported in news media 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Lastly, in Table A30 I consider the robustness of the enforcement results to different distance radii around the field centroid, from 5-30 km in intervals of 5km. The control group mean is given in each specification so that the results can be rescaled. The results are positive, large, and generally significant. It does not appear that spurious correlations or measurement error induced the definition of oil theft enforcement is driving our results.

Table A30: The effect of divestment: robustness to distance radius

Outcome	Oil theft enforcement					
	(1)	(2)	(3)	(4)	(5)	(6)
Local firm	0.295 (0.181)	0.709** (0.352)	1.163* (0.688)	2.144* (1.120)	3.281** (1.379)	4.743*** (1.689)
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
Distance (km)	5	10	15	20	25	30
Control group mean	0.209	0.752	1.739	3.049	4.529	6.330
Observations	3183	3183	3183	3183	3183	3183
R ²	0.378	0.476	0.578	0.627	0.673	0.708

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. Enforcement is the total number of non-militant anti-oil theft enforcement actions within a given distance radius of the field centroid. Distance radius is indicated in table footer. Controls are latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital, all measured in km. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

F Additional results

F.1 Political connections

Section 7.2 shows that firms' connections to the Nigerian security forces are associated with substantially lower levels of oil theft. In Table A31, I consider additional outcomes. Each column pair represents estimates from a TWFE regression of the outcome (indicated in table header) on an indicator for field-level connections to the security forces, with and without interacted controls.

Table A31: The effect of security connections on other outcomes

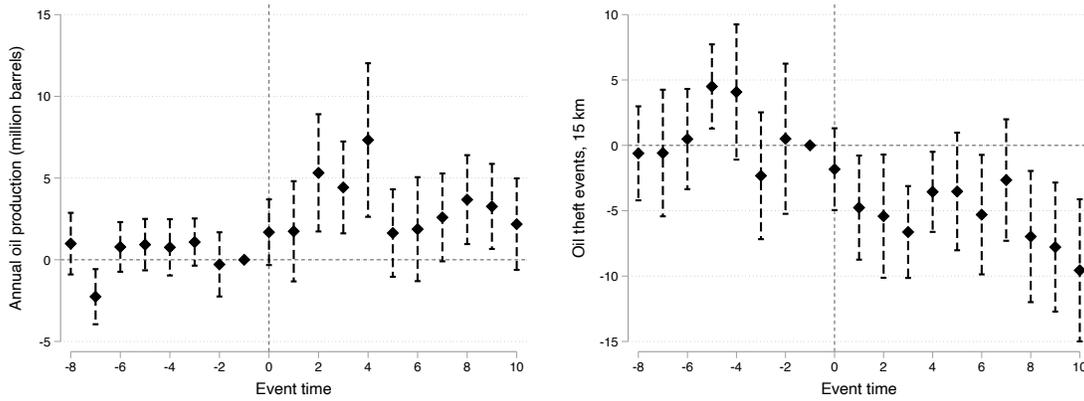
Outcome	Oil output		Oil spills		Oil theft		Enforcement	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Security forces	1.588*	1.495**	2.249***	2.081***	-3.523**	-4.068***	0.937	1.261**
	(0.888)	(0.742)	(0.733)	(0.692)	(1.412)	(1.492)	(0.691)	(0.593)
Field FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	2476	2476	3183	3183	3183	3183	3183	3183
R ²	0.861	0.878	0.590	0.650	0.711	0.754	0.476	0.577

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 indicated in table header, and measured as the total number of events 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Columns (1)-(2) show that security connections are associated with an increase in crude oil output of 1.5-1.6 million barrels per field annually, or 54% of the sample mean. These estimates are significant at 5% with control variables, and 10% without. Columns (3)-(4) investigate the impact on oil spills. Connected firms increase oil spills by 2.1-2.3 additional spills per field annually, significant at 1%. Column (6) reprints column (6) of Table 8, while column (5) shows the estimate without interacted controls. Finally, columns (7)-(8) use oil theft enforcement actions as the main outcome, and show a 0.9-1.3-event increase in state law enforcement, significant at 5% with controls.

Figure A22 estimates TWFE event studies, where the event-year is defined as the first year in which an oilfield obtains a boardmember or shareholder connected to the security forces. The results indicate that for both oil output (left panel) and oil theft (right panel), pre-trends are broadly parallel.

Figure A22: TWFE event-study: security connections



Note: Figure shows coefficients from TWFE event-study regressions for oil production (left panel) and oil theft (right panel). Standard errors are clustered at the field-level. Treatment timing is defined as the year an oilfield obtains its first boardmember or shareholder connected to the security forces. 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.

F.2 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, increasing bargaining space. The Nigerian oil market exhibits substantial variation in ownership agreements (see Figure A4), and local operators may have greater ownership stakes for several reasons: *i*) multinational divestment may lead to consolidation of stakes in joint ventures, and *ii*) 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.

It is therefore 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, in Table A32 I re-estimate the main TWFE regression at the concession-year level,⁷¹ where the outcome variable is either the concession equity Herfindahl-Hirschman Index (HHI), which measures overall consolidation, or the operator’s stake. Columns (1)-(2) estimate the model with the HHI outcome, while

⁷¹This is the level for which detailed ownership data are available for the period of 2013-2018.

Table A32: The effect of local ownership on equity consolidation

Outcome	HHI		Operator stake	
	(1)	(2)	(3)	(4)
Local operator	0.128*** (0.040)	0.087** (0.039)	0.196*** (0.048)	0.128*** (0.045)
Block FE	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Field controls × Year FE	No	Yes	No	Yes
Control group mean	0.520		0.440	
Observations	541	541	541	541
R ²	0.082	0.936	0.128	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. Controls are indicators for joint-venture, sole-risk, area, onshore, block area, number of fields, and number of wells. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

columns (3)-(4) use operator share. Specifications (2) and (4) control for concession characteristics, including 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 a potentially important mechanism driving local advantage.

F.3 Corruption laws

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, the prospect of legal liabilities could plausibly deter multinationals from bargaining with law enforcement, even at arms length. If this is the case, we should observe that among multinationals, exposure to these laws should explain variation in levels of theft. By restricting the sample to multinationals, I remain agnostic about the content, quality, and enforcement of Nigeria's own anti-corruption laws.⁷²

⁷²This is preferable to assessing the effectiveness of these laws, which legal analysis suggest are basically ineffective (Aigbovo and Atsegbua 2013).

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

The results of this estimation are in Table A33 for oil output, theft, and local conflict outcomes. In general, foreign corruption laws have limited effect on the actual production decisions of the firm: the estimate with controls in column (2) is near zero and insignificant. However, in columns (3)-(4), 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 3.8-7.6-incident increase in theft, significant at the 1% level. A similarly large increase of 0.7-1.3 conflict deaths is shown in columns (5)-(6).

Table A33: The effect of corruption costs on output and criminality

	(1)	(2)	(3)	(4)	(5)	(6)
Outcome	Output		Oil theft		Conflict deaths	
Home-country corruption law	0.684** (0.279)	0.069 (0.330)	7.593*** (0.925)	3.808*** (0.908)	1.260*** (0.217)	0.744*** (0.265)
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
Control group mean	2.069		8.826		0.503	
Observations	2111	2111	2679	2679	2679	2679
R ²	0.873	0.888	0.729	0.766	0.277	0.384

Standard errors in parentheses are clustered at the field level. Sample is all untreated field-years from 2006-2016 (i.e., operated by multinational firms). Output is measured in millions of barrels of oil per year. Oil theft is the total number of sabotage spills within 15 km of the field. Conflict deaths is the total number of violent conflict-related fatalities within 15 km of the field as reported by local news media sources. 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

G Alternative explanations

G.1 Spatial spillovers

Gangs may not operate as local monopolists, as assumed in the model. In a general equilibrium setting, gangs may optimally choose targets for theft across all possible oil fields. 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.⁷³ In contrast, if security is at least partially non-excludable, increased anti-crime enforcement by security forces may have positive spillovers to nearby multinational firms. 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 et al. 2014 and Diamond and McQuade 2019). In the stacked dataset (see Appendix D.3), 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 90. 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 90 kilometers away from a treated field acts as the “pure” control group. I define theft outcomes in a 10 km radius around the field in order to minimize overlap which induces a mechanical spatial correlation in outcomes and therefore spurious spillover effects.

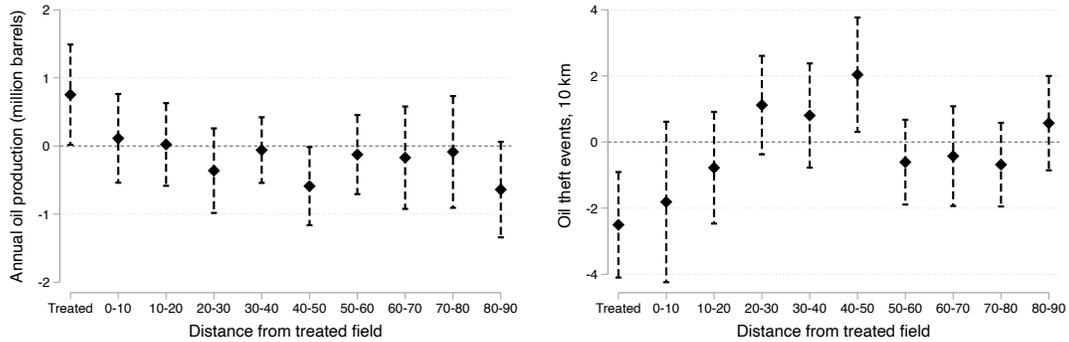
The results are in Figure A23, which plots the treatment effect, as well as coefficients at each ring from 0-10 to 80-90 km, for output and oil theft (left and right panels, respectively).⁷⁴ In both cases, the main treatment effects remain strong. This implies that mechanical spatial correlation in theft outcomes is not generating spurious treatment effects, since we have large treatment effects when the control group is defined as only fields further than 90 kilometers away. Furthermore, the output effect does not appear driven by declines on control fields. In fact, there is minimal evidence of substantial spatial spillovers across either outcome for nearby or faraway fields. There *are*, however, statistically significant spillovers in the 40-50 km bin, suggesting a crime displacement effect resulting in less output. This is reasonable, since positive security externalities and increasing costs of transport might limit displacement effects

⁷³This is similar to the standard displacement effect in the economics of law enforcement ().

⁷⁴Note that the 0-10 km coefficient in the oil theft panel will exhibit mechanical spillovers because of spatial correlation and should be disregarded.

for nearby and faraway fields, respectively. As such, we might expect negative spillovers to occur in the middle range, which we observe. Nonetheless, they do not affect the main treatment effects.

Figure A23: Spatial spillovers



Note: Figure plots coefficient estimates of treatment effect and spillover effects for output (left panel), and oil theft (right panel), defined as the total number of sabotage spills within 15 km of the field. Estimates are derived from a stacked difference-in-differences regression of the outcome on a dummy for post-treatment interacted with indicators for the treatment and “ring” distances from the nearest treated field. Omitted control group is untreated fields further than 90km from the nearest localized field. All specifications include stack, time, and field fixed effects and their interactions, as well as interacted controls for latitude of the field centroid, distance to coast, distance to Niger River, and distance to the capital. Standard errors clustered at the field level.

G.2 Discount rates

The local advantage in production may be driven not by organized criminal activity and law enforcement corruption, but rather by different optimal extraction profiles given underlying time preferences. This is a plausible mechanism if local companies have shorter time horizons than multinationals. I argue that 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.

However, I test this argument directly by estimating extraction profiles for local and multinational fields. In petroleum engineering, oil production typically follows what is called a “decline curve,” which models oil output as an exponential decay function over time (). The curvature of this extraction profile suggests an implied discount rate, given field characteristics – steeper declines suggest over time less patience. As such, impatient local firms should extract more oil earlier in the life cycle of the field.

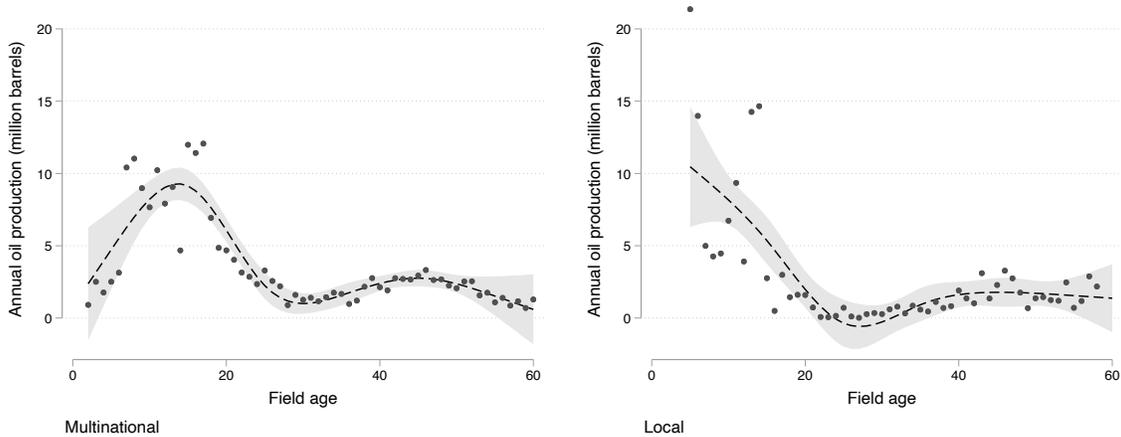
I estimate and plot decline curves in Figure A24 for the subsamples of multinational and locally-operated fields separately. Instead of directly estimating the parameters of an expo-

ponential decline curve, I model output as a flexible nonlinear function of age, following

$$y_{it} = \alpha + g(a_{it}) + \epsilon_{it}$$

I estimate g using a cubic spline with 7 knots spaced evenly every ten years.⁷⁵

Figure A24: Extraction curves by type, cubic spline



Note: Figure plots extraction curves – the level of oil output by the age of field – for subsamples of multinational (left panel) and local fields (right panel). Dots indicate mean output by age for each subsample. Fit is estimated using a cubic spline with the following knots: $[0, 5], [5, 15], [15, 25], [25, 35], [35, 45], [45, 55], [55, 60]$. Field age is defined as the difference between the current year and the year of the first well drilling.

The results indicate a clear decline curve for both multinational local firms. These curves both peak between 10-20 years of field life near 10 million barrels annually and decline steeply thereafter, approaching zero between 20 and 30. Both curves also suggest a small revival of in the later years of the field lifecycle, perhaps driven by new well drilling.⁷⁶ The only qualitative difference between the patterns of the extraction profiles by firm type is that multinational fields experience initial growth in production, as output comes online up in the early years of the field. This initial growth is absent in the local fields, likely because local firms tend to acquire fields that have already been developed, rather than develop new fields. As such, there are no field-year observations with ages below five in the local sample. The early part of the curve is therefore missing in this subsample.

Table A34 estimates quantitatively whether the slope of the extraction profile varies systematically and significantly across local and multinational fields. For ease of interpretation, I use a linear spline with the same knots as in Figure A24. The coefficients represent the slopes

⁷⁵The intervals for the cubic polynomial are $[0, 5], [5, 15], [15, 25], [25, 35], [35, 45], [45, 55], [55, 60]$

⁷⁶Note that decline curves are typically modeled at the well-level, though we only have output data at the field-level.

Table A34: Extraction curve estimates, linear spline

	MNC	Local	Diff
	(1)	(2)	(3)
Knot 1, [0,5]	2.280 (2.999)		
Knot 2, [5,15]	0.429 (0.553)	-0.485 (0.416)	-0.914 (0.689)
Knot 3, [15,25]	-0.831** (0.391)	-0.657 (0.515)	0.175 (0.642)
Knot 4, [25,35]	-0.037 (0.086)	0.181 (0.160)	0.217 (0.181)
Knot 5, [35,45]	0.169** (0.067)	0.145 (0.116)	-0.025 (0.134)
Knot 6, [45,55]	-0.154** (0.074)	-0.125 (0.132)	0.029 (0.150)
Knot 7, [55,65]	-0.178 (0.221)	0.564 (0.761)	0.742 (0.848)
<i>F</i> -statistic			1.516
<i>p</i> -value, joint test			0.173
Observations	2036	356	2392
R^2	0.070	0.153	0.078

Standard errors in parentheses are clustered at the field level. All models estimate linear spline models of output as a function of field age, with knots indicated in table. Subsample is in table header. Column (3) shows difference in slopes from an interacted specification. Column (2) drops the first spline because it is outside of the support of age for the local subsample. *F*-statistic and *p*-value refer to the joint test on all interaction terms. Output is measured in millions of barrels of oil per year. Field age is defined as the difference between the current year and the year of the first well drilling. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

of local linear regressions within each knot, estimated separately for multinational fields (1) and local fields (2). Column (3) estimates a model that interacts $g(a_{it})$ with $local_{it}$ to identify the difference in slopes for each knot. I also include the results of an *F*-test of joint significance across all interaction coefficients in column (3). Note that coefficients for the first knot in columns (2) and (3) are missing because of the lack of support of a_{it} in this range for the local field subsample. The results indicate that while the piecewise slopes of the decline curve not identical, these differences are not statistically significant, either individually or jointly.

G.3 Grievance

Criminal and militant activity may be driven by grievance rather than economic motives (Buhaug et al. 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 21% of all fields at any point during the sample period. In Table A35, 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.

Table A35: The effect of divestment on riots and protests

Outcome	All protests		Oil protests		Riots	
	(1)	(2)	(3)	(4)	(5)	(6)
Local firm	0.273 (0.168)	0.144 (0.170)	0.004 (0.017)	-0.014 (0.022)	0.193 (0.306)	0.155 (0.318)
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	3183	3183	3183	3183	3183	3183
R ²	0.353	0.410	0.160	0.238	0.401	0.452

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. Outcome variable is given in table header, defined as the total number of incidents 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

G.4 Local employment and welfare

Part of the rationale behind indigenization is that local firms may generate positive spillover effects 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. All else equal, this increases the bargaining space between firms and law enforcement, since gangs have less profit with which to offer competing bribes.

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-age⁷⁷ 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

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

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 capita 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, annual rainfall, and a rural indicator. 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 A37. 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 occur.

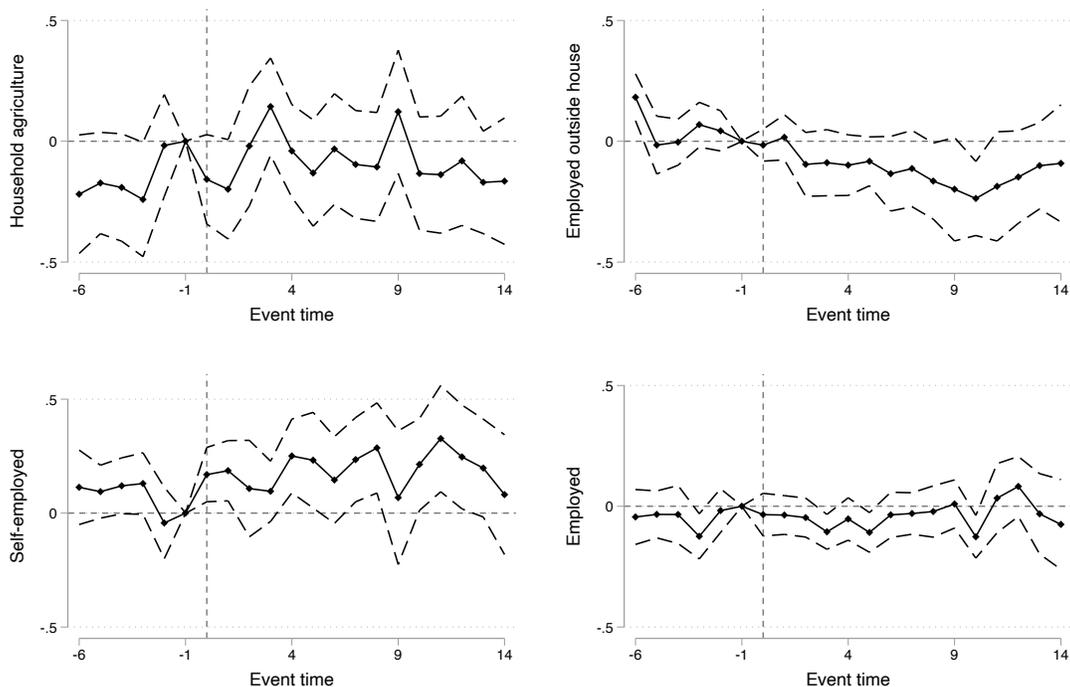
The results show no effect on the level of employment (Panel A). For the composition of employment, I do not find any statistically significant changes in employment outside the home (Panel B) or employment in household agriculture (Panel D), although the point estimate for both of these outcomes are consistently negative. However, there does appear to be an increase in self-employment (Panel C) by roughly 9-10 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. In Panel E, I test the impact of localization on log household per capita consumption. Again, there are no statistically significant effects. Overall, there is no evidence that localization creates meaningful positive economic spillovers for nearby oil-producing villages.

I test for parallel pre-trends in Figure A25. All results suggest that pre-trends are essentially flat and insignificant for each outcome considered in Table A37. 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.

Opportunity costs for young men in particular are likely to determine wages offered by organized crime. If employment effects are heterogeneous across demographics, then it may

⁷⁷Defined as ages 15-60.

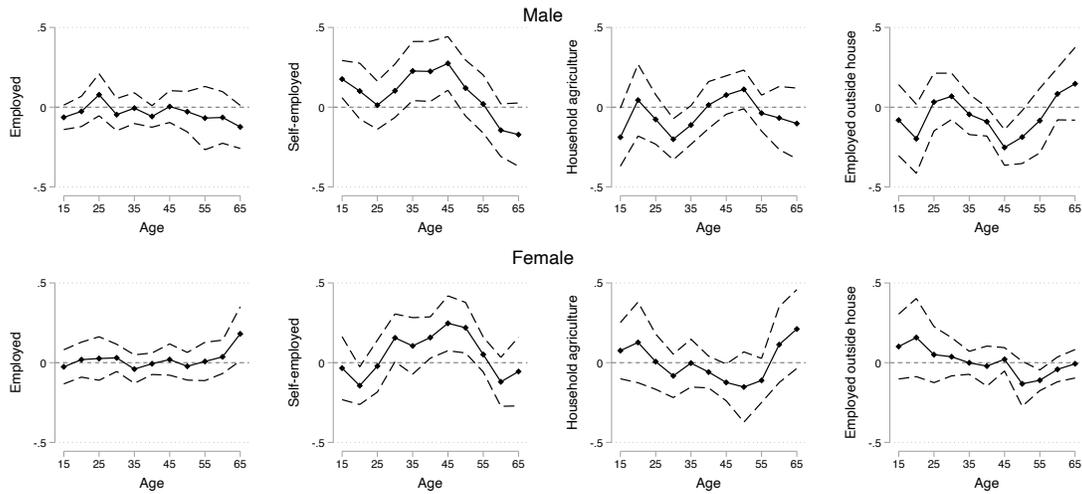
Figure A25: Local ownership and local employment, parallel trends



Note: Figure displays coefficients of TWFE event-study regressions of employment outcomes on pre-and-post treatment indicators for localization, conditional on unit and year fixed effects 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.

be that the aggregate zero effects are masking effects on the demographic groups relevant for the gangsters' cost structure. To test this hypothesis, I re-estimate the employment equation of each outcome by ten-year age bins and gender. The results are shown in Figure A26, which plots coefficients by age bin and gender for each outcome. The results indicate robust zeroes along each outcome and for each age group and gender. The only exception to this pattern is an increase in self-employment among middle-aged men and women, which is offset by a reduction in agricultural employment for the same demographic groups. However, for both men and women, the aggregate employment effects are zero at all ages. It is therefore unlikely that the effect of localization on theft and output is operating through opportunity cost mechanisms.

Figure A26: 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. All estimates include controls for distance to road, population center, market, border, and administrative center, a rural dummy, slope, elevation, and mean annual temperature and precipitation.

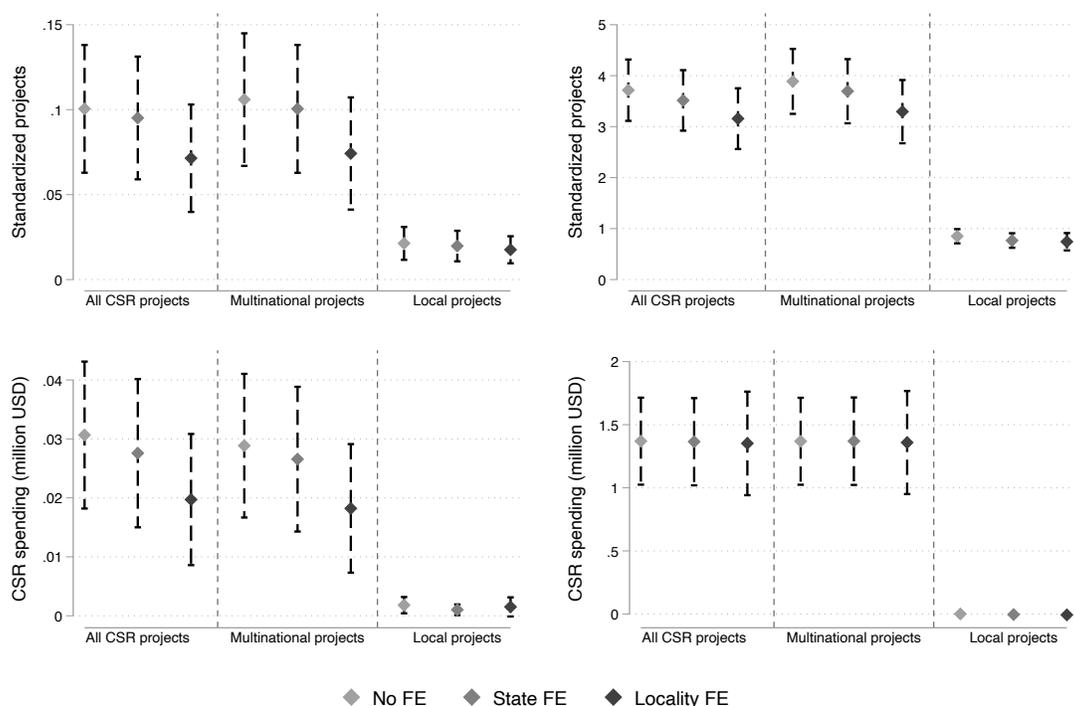
G.5 Corporate social responsibility

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 may observe a positive correlation between conflict and CSR expenditures.

Figure A27: CSR projects and local conflict



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 (top panel) or total expenditure (bottom panel), 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 (right panel) or the cumulative number oil-related militant attacks from 1997-2015 (left panel). Model specification is indicated in subfigure headers. Models are either unconditional or include state or locality fixed effects, as indicated in legend.

Figure A27 plots coefficients from these regression models. The top panel uses standardized CSR projects as the outcome to account for the fact that local firms are generally smaller and therefore have fewer projects overall, while the bottom panel use total CSR expenditure in millions of USD. The left panel use cumulative attacks up to 2015 on the righthand side, while the right panel uses attacks in 2015. For each specification, I estimate the bivariate relationship unconditionally, and with state or locality fixed effects. In general, there is evidence suggestive of targeting – prior 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 multinationals. 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 than local firms because they are unable to leverage political connections to bargain directly with gangs.

G.6 State violence

The welfare consequences of localization may be ambiguous or even negative if there are negative spillovers of increased law enforcement activity on local populations. Increased presence of the armed forces and militarization of Niger Delta communities at the behest of local oil firms may lead to more violence against civilians and human rights abuses. I test for these spillovers in Table A35 by estimating the main TWFE models using various measures of state violence from ACLED as the outcomes. Columns (1)-(2) use total incidents of state violence against civilians, while columns (3)-(4) look at civilian casualties from these events. I find no evidence of an increase in state violence against the local civilian population. If anything, there is a slight *decrease* in the level of state violence against civilians, significant at 10% in the specifications with control variables.

Table A36: The effect of divestment on state violence against civilians

Sample Outcome	Full		Onshore	
	Events	Deaths	Events	Deaths
	(1)	(2)	(3)	(4)
Local firm	0.070 (0.092)	-0.124* (0.073)	0.061 (0.101)	-0.199** (0.101)
Field FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Controls × Year FE	Yes	Yes	Yes	Yes
Observations	3183	3183	2296	2296
R^2	0.479	0.213	0.489	0.261

Standard errors in parentheses are clustered at the field level. Sample is the panel of 314 oilfields from 2006-2016. Outcome variable is the total number incidents or fatalities from state violence against civilians 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 A37: Local ownership and local employment

Sample	All				Onshore	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Employed</i>						
Local firm	0.024 (0.028)	0.024 (0.030)	0.001 (0.025)	0.045** (0.019)	0.022 (0.032)	0.000 (0.029)
Observations	16827	16827	16827	16827	15616	15616
R ²	0.033	0.034	0.041	0.039	0.029	0.038
<i>Panel B: Employed outside home</i>						
Local firm	-0.020 (0.020)	-0.017 (0.019)	-0.033* (0.017)	-0.023 (0.018)	-0.021 (0.018)	-0.014 (0.020)
Observations	9225	9225	9225	9225	8551	8551
R ²	0.107	0.109	0.158	0.132	0.095	0.139
<i>Panel C: Self-employed</i>						
Local firm	0.098** (0.037)	0.097** (0.037)	0.108** (0.044)	0.096*** (0.033)	0.109** (0.044)	0.102** (0.049)
Observations	9225	9225	9225	9225	8551	8551
R ²	0.070	0.071	0.122	0.106	0.068	0.123
<i>Panel D: Employed in household agriculture</i>						
Local firm	-0.050* (0.029)	-0.050 (0.031)	-0.025 (0.037)	-0.008 (0.037)	-0.078** (0.031)	-0.055 (0.038)
Observations	9225	9225	9225	9225	8551	8551
R ²	0.152	0.153	0.271	0.183	0.151	0.277
<i>Panel E: Household consumption</i>						
Local firm	0.035 (0.072)	0.038 (0.068)	0.038 (0.073)	0.028 (0.045)	0.038 (0.079)	0.030 (0.069)
Observations	5119	5118	5119	5119	4750	4750
R ²	0.243	0.244	0.292	0.270	0.250	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 × 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, except Panel E, which is at the household-level. 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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.