Delta boys: Bargaining, war, and black market oil in Nigeria*

Jonah Matthew Rexer† Even Comfort Hvinden‡

September 22, 2020

Abstract

We study how a ruling elite facing conflict over natural resources allocates rents to rebel groups and how the structure of the elite's settlement with rebels affects post-conflict resource theft. Using original data on the locations, alliances, black market activities, and attacks of militant commanders in the Niger Delta conflict, we find that a peace deal led to large declines in violent attacks on the oil sector, but also sustained growth in the black market for stolen oil. We use a model of dynamic bargaining under imperfect information and limited commitment to explain why inefficient conflict and oil theft persist in equilibrium. The model predicts that the government may optimally allow theft of resources by strong rebels in locations with low opportunity costs of black market activity. We test and find support for these propositions in the data. Our analysis highlights how the industrial organization of black markets and military dynamics jointly shape incentives for participants in resource conflicts.

*We thank Santosh Anagol, Jørgen Juel Andersen, Arthur van Benthem, James Fenske, Guy Grossman, Magnus Våge Knutsen, Elisa Macchi, Thomas McKay, Ragnar Torvik, and seminar participants at the University of Pennsylvania, the Wharton School, the Norwegian University of Science and Technology, the BI Norwegian Business School, and the 2020 Development and Political Economy Graduate Student Conference for useful comments. The conflict bargaining model was considerably improved by feedback from Joseph Harrington Jr. This project was funded by the Kleinman Center for Energy Policy at the University of Pennsylvania, the Centre for Applied Macroeconomics and Commodity Prices at the BI Norwegian Business School, and a grant from Ann Harrison. All errors are our own.

†Wharton School of Business, University of Pennsylvania.

‡Department of Economics, BI Norwegian Business School.

1 Introduction

The ruling elites in institutionally weak, resource-rich states frequently face the threat of violent insurrection over the distribution of resource rents.[1] To resolve armed conflict and deter theft of natural resources the elites must generally implement a rent-sharing agreement with rebel groups.[2] Yet for resource-cursed economies, peaceful rent-sharing is fragile and black markets are rife.[3,4] How can we explain the persistence of violence and criminality in these settings, despite strong incentives to restrain costly conflict and resource theft?

Elites and rebel groups in poor countries typically face environments replete with bargaining frictions. We hypothesize that asymmetric information ex-ante, the inability to commit ex-post, and the lure of black market rents interact to hold up mutually beneficial bargains. In particular, we argue that tacit acceptance of resource theft, enforced by credible threats of violence, constitutes a de facto rent-sharing mechanism for ruling elites that maintains peace while allowing criminality to flourish.

We test this proposition in a case study of the Niger Delta conflict, a decade-long insurgency by militant rebels in Nigeria's oil rich south accompanied by a thriving black market for oil stolen by rebels and criminal organizations. In 2009, a peace deal (hereafter, the amnesty) provided legal immunity to combatants – and sizable private payoffs to some rebel leaders – in exchange for halting attacks on oil infrastructure. Despite success in reducing violence, the Niger Delta nonetheless experienced substantial growth in oil theft following the amnesty, leading some observers to question whether Nigeria's rulers were "[rewarding] bad behavior" (Economist 2016). We formalize the problem of dynamic bargaining over resource rents in the presence of a post-conflict black market. We show how the inability to commit,

[1] An extensive literature documents robust links between violence and resource endowments, that fighting incidence rises with the price of contested minerals and that relative prices matter in structuring the benefits and costs of conflict, see Dube and Vargas (2013), Bazzi and Blattman (2014), Berman et al. (2017), Nwokolo (2018), Fetzer and Kyburz (2018). Sierra (2019) studies the impact of resource price shocks on state capacity investment by armed actors in Eastern Congo. There is also a literature that studies civil war onset through country-level regressions, typically finding significant correlations between resource endowments and conflict; see, among others, Fearon and Laitin (2003), Bannon and Collier (2003), Collier and Hoeffler (2004), Collier and Hoeffler (2005), Hegre and Sambanis (2006), and Collier, Hoeffler, and Rohner (2009), and Sala-i-Martin and Subramanian (2013) for the case of Nigeria specifically.

[2] Inefficient competition over resource rents has been identified as the key mechanism behind the resource curse, see Acemoglu, Verdier, and Robinson (2004), Hodler (2006) and Robinson, Torvik, and Verdier (2006), among others.

[3] In a data set of 282 civil conflicts spanning 1946-2006 from the Uppsala Conflict Data Program, resource-related conflicts are much less likely to end in government victory, and more likely to result in a peace deal. And yet, despite strong incentives to end costly conflict, these deals often fail to establish durable peace – resource conflict episodes are more persistent than other forms of civil war by nearly 4.5 years (Rustad and Binningsbø 2012). Resource conflicts also have more reversion to fighting after a deal – 65% of the resource conflicts featured a broken deal, compared to 40% of the non-resource episodes.

[4] The value of crude oil stolen globally has been estimated to be in the region of 100 billion U.S. dollars per year, see Ralby (2018) and Ralby and Soud (2018). Extensive black markets for stolen crude reportedly exist in, among others, Azerbaijan, Ghana, Morocco, Mozambique, Nigeria, Thailand, Turkey and Uganda. Globally, illegal resource markets in gold, rare minerals, precious stones, illicit drugs, and forestry resources are common in conflict-affected countries.

asymmetric information about military capabilities, and low barriers to entry in the black market combine to generate conflict and resource theft in equilibrium. We derive testable predictions on how these outcomes vary with endogenous peace agreements and exogenous features of the strategic environment over time. To evaluate our theoretical predictions empirically we collect highly detailed data on the location, amnesty status and alliance network of militant commanders, as well attacks on the oil sector, crude oil production, and black market activity in nearby villages.

To motivate our theoretical framework we establish that the amnesty coincided with an immediate and persistent drop in violence targeting the oil sector and was followed by a sustained growth in oil theft. However, the increase in oil theft was highly heterogeneous across locations, muted in areas where local commanders were awarded relatively generous post-amnesty settlements.

To estimate a discontinuous change in the time-series of attacks at the time of amnesty we employ a regression discontinuity in time (RDiT) specification (Hausman and Rapson 2018). We show that oil-related militant attacks fell from 3-5 per month at the exact date of the amnesty announcement to nearly zero in the immediate post-amnesty period. The estimate is robust to polynomial time trends, varying event windows, month- and year fixed effects, outlier sensitivity, and optimally-selected autoregressive terms to capture conflict persistence. We also conduct several falsification tests, showing that the effect is driven entirely by amnestied areas, does not obtain for placebo event-dates, and is absent for non-oil related conflict. Lastly, we show that our event date identifies a structural break in the time series (Perron 2006.)

We use data on oil theft from Nigerian environmental regulators to estimate in differences-in-differences regressions which show that amnesty differentially increases oil theft in amnestied areas by 40% of the sample mean. Event-study regressions reveal both parallel trends and dynamic treatment effects, with oil theft growth largest 2-6 years after amnesty. The results are robust to differential oil price responses, controls interacted with time dummies, different definitions of an amnestied area, and village-specific time trends. The results are also robust to several definitions of the control group, including a synthetic differences-in-differences estimator (Arkhangelsky et al. 2019). A specification curve (Simonsohn, Simmons, and Nelson 2019) reveals that the vast majority of point estimates from reasonable specifications are positive and significant. We also rule out several alternative explanations for the results. Amnesty might simply have improved the ability of government regulators to collect data on oil spills in conflict areas, creating the illusion of growth in theft. We show that non-sabotage oil spills – subject to the same data limitations – do not exhibit differential growth. We also normalize theft by output to show that the results are not driven by a post-amnesty rebound in oil production. Lastly, we rule out that the trend is driven by some source other than organized crime, such as local communities sabotaging oil infrastructure in protest or for compensation. Disaggregating impacts by asset type, we show that the increase in sabotage events is concentrated only on the highest-value pipelines, suggesting theft of oil by profit-maximizing criminal groups, rather than unorganized community sabotage.

Third and finally we show that the growth in oil theft was muted in locations where commanders were awarded lucrative "pipeline surveillance contracts" (Stakeholder Democracy 2019c, hereafter security contracts) in which the government transferred sizeable sums in exchange for protecting local pipelines from attack.[5] In particular, areas controlled by rebels who received security contracts see a nearly 75% reduction in oil theft relative to amnestied rebels that did not, suggesting that the contracts were indeed effective at moving rebels out of the black market. The data suggests that the amnesty and pipeline surveillance contracts successfully reduced violence and oil theft, raising the question of why they were not granted more widely. Yet clearly the estimates do not have a direct causal interpretation because the receipt of amnesty and security contracts is endogenously determined. To shed light on the selection mechanism we derive a model of dynamic conflict bargaining.

In each stage game the ruling elite proposes a division of rents between itself and the rebel. Rebels may then reject the deal, yielding a period of conflict, or accept it, and then choose whether or not to steal oil. If there is rebel theft the government may cease payments or attack. The rebels faces a positive probability of military defeat when fighting, the event of which ends the game. Rebels are a dominant firm that face a competitive fringe of local criminal entrepreneurs in the black market for stolen oil. When rebels do not steal oil the fringe increases their theft proportionally to local cost conditions. The total reduction in oil theft induced by providing incentives to rebels depends on the endogenous response of the fringe.

The full information game features a unique subgame perfect equilibrium, up to off-path actions. However, variation in relative rebel military capability and the efficiency losses from oil theft gives rise to three distinct forms of equilibrium behavior. In equilibria *without incentives* elites allow oil theft but demand bribes from the rebels. This behavior obtains when the competitive fringe is highly responsive, so that oil theft in any case is expected to be high, or when the efficiency losses from rebel oil theft are low. In this case, elites prefer to deal with a dominant firm from which they can recapture surplus. If instead the efficiency losses from theft are high or the fringe is only moderately responsive, the ruling elite will provide costly incentives to deter rebel oil theft. If payments can be promptly halted, elites grant rebels a rent proportional to their income from theft and punish deviations by a temporary suspension of payments. We call these incentives *low-powered*, as in equilibrium the rebels are indifferent between stealing oil and not. If payments are not easily halted, elites

[5] In an illustration of the sums involved, the World Peace Foundation (2018) reports that Government "Tompolo" Ekpemupolo, a leading militant commander and surveillance contract recipient, used his funds in 2012 to purchase six *Hauk*-class motor torpedo boats and a large support craft formerly of the Royal Norwegian Navy for some $12 million U.S dollars.

must punish oil theft by ceasing payments *and* fighting, since military intervention is the only viable sanction. For the threat of fighting to be credible, subgame perfection demands that the government's offer to the rebel is proportional to the their cost of fighting. In this equilibrium with *high-powered* incentives, rebels face generally greater payoffs, more severe punishment, and strictly prefer not to steal oil. A key implication of our model is that when the black market fringe is responsive and rebels are militarily capable the government may optimally allow rebel oil theft even when this activity entails considerable efficiency losses.

In the full information setting there is no conflict. We therefore study perfect Bayesian equilibria of a game where the rebel is privately informed of their survival probability, rendering the ruling elites uncertain of the return to fighting. When coupled with the elite's inability to commit, this uncertainty over the prospects of military victory creates powerful incentives for weak rebels to mimic the strong, rejecting inferior offers and fighting in the hope of surviving to enjoy greater future pay-offs – a mechanism known as the *ratchet effect* in the literature on incentive contracts and non-democratic politics.[6] We show that the game admits an equilibrium unique up to the value of rejected offers and off-path actions. Weak rebels anticipate that for every period of fighting they survive, the elite revises upwards their belief of the expected cost of conflict. After a finite number of fighting periods the governing elite prefers to offer any rebel the strong type's subgame perfect equilibrium contract to further conflict. If weak rebels believe the amount of fighting required to achieve a generous payoff is small they demand a steep price to reveal their type. The government understands that if the compensatory transfer is too generous the strong will mimic the weak. In that case there exists no offer that separates rebel types. Restricted to pooling offers, an optimistic government initially prefers to fight, and there is conflict on the equilibrium path. In the baseline model rebels either receive generous amnesty or are defeated, a feature arguably inconsistent with the observed heterogeneity in amnesty outcomes. Thus we consider a model with two-sided uncertainty in which rebels form higher-order beliefs on the government prior. There exist unique perfect Bayesian equilibria where weak types first fight and nevertheless accept a separating offer because they learn along the equilibrium path that the government has greater resolve than first anticipated.

The model predicts that militarily powerful rebels are more likely to receive amnesty, that oil theft is elevated in areas with competitive cost conditions, and that the ruling elite is more likely to tacitly accept oil theft from militarily powerful rebels in locations with competitive black markets. When local conditions are such that elites' equilibrium behavior switches from incentive provision to minimal amnesty, we expect non-linear changes in aggregate oil

[6]See for instance Freixas, Guesnerie, and Tirole (1985), Laffont and Tirole (1988), and Laffont and Tirole (1990). See also Weitzman (1976) and Holmstrom (1982). Information asymmetries in no-commitment environments are a key ingredient to non-democratic politics in general, and conflict bargaining in particular, see Fearon (2007), Chassang and Padro-i-Miquel (2009), and Dal Bó and Powell (2009). See also Spaniel and Bils (2018) and Shirkey (2016).

theft. We test and find support for these predictions in the data.

To verify whether heterogeneous military capabilities are associated with differentiated amnesty outcomes, we first construct and validate a measure of rebel military strength based on the network of local alliances along oil pipelines. Groups with a denser network of local allies plausibly have greater ability to provide mutual support, share information, and coordinate attacks on oil infrastructure. We show rebel groups with more local allies destroy significantly more oil output during the conflict, suffer battlefield defeat at lower rates, and are more likely to be amnestied and receive pipeline surveillance contracts. The results are robust to state fixed effects, measurement assumptions, and various geospatial controls. These cross-sectional correlations are consistent with predictions of canonical conflict bargaining models, and suggest differentiated amnesty based on military capability. Moreover, a central assumption of our model is that rebels retained their military capability post-amnesty. The implicit threat of violence generates the incentives for ruling elites to maintain costly incentive provision. In May 2016 the newly elected President Muhammadu Buhari announced a 70% reduction in amnesty payments, providing a natural experiment to test this hypothesis. Again using an RDiT model, we estimate the discontinuous change in rebel attacks following this event, and robustly find an increase in violence of roughly similar magnitude to the reduction in rebel activity following the amnesty.

Second, we show that the largest increases in post-amnesty oil theft are concentrated in areas where the costs of oil theft are low. We measure the cost of oil theft by labor costs in the illegal sector, using the median wages of young men. We find that low-wage labor markets experience much larger post-amnesty rises in theft. This does not simply proxy for level of development – placebo tests using wages of other demographic groups reveal that these patterns of heterogeneity are muted for old men, and absent for young women. Since a substantial share of stolen crude is exported, we also proxy for theft costs using export cost as measured by distance to the Atlantic coast. This analysis reveals similar patterns – growth in the black market is concentrated in lower-cost areas near the coast.

To deter oil theft with high-powered incentives the governing elites must pay the rebels in proportion to their military capability to render the off-equilibrium threat of fighting credible. When faced with militarily powerful rebels the government may then optimally provide rebels with minimal amnesty agreements that implicitly tolerate oil theft, efficiency losses notwithstanding. Using our local-alliance measure of military strength, we find that an additional ally increases the treatment effect of amnesty on oil theft by between 25-46% of the main effect. This is consistent with the prediction that militarily capable commanders might be optimally allowed to steal oil after amnesty. In a quadruple-difference estimation, we find that the role of military strength diminishes as theft costs rise, because the fringe response falls, increasing the value of a deal for government. This highlights the important interaction between the industrial organization of the black market and the military dynamics of the conflict in determining post-amnesty illegal activity.

Finally, we show that the patterns of heterogeneity in the data are driven by a threshold effect, consistent with the proposed equilibrium switching mechanism. Quantitatively, the effect of amnesty on oil theft at the 5th percentile of the wage distribution is 5 times larger than at the median, after which it falls to zero. In addition, we find that the effect of amnesty on theft is 4.7 times greater in the subsample of villages with the strongest militants than in aggregate; for the rest of the distribution it is bounded between zero and the aggregate effect size. These nonlinearities strongly suggest threshold points at which the ruling elites optimal policy shifts from costly incentive provision to tacit acceptance of oil theft.

Our work adds to a burgeoning empirical literature studying the endogenous formation and post-conflict effects of peace agreements. Both Konig et al. (2017) and Miceli (2017) highlight the importance of alliance networks among combatants in structuring the incentives for peace. Dancy (2018) studies the effectiveness of legal amnesties in peacemaking. A substantial literature also analyzes post-conflict conditions to sustain peace (Collier, Hoeffler, and Soderbom 2008, Fearon, Humphreys, and Weinstein 2009, Rustad and Binningsbø 2012). However, these studies do not offer empirical tests of the conflict bargaining hypotheses, and therefore fail to answer the fundamental question of why bargains emerge and endure. A related contribution is who Francois, Rainer, and Trebbi (2014) study ethnic power-sharing in African governments and highlight as we do the importance of credible threats of violence – in their case, coups – for enforcing bargains. Our results provide substantial support for the uncertainty resolution mechanism at the core of conflict bargaining models (Fearon 1995, Fearon 2007, Ellingsen and Miettinen 2008, Wolford, Reiter, and Carrubba 2011, Thomas, Reed, and Wolford 2016, Ghosh, Gratton, and Shen 2019) while highlighting the role of black markets as a mediator between battlefield and post-conflict outcomes.

A wide-ranging theoretical literature emphasizes the role of inefficient rent sharing between elites as the mechanism sustaining kleptocracy (Grossman 1999, Acemoglu, Verdier, and Robinson 2004) and the resource curse (Ross 1999, Hodler 2006, Mehlum, Moene, and Torvik 2006, Robinson, Torvik, and Verdier 2006). Our analysis contributes to this literature by shedding light on how efficiency losses in rent-sharing are shaped by fundamental conditions. An extensive literature, reviewed in Ross (2015) and Nillesen and Bulte (2014), links the value of natural resources to civil conflict at national and subnational levels. However, this literature tells us little about the mechanisms that jointly determine conflict and black market outcomes, and, despite their important role in resource conflicts, generally overlooks black markets entirely. At the same time, a growing literature studies the industrial organization and political economy of black markets more broadly, focusing primarily on drug trafficking (Angrist and Kugler 2008, Dell 2015, Gehrin, Langlotz, and Kienberger 2018). Along with Romero and Saavedra (2019), we extend this literature to the study of illegal natural resource extraction. This literature has studied the role of law enforcement policy, formal sector taxation, and price shocks in structuring black market incentives. We are the first to study the political economy of conflict and illegal activity jointly, highlighting the importance of detailed local data on both conflict participants and black market outcomes.

The paper proceeds as follows. Section 2 details the Niger Delta conflict, while 3 explains our data sources and measurement assumptions. In Section 4, we provide reduced-form estimates of the causal effect of amnesty on conflict and criminal activity to motivate our theoretical framework. We introduce our conflict bargaining model in Section 5 and derive predicted behavior. Finally, we test for heterogeneous outcomes along the model's predicted dimensions in 6. Section 7 concludes with a discussion of avenues for future research.

2 Background

2.1 The Niger Delta crisis and amnesty

The Niger Delta is Nigeria's oil-rich southern region, responsible for essentially all of the country's 2.3 million barrels per day of oil output. Beginning in the early 2000s, longstanding unrest over the distribution of oil revenues erupted in armed conflict as militant groups took aim at the Nigerian federal government and multinational oil companies. Initially, militants operated independently from camps deep in the mangrove forests of the coastal Niger Delta, favoring abductions of foreign oil workers, bombings of oil pipelines, and direct engagement with the Nigerian military. In 2005, the Movement for Emancipation of the Niger Delta (MEND) united key rebel commanders, leading to a period of escalation in the conflict, most notably the 2008 Bonga offshore platform attack, which took 10% of Nigerian oil production offline in a single day. During the height of the conflict from 2005-2009, annual Nigerian oil production fell by 15%. While all groups professed demands for greater resource control for the region, a thin line divided most militant groups from organized criminal gangs trafficking in guns, drugs, and stolen oil (Asuni 2009).

In July of 2009, the Nigerian federal government announced that members of Niger Delta militant groups would receive amnesty in return for disarming and demobilizing. The amnesty covered roughly 20,000 combatants – primarily young men – 86% of whom came from the core Niger Delta states of Delta, Bayelsa, and Rivers. These fighters – colloquially known as "boys" – received large monthly cash stipends, participated in numerous reintegration programs, and were given scholarships for education and training. In our sample of 41 militant commanders, 18 did not receive the amnesty. Of these, 7 were defeated prior to amnesty, while the remaining 11 either were not included, rejected amnesty, or are of unknown status. In 2012, additional payments were made to 7 top-tier militants commanders in the form of "pipeline surveillance" contracts. These contracts, granted by government to security firms controlled by militant leaders, were ostensibly used to safeguard pipeline infrastructure from sabotage, but in reality served as another mechanism of transferring rents to militant leaders (Stakeholder Democracy 2019c).

2.2 The black market for stolen oil

Theft of crude oil for sale on the black market – known as oil bunkering – is a fixture of Nigeria's petroleum sector. The region's 5000 kilometer network of oil pipelines traverses militant-controlled swamplands, effectively at the mercy of bunkering gangs. Oil thieves cut into pipelines using hacksaws, installing a valve connected to an illegal pipeline – known as a "hot tap" – which siphons oil to a nearby barge. Stolen oil is then sold either to local artisanal refineries for processing and sale on the local market, or to oil tankers in nearby offshore waters for export. The size of the black market is widely debated; in 2016, theft losses were estimated at 4.2 billion US dollars, or 15% of Nigeria's total production (NEITI 2016). The market has a two-tiered structure, with larger operators exporting stolen crude and smaller players selling to local illegal refineries (Katsouris and Sayne 2013, Stakeholder Democracy 2019a, and Stakeholder Democracy 2019c). The balance has shifted towards the local market, which comprised 75% of black market sales in 2017. Collusion between security forces, sub-contractors for oil companies, and allegedly local politicians allows the black market to function relatively unfettered (Stakeholder Democracy 2019b).

Militant groups play an important role in this market. Initially, bunkering emerged among militant groups in order to finance their military endeavors (Watts 2007). In the post-conflict period, however, the black market has grown substantially in size (See Figure A1, Appendix A), driven in part by entry of criminal gangs and black market entrepreneurs. Some ex-militants act as patrons to new gangs who wish to tap in areas still under their influence, while others have elected to participate directly; some have even left the market completely, entering politics or legitimate business.

2.3 The effect of amnesty

The amnesty for Niger Delta militants marks a considerable shift in the dynamics of violence and criminality in the region. The amnesty was preceded by a surge in militant attacks, and followed by sharp drop in violence and gradual increase in oil theft . Consider Figure 1, which plots trends of militant attacks (Panel A) and oil theft (Panel B) in amnestied regions relative to non-amnestied ones. These trends are estimated by quarterly event-study regressions of the outcomes on dummy variables for periods pre and post amnesty, interacted with an indicator for whether a village is within 30 kilometers of an amnestied militant camp.[7]

[7] All estimates control for location and time fixed effects, with standard errors clustered at the village level and the initial quarter serving as the omitted reference period.

FIGURE 1 HERE

Panel A shows that amnestied regions experience substantially different trends with respect to militant attacks. In particular, violence spikes in months preceding the amnesty deal, and falls dramatically thereafter, differentially so in militant-affected regions. Dummy variables for pre-event periods are positive, significant, and increasing as time to amnesty falls, while those for post-amnesty periods are negative or zero. The results demonstrate a sharp drop in conflict at exactly the time of the amnesty event. Panel B plots the event-study coefficients for oil theft. Differential trends in oil theft between amnestied and non-amnestied areas are zero and insignificant for most of the pre-amnesty quarters. However, oil theft rises steadily in the post amnesty period in amnestied areas. This difference spikes nearly 5 years after the amnesty and then declines thereafter.

3 Data and measurement

3.1 Data

We briefly present our data sources and variable construction. A detailed derivation of our measures and discussion of potential measurement error and robustness to arbitrary choices can be found in Appendix C.

Militant events: We use the Armed Conflict Location Event Dataset (ACLED) to measure attacks on the oil sector perpetrated by militant groups, e.g. bombings of major oil infrastructure, kidnappings and killings of oil workers, and battles against the Nigerian military. To identify all such events from 1997-2017, we conduct keyword searches, and further subset events containing the keywords to include only those attacks that were perpetrated by political militias or rebel groups.

Militancy and amnesty: Data on militant commanders was collected by the authors from several sources. In 2018, we visited Warri, Delta State, one of the Niger Delta crisis' epicenters. We first collated a list of militant commanders from previous qualitative work on the program, including Ugwu and Oben (2011), Ojakorotu and Dodd Gilbert (2010), and then consulted with AA Peaceworks (AAPW), a highly informed local non-profit organization. For each militant commander, AAPW provided the following information: *i*) the group that this commander was affiliated with, *ii*) the location of their camp(s), usually denoted by the exact creek or a nearby village, and *iii*) whether they accepted amnesty. Gaps in the data and verification of accuracy were addressed by consulting Nigerian newspapers. We supplement this dataset with a list of pre-amnesty militant camps collected in a similar exercise by Blair and Imai (2013) to yield 69 militant camps led by 41 unique commanders.

Using searches of local Nigerian newspapers, we further identify 7 militant commanders on our list who received government contracts to perform security services in the oil sector. We then define three outcomes at the militant camp level: *i*) if the camp received any amnesty at all, *ii*) if the camp received a "generous" amnesty, which includes pipeline security contracts, and *iii*) if the camp was defeated in battle. For each camp, mapped in Figure A2, we calculate the distance between that camp and all 11,607 villages in the sample. An area is considered amnestied for the purposes of the difference-in-differences analysis if it falls within a 30 kilometer radius of an amnestied camp. Areas are then further designated as receiving a generous amnesty if the commander of this camp received security contracts.

Military strength: The extent of damage that a militant group is able to inflict on oil production – a key parameter in the model – is unobserved. To proxy for the underlying strength of a militant camp, we note that the Niger Delta conflict is known for a complex web of alliances between militant commanders and that rebel strength is inherently localized by its dependence on access to physical infrastructure targets. We measure camp-level military strength by the intersection of militant alliance and physical pipeline infrastructure networks; our measure is the number of allied camps within 10 kilometers along a pipeline. The network of alliances and pipelines is visualized on Figure 2. This map plots the pipeline infrastructure, with each militant camp attached to its nearest pipeline. Red triangles indicate oil export terminals, the geographic endpoints of a pipeline. Blue triangles indicate the location of militant camps along pipelines and are scaled by the number of local allies.

FIGURE 2 HERE

We describe and validate this measure in detail in Section 6.1. As a first pass, however, we show in Figure A13 that the number of local allies around a militant camp is highly correlated with the decline in oil production in that camp's area of control during the peak years of the conflict – while the number of non-allies exhibits no such correlation – suggesting that our measure is a reasonable proxy of underlying capacity to destroy output.

Oil production and infrastructure: Data on oil output comes from the administrative records of the Department of Petroleum Resources (DPR) and the Nigerian National Petroleum Corporation (NNPC), made available to the authors by these government agencies. These data contain field-level monthly production quantities for all oil fields in Nigeria from 1999-2015, excluding 2009, where data is unavailable. Data on the geographic location of oil and gas infrastructure comes from the DPR and Google Maps. Out of 337 oil fields that ever produced in the administrative data, we are able to georeference 314, or 93%. The data also feature a pipeline network of 4,284 kilometers. A map of the oil infrastructure, along with the locations of militant camps and their amnesty status, can be found in Appendix A, Figure A2.

Oil theft: Information on the time, location, and details of 11,327 georeferenced oil spills covering 2006-2017 comes from the administrative records of the the National Oil Spill Detection and Response Agency (NOSDRA), made publicly available on their Oil Spill Monitor. Bunkering incidents are defined as spills that occurred as an act of sabotage.

4 Motivating evidence

The event-study in Figure 1 (Section 2.3) suggest that the amnesty plausibly caused a decrease in violence and an increase in oil theft. To build confidence in the presence of a causal effect we establish two robust patterns in the data. First we use a regression discontinuity in time (RDiT) approach on a monthly time series to demonstrate a sizable and durable reduction in militant violence targeted at the oil sector immediately after the amnesty (Section 4.1). Second we show in a difference-in-differences analysis (DD) that the incidence of black market oil theft increased substantially and differentially in locations controlled by amnestied rebels, but not in areas where the rebel received costly surveillance contracts from the government (Section 4.2). Batteries of robustness tests lend further credence to the notion that observed patterns of violence and oil theft indeed reflect a causal effect of amnesty. Yet because rebel groups were amnestied endogenously the magnitude of estimated coefficients do not have a straightforward causal interpretation. The patterns of militant violence, government incentive provision, and oil theft combine to motivate our theoretical framework (Section 5) which formalizes the selection mechanism.

4.1 Conflict and amnesty

To investigate the dynamic short-run relationship between militant attacks and amnesty we estimate a regression discontinuity in time (RDiT, Hausman and Rapson 2018) using a monthly time series of militant attacks on the oil sector.[8] We use an RDiT model instead of a DD because the assumption of parallel trends for conflict is unlikely to be satisfied, for instance if militant attacks reflect strategic signaling behavior. Therefore, we estimate

$$m_t = \alpha_1 + \rho m_{t-1} + \theta 1(t \geq \tau) + g(t) + 1(t \geq \tau) h(t) + \delta_m + \epsilon_t$$

for $t \in [\tau - \Delta, \tau + \Delta]$, where Δ is the bandwidth, or event-window, which we vary across specifications, and $g(t)$ and $h(t)$ are flexible functions of time that vary with the post-amnesty indicator $1(t \geq \tau)$. To account for autocorrelation in the outcome, we allow for an AR(1) process by including m_{t-1}, with ρ the autocorrelation coefficient.[9] Finally, δ_m is a fixed effect for either month-of-year or year. The former captures seasonal within-year cycles in oil

[8]This method has also been called both an interrupted time series or event-study model. However, following convention in applied microeconomics, we use the term "event-study regression" to refer to difference-in-differences with dynamic treatment effects in a panel data setting.

[9]Higher-order processes do not materially affect the results.

attacks that might be correlated with the month of amnesty, while the latter accounts for year-specific spikes in attacks that may also be correlated with amnesty timing.[10] The identifying assumption is that the time trend of attacks would be continuous over the event-date in absence of amnesty. Practically, this requires that other announcements or events DD not occur in the month of amnesty to cause a precipitous drop in conflict. We also require that the amnesty has immediate effects on attacks, since the treatment effect identified in the RDiT is inherently short-term.

The result of the event-study is presented graphically in Figure 5, which plots monthly outcomes over time with a flexible local polynomial fit estimated separately on either side of the July 2007 cutoff, for an event window of $\Delta = 30$ months.

FIGURE 5 HERE

The results indicate a clear and statistically significant reduction in militant violence immediately following the amnesty. Given a relatively small window, and the observed discontinuity at the exact event-date, it is unlikely that temporally correlated non-amnesty shocks are causing the drop.

Following recommendations in Hausman and Rapson (2018), we estimate the parametric event study results with numerous robustness tests in Table 1. Given the small number of aggregate time-series observations, there may be worries about the sensitivity of estimates to model specification. We consider numerous specifications of the event-study estimating equation, allowing the polynomial to vary from linear to third-order, and test the following event windows: the full sample, and within 30, 20, and 10 months. We also estimate AR(1) specifications to account for autocorrelation in the time series. Table 1 contains regression results for all of these specifications.

TABLE 1 HERE

Overall, amnesty had the effect of reducing militant attacks by roughly 3-7 events per month, a nearly complete reduction from the pre-amnesty mean of 4.9. The results appear robust to the choice of time trend and the inclusion of lagged terms or seasonal effects.

In robustness tests in Section E.1, we consider whether the results are robust to: *i*) many permutations of the main specification (Table 1), *ii*) the use of a control group/differences-in-discontinuities estimation (Figure A5), *iii*) placebo tests for all possible event-dates (Figure A6), *iv*) tests for outliers and influential observations (Figure A7), *v*) placebo-tests of non-oil-related conflict outcomes (Table B3), *vi*) structural break tests (Table B4 and Figure A8), and *vii*) AR(p) processes with optimal lags (Table B5). The estimated short-run event-study

[10] Note that since data is at the month level, they are collinear when included simultaneously with the functions of t, and so are included separately.

effect appears robust to all of these concerns.

4.2 Oil theft and amnesty

We estimate the differential evolution of oil theft in locations containing amnestied and non-amnestied rebels using the following difference-in-differences regression

$$y_{it} = \alpha + \psi T_i^d Post_t + \xi_t + \xi_i + \xi_i t + X_{it}' \beta + \varepsilon_{it} \quad (1)$$

where y_{it} is the level of oil theft in village i at time t, $T_{it}^d = T_i^d Post_t$ is an indicator of whether militant groups in the municipality were amnestied by period t, measured by an indicator for being d kilometers from an amnestied militant camp after 2009, where $d = 30$ is used for the bulk of the analysis.[11] The ξ_t and ξ_i are time and village fixed effects, and $\xi_i t$ allows for unit-specific linear time trend, which we include in some specifications. The vector of controls X_{it} contains initial conditions interacted with time dummies, including gridded population density in 2005 and distances to the nearest oil field and/or pipeline, state capital, Niger river and coast. The random disturbance ε_{it} is assumed, for identification, orthogonal to the indicator T_{it}^d, amounting to the standard parallel trends assumption required for DD to deliver consistent estimates of ψ.

Table 2 presents the main results of the DD estimation. In column (1), the baseline estimate of the standard TWFE model indicates that post-amnesty, sabotage events were differentially 0.069 events higher in amnestied villages. The estimates are significant at the 1% level. This is a very large effect, 1.38 times the sample mean, since sabotage is relatively rare in the sample. The estimates tell us that oil theft increased after the amnesty, and that this increase was indeed significantly concentrated in amnestied areas.

TABLE 2 HERE

In the DD analysis, we might worry about observable and unobservable time-varying, village-specific confounds. In particular, oil producing communities affected by militancy may be experiencing differential trends over the period in which the amnesty is implemented. Economic growth, trends in oil prices, or political crises affecting oil-producing communities may lead to changes in the level of oil theft in oil-rich villages for reasons unrelated to the amnesty. Several permutations to the main specification help rule out these alternate explanations. We include the controls X_{it} in column (5) and (6), and location specific time-trends $\xi_i t$ in columns (7) and (8), yielding broadly similar conclusions to the baseline model; the results remain significant at the 1% level. This suggests that – to the extent location-specific trends can be captured with a parametric linear specification – they are not

[11] We test robustness to different thresholds in Figure A21.

driving the results, nor are pre-treatment observables that might determine bunkering trajectories over time.

To relax parametric assumptions and rule out flexible differential trends in oil-intensive regions, we include the interaction between oil-production and time dummies in X_{it}, as measured by distance to oilfields and distance to oil pipelines. In columns (2) and (3) we explicitly include the coefficients for these oil-intensity covariates interacted with the post-amnesty indicator. While distance to oilfields is small and insignificant, distance to oil-pipelines – the primary sites of oil theft – is negative and significant at 1% in column (3). Villages closer to pipelines, with lower costs of oil theft, experience differentially larger growth in theft during the post-amnesty period. This suggests that part of the trend in Figure A1 is due to secular growth in theft which is naturally concentrated in places near oil pipelines. Still, the main coefficient remains positive and significant, indicating that the amnesty effect is present even holding constant this differential trend in pipeline communities.

Several potential confounders complicate the causal interpretation of ψ. In Section E.2, we consider the robustness of these main results in Table 2 to differential responses to in oil price trends (Table 2), control group definition (Figure A19 and Table B7), rebounding oil output (Figure A20), systematic measurement error in theft (Table B8), and geographic definition of the amnesty treatment (Figure A21). Figure A22 plots a comprehensive summary of 120 different robustness specifications. Beyond these specific tests, the parallel trend assumption might be violated for any number of reasons outside those explicitly controlled for in Table 2, perhaps due to unobservable differential trends between amnestied and non-amnestied territories. We evaluate the evidence of pre-trends in the standard event-study regression

$$y_{it} = \alpha + \sum_{\tau} \psi_\tau T_i^d D_\tau + \xi_t + \xi_i + X_{it}'\beta + \varepsilon_{it} \tag{2}$$

The specification is identical to the TWFE model (1) except that the time-varying treatment indicator T_{it}^d is replaced by T_i^d, an indicator for amnesty treatment within d kilometers, interacted with leads and lags of the treatment period, D_τ.[12] The ψ_τ are coefficients for each year τ before and after the amnesty, with $\tau = 0$ the omitted reference year. Figure 6, Panel A presents the results, plotting the ψ_τ by year. The pre-amnesty coefficients are very near zero and precisely estimated. The similar pre-amnesty trends indicate that, in contrast to violence outcomes, differential trends in bunkering are unlikely. The post-amnesty coefficients rise steadily and become significantly different first by 2013. Mirroring the raw time series in Figure A1, post-amnesty oil theft from amnestied groups grows at an increasing rate, also peaking in 2014.

FIGURE 6 HERE

[12] Not to be confused with the event-study or interrupted time series in Section 4.1.

A final potential concern is that the recorded sabotage oil spills reflect wanton vandalism by local communities rather than theft by criminal organizations. We show in E.3 that the effect of amnesty on pipeline sabotage is concentrated almost entirely on remote and fast-flowing "trunk" lines, large and economically valuable targets that carry crude oil from intermediate flow-stations out to export terminals. In contrast, the effects are minimal or non-existent on smaller lines and wellheads, suggesting an economic motive. The weight of evidence suggests it is highly unlikely that the observed increase in oil theft was an outcome independent of amnesty receipt.

4.3 Oil theft and security contracts

Providing certain rebel groups with lucrative security contracts may represent an attempt on the part of government to provide additional incentives to restrain post-amnesty rebel oil theft. If so, a natural question is whether these contracts effectively incentivized good behavior, notwithstanding the large rise in theft identified in Section 4.2. Figure A10 plots annual average village-level oil theft incidence for three groups of villages – those located within the territory of a rebel group that received amnesty and security contracts, those within the territory of a rebel group that received only amnesty, and those either located within the territory of a non-amnestied rebel group or outside of militant control altogether (the control group). The three groups evolve similarly pre-amnesty, while the post-amnesty spike is concentrated primarily in the amnesty-only group. The increase in oil theft is substantially smaller among those amnestied with security contracts, and flat in villages unaffected by amnesty. These results suggest that more generous amnesty terms provided effective incentives for rebels not to steal. Quantitatively, the average DD effect of amnesty of 0.069 in Table 2, column (1) rises to 0.10 among groups that received only amnesty and no security contract, significant at 1%. The effect of amnesty on oil theft in areas controlled by groups that received a security contract falls to 0.03. The effect among security contract recipients is therefore 73% smaller than the effect among amnesty-only villages, a difference significant at the 10% level.

5 Model

The analysis Section 4 above establishes that violence targeting the petroleum sector, the government's choice of incentive provision, and the increase in oil theft are likely interdependent, endogenous outcomes. We now develop a simple model of the Niger delta conflict where militant violence, government incentive provision and oil theft are jointly determined, capturing the idea that rebels fight to credibly signal their strength and extract concessions from the governing elites. In locations where militarily powerful rebels operate alongside active black markets this concession may take the form of tacitly accepting oil theft from the rebel group due to excessive cost of credible deterrence. These hypotheses are formalized in a dynamic game of no-commitment bargaining between the governing elite and a rebel group. The model's purpose is both to highlight the strategic mechanism sustaining violence and oil theft as equilibrium outcomes and to improve external validity by abstracting from the local setting. To this end our model is deliberately reduced-form and simple with respect to timing, preferences, and production functions. We abstract from the dynamic impact of recovering oil production, the organic growth of a black market and the strategic interaction between rebel groups.

We proceed as follows. The stage game and parameter restrictions are introduced in Section 5.1. We then derive the subgame perfect Nash equilibrium (SPE) and perfect Bayesian equilibria (PBE) of the dynamic game under full- and asymmetric information, respectively, in Sections 5.2 and 5.3. Section 5.4 concludes by presenting testable predictions on the relationship between military capability, local costs of oil theft, receipt of amnesty and oil theft incidence. These predictions are then taken to the data in Section 6.

5.1 The stage game

The extensive form of the stage game is shown in Figure 3 and proceeds in three steps. The government first makes an offer of x from a surplus of $\pi > 0$. The rebel either i) accepts the offer and abstains from oil theft, ii) accepts the offer and partakes in oil theft, or iii) rejects the offer outright and fights. If the rebel does not steal or the offer is rejected outright, payoffs are realized and the stage game ends. If the rebel steals, the government may either accept theft or cease payments, and in the latter case may also choose to fight. Both government and rebel have linear utility in their payoff. Conditional on an offer the game admits five feasible outcomes. We now introduce restrictions that impose an ordinal ranking over the five outcomes.

FIGURE 3 HERE

When the offer is accepted without theft the payoffs are $(\pi - x;\ x)$ for government and rebel. We allow x to be negative, in which case it should be interpreted as a bribe from rebel to government.[13] If the offer is rejected outright by the rebel there is fighting with payoffs $(\pi - \kappa;\ \alpha)$ where $\kappa \geq \alpha$ measures the government's cost of fighting, which we call the rebel battlefield damage, and $\alpha > 0$ the rebel's conflict payoff. If there is oil theft, let the market value of oil stolen by the rebel be μ and the rebel's income from oil theft be $\rho \leq \mu$. The government's net losses from rebel theft depend on the response of other actors in the black market for stolen oil. Recall from Section 2.2 that black markets for crude oil are in general characterized by dominant rebel groups operating alongside numerous smaller

[13] See reports discussed in Section 2.2, which discuss the rampant incidence of kickbacks and bribes to government officials in the black market for stolen oil.

competitors. We represent this market using the simplest possible dominant firm, competitive fringe structure. Let the government's losses from rebel theft be $(1-\gamma)\mu$, where γ measures rebel displacement of the fringe. We provide a simple microfoundation for γ in Appendix D.1 to capture the idea that fringe activity falls in local cost conditions when there is decreasing returns to scale in aggregate oil theft. We thus have that rebel oil theft affects government payoffs directly, through losses $\mu-\rho$, and indirectly, by displacing fringe theft to the value of $\gamma\mu$. Let

$$\underbrace{\mu-\rho}_{\text{lost surplus from rebel theft}} \quad - \quad \overbrace{\gamma\mu}^{\text{displaced fringe}} \quad = \quad \underbrace{(1-\gamma)\mu-\rho}_{\text{net loss from rebel theft}} \quad =: \ell$$

be the loss from rebel oil theft μ net of fringe displacement $\gamma\mu$ and rebel income ρ, to play a central role in determining equilibrium behavior. Consider the limiting cases to build intuition: If the competitive fringe is fully responsive $\gamma \to 1$ there is no efficiency loss from rebel oil theft and $\ell \to 0$. If the rebel is a local monopolist $\gamma \to 0$ the efficiency loss goes to its maximum value of $\mu-\rho$. We assume that $\mu \leq \kappa$, so the government's cost of fighting a rebel weakly exceeds losses from oil theft, and $\rho \geq \alpha$ rebel oil theft income weakly exceeds their payoff under conflict. When oil theft is allowed the pay-offs are $(\pi - x - (1-\gamma)\mu;\ x+\rho)$. When the payment is ceased the government recovers a share $1-\psi \in (0,1)$ of the offer x. Hence payoffs are $(\pi - \psi x - (1-\gamma)\mu;\ \psi x + \rho)$ when theft occurs and the transfer is ceased. We interpret ψ as representing the time spent detecting, attributing and responding to oil theft by halting amnesty payments. If $\psi \to 0$ the entire payment is recovered and rebels are left only with income from resource theft, and conversely if $\psi \to 1$ the rebels enjoy both the payment and oil theft income. Hence ψ governs the government's ability to provide rebels with monetary incentives. Finally, when payment is ceased and there is fighting the payoffs are $(\pi - \kappa - \psi x;\ \alpha + \psi x)$.

To summarize, the restrictions imply $0 < \alpha \leq \rho \leq \mu \leq \kappa$. Rebels enjoy a positive payoff α from fighting, which is inferior to their income ρ from oil theft. The market value of stolen oil extracted by the rebels is μ, worth at least the rebel income ρ, but plausibly μ substantially exceeds ρ.[14] The inequality $\mu \leq \kappa$ is consistent with the notion that rebels intentionally not only extract, but also deliberately destroy surplus when fighting.

[14] Field research suggests that the market value of lost and stolen oil may substantially exceed the income derived from theft, see Section 2.2 for details. Spills and leakage ensure that not all the oil taken from the pipeline reaches the market. The price rebels receive for both refined- and unrefined products is also lower than international benchmarks, with stolen crude sold at a steep discount in its unrefined state. Artisanal refining operations enjoy smaller returns to scale compared to sophisticated industrial refineries and yield a lower-quality product. Finally, the transportation and labor costs associated with marketing stolen crude are considerable.

We now formally define the stage game. Let

$$\mathcal{A}_R := \{\text{accept offer and do not steal, accept offer and steal, reject offer and fight}\}$$

be the action space of the rebel. Turning to the government, note that it is without loss of generality to consider only offers $x \in \mathcal{X} := [\alpha - \rho, \pi]$ as outright rejection and fighting is a strictly best response for the rebel to any offer $x < \alpha - \rho$ in the stage game, and the payoffs under outright rejection are independent of the offer. We define

$$\mathcal{A}_G := \mathcal{X} \times \{\text{cease payments and do not fight, allow theft, cease payments and fight}\}$$

as the action space of the government and $\mathcal{A} := \mathcal{A}_G \times \mathcal{A}_R$ the action space of the game. The space of rebel stage game strategies $\mathcal{R}$ is the set of all mappings from offers to actions $a_R : \mathcal{X} \to \mathcal{A}_R$. The government strategy is an offer $x \in [\alpha - \rho, \pi]$ and a function $a_G : \mathcal{A}_R \times [\alpha - \rho, \pi] \to \mathcal{A}_G$ that maps from offers and rebel actions to their action space, with the strategy space $\mathcal{G}$ the set of all such pairs $\{a_R, a_G\}$. We may define the stage game as a mapping from strategy profiles to individually rational payoffs $\Gamma : \mathcal{R} \times \mathcal{G} \to [\pi - \kappa, \pi] \times [\alpha, \pi - \kappa]$.

5.2 Subgame perfect equilibria

The dynamic game $\Gamma^\infty(\Delta)$ is played over an infinite, discrete time horizon, indexed over $t \in \{0, 1, \ldots\}$ and commonly discounted by $\delta := \exp(-\Delta\nu)$, where $\nu > 0$ is the discount rate, Δ is the time between offers and $\exp(\cdot)$ is the natural exponential function. The rebel survives a period of fighting with probability $\theta := \exp(-\Delta\sigma)$, where $\sigma > 0$ is the time rate of defeat. A rebel defeated in battle exits the game, yielding final payoffs $(0, \pi)$ for the rebel and government, respectively.[15] If there is fighting on the equilibrium path a rebel exits the game with positive probability. The game has a unique subgame perfect equilibrium with the Markov property of history independence, differing only in actions off the equilibrium path.[16] We now briefly describe how the equilibrium actions depends on parameter values and then formally derive conditions under which the three distinct behaviors arise.

Equilibrium without incentives. We first show that if efficiency losses from rebel oil theft

[15] Defining probabilities and discount rates in terms of the exponential function allows the time between offers to be tractably varied while holding the preferences and strategic environment constant.

[16] Let a history h_t be a list of action profiles denoting the actions taken from periods 0 to $t - 1$. The space of period-t histories $\mathcal{H}^t := \mathcal{A}^t$ is thus given by the t-fold product of the action space, where $\mathcal{H}^0 = \emptyset$. A time-$t$ strategy in the dynamic game is a mapping from history to a stage game strategies $a_G : \mathcal{H}^t \cup \mathcal{A}_R \to \mathcal{G}$ and $a_R : \mathcal{H}^t \cup \mathcal{X} \to \mathcal{R}$ for government and rebel. Let $\mathcal{R}^t$ and $\mathcal{G}^t$ denote the set of all such mappings; they are the time-t strategy spaces. We may define the dynamic game as a mapping from the strategy space to the set of individually rational payoffs $\Gamma^\infty(\Delta) : \mathcal{R}^t \times \mathcal{G}^t \to [\pi - g, \pi] \times [r, \pi - g]$. Strategies with the Markov property are independent of history $a_G : \mathcal{A}_R \to \mathcal{G}$ and $a_R : \mathcal{X} \to \mathcal{R}$. These strategies do not correspond to the stage game, as the set of feasible payoffs is expanded in the dynamic setting, with both rebel and government anticipating the payoffs from a possible rebel defeat when fighting.

are low $\ell \to 0$ the unique equilibrium features rebel oil theft and a transfer from rebel to government. We refer to this equilibrium as a minimal amnesty. The government optimally does not provide incentives when oil theft from the competitive fringe is responsive $\gamma \to 1$, if rebels are able to sell the stolen oil at a low discount $\rho \to \mu$, or if the cost of incentive provision is greater than the efficiency losses from oil theft. The intuition is that the government, faced in any case with considerable oil theft, prefers to deal with a dominant firm from which they can recapture lost surplus through bribes.

Low-powered equilibrium. Second we consider the case where efficiency losses from rebel oil theft are high $\ell \to \mu - \rho$ and the government can rapidly halt transfers if rebels steal, $\psi \to 0$. The rebel will receive a rent proportional to their income from oil theft and the government threatens a temporary suspension of payments should the rebels steal. We call this a low-powered equilibrium as the rebel is given weak incentives not to steal and is indifferent between participating in oil theft and not.

High-powered equilibrium. Finally we study a case where efficiency losses from rebel oil theft are again high $\ell \to \mu - \rho$ but the government is unable to rapidly halt transfers if rebels steal, $\psi \to 1$. To provide incentives not to steal the government must credibly threaten to fight rebels in the event of oil theft. Subgame perfection renders the rebel payoff proportional to their military capability net of efficiency losses. The intuition is that costly fighting entails a positive probability of rebel defeat, so the government's value of defeating the rebel is increasing in the equilibrium payment. Fighting a militarily strong rebel is rendered worthwhile by the anticipation of not having to pay exorbitant transfers in the future.

We now derive sufficient conditions under which these strategies constitute a unique equilibrium. Define the reservation values as the time-averaged payoff from fighting forever, which are

$$r := (1-\delta)\sum_{t=0}^{\infty}\delta^t\theta^t\alpha = \frac{1-\delta}{1-\delta\theta}\alpha \tag{3}$$

for the rebel and

$$\pi - g := (1-\delta)\left(\sum_{t=0}^{\infty}\delta^t\pi - \delta^t\theta^t\kappa\right) = \pi - \frac{1-\delta}{1-\delta\theta}\kappa \tag{4}$$

for the government respectively.

Proposition 1. ***Equilibrium without incentives.*** *Suppose $\ell \leq 0$, so that $\gamma \geq \gamma' := (\mu-\rho)/\mu$, and consider the following strategy profile: The government offers*

$$m := r - \rho < 0 \tag{5}$$

for any history, where r is the rebel reservation value (3) and allows theft. The rebel rejects and *fights for all $x \in (-\infty, m)$ accepts and steals for all $x \in [m, \rho/(1-\psi))$ and accepts without theft for all $x \in [\rho/(1-\psi), \infty)$. Then this strategy profile is the unique subgame perfect equilibrium.*

The proof is in Appendix D.2. Notice that m is negative by the assumption that $\alpha < \rho$, so the rebels are transferring surplus back to the government. An amnesty without incentives maximizes the total available surplus and transfers to the government all but rebel reservation value. The equilibrium is unique because an offer m is strictly preferred one-shot deviation by the government in any other strategy profile. We turn to an environment where there are efficiency losses from oil theft, $\ell > 0$ and recovery of amnesty payments is low $\psi \to 1$, so the government finds it optimal to create high-powered incentives. The following proposition shows that when recovery is high $\psi \to 0$ the equilibrium will feature low-powered incentive creation for rebels of sufficiently militarily power.

Proposition 2. ***Equilibrium with low-powered incentives.*** *Suppose $\ell > 0$, $\psi \to 0$ and $\kappa \geq \kappa' = \sigma/(\sigma+\nu)\rho + (1-\gamma)\mu$. Let*

$$c^{\ell} := \frac{\rho}{1-\psi} \tag{6}$$

and consider a strategy profile where the government offers $x = c^{\ell}$ for any history and ceases the transfer without fighting if the rebel steals. Rebels reject and fight for all $x \in (-\infty, r-\rho)$ accepts and steals for all $x \in [r-\rho, c^{\ell})$ and accepts without theft for all $x \in [c^{\ell}, \infty)$. Suppose rebel income from oil theft is exceeded by the government's average cost of fighting $\kappa(1-\delta)/(1-\delta\theta) \geq \rho$ and efficiency loss from oil theft $\rho < \ell$. Then this strategy profile is the unique subgame perfect equilibrium.

The proof is in Appendix D.3. The equilibrium features low-powered incentives when rebels are costly to fight relative oil theft losses and amnesty payments can readily be halted. If the battlefield damage satisfies $\kappa \leq \kappa'$ the government prefers to fight when ceasing amnesty payments off the equilibrium path, in which case incentive provision is high-powered. If rebels inflict sufficiently low battlefield damage relative to oil theft income $\kappa(1-\delta\theta)/(1-\delta) < \rho$ the government strictly prefers to fight rather than provide incentives. Notice that equilibrium punishments must be within the stage game. This is because there in the following stage always exists an offer supported by within-period punishments that the government prefers to punishing an additional period. The equilibrium is therefore uniquely determined by the lowest-cost incentive which deters oil theft and is consistent with subgame perfection. We now show that when recovery is low $\psi \to 1$ and the government's excess cost of fighting rebels over efficiency losses from rebel oil theft is sufficiently bounded the equilibrium will feature high-powered incentive creation.

Proposition 3. ***Equilibrium with high-powered incentives.*** *Let $\ell > 0$, $\psi \to 1$, suppose rebel battlefield damage is bounded by*

$$\kappa \leq \kappa'' := \frac{\delta(1-\theta)}{1-\delta\theta}\alpha + (1-\gamma)\mu + \frac{\delta(1-\theta)}{1-\delta}\ell \tag{7}$$

and consider a strategy profile where for any history the government offers

$$c^h := \max\left\{\frac{1-\delta}{1-\delta\theta-(1-\delta)\psi}\alpha, \frac{1-\delta}{\delta(1-\theta)}[\kappa-(1-\gamma)\mu]\right\} \tag{8}$$

and ceases amnesty payments and fights if the rebel steals. Rebels reject and fight for all $x \in (-\infty, r-\rho)$ accept and steals for all $x \in [r-\rho, c^h)$ and accept without theft for all $x \in [c^h, \infty)$. Then this strategy profile is the unique subgame perfect equilibrium.

The proof is in Appendix D.4. The government's threat of fighting is subgame perfect if the transfer is costly enough to render them indifferent between the losses from fighting and oil theft. The key intuition is that when the government fights, they may prevail militarily, the return to which is increasing in the contract value they otherwise would have to pay. The more costly a rebel is to fight, then, the more generous must the transfer be to sustain credibility of the government's off-equilibrium threat. Thus the surplus transferred is weakly increasing in the rebel's military capability. Importantly the cost of credible deterrence increases also in fringe activity γ, by a reduction in the efficiency loss of rebel oil theft. The implication is that

The proofs of Propositions 1-3 show that transitions between equilibria with minimal, low or high-powered incentive provision are governed by relative parameter values. The stated limiting conditions are sufficient but not necessary. To gain some intuition for the non-limiting cases consider Figure 4 which plots parameter combinations for which the equilibrium features minimal, low- and high-powered incentives, denoted in white, gray- and black. Panel (a) considers rebels of intermediate military strength κ, fixed and jointly vary recovery rates ψ (horizontal axis) and fringe activity γ (vertical axis). The minimal amnesty is optimal when efficiency losses from rebel theft are low relative to the cost of enforcement. Notice the limiting condition $\gamma \geq \gamma' = (\mu-\rho)/\mu$ from Proposition 1 on the vertical axis. If the government can provide responsive monetary incentives and efficiency losses are high, $\psi \to 0$ and $\gamma < \gamma'$ the government optimally provides low-powered incentives, granting the rebels a rent. Note that an active fringe γ increases the cost of a high-powered contract, rendering the minimal amnesty optimal also with high efficiency losses $\ell > 0$ by violating (7) for a given κ. In panel (b) we consider an environment where the fringe productivity γ is intermediate and jointly vary recovery rates ψ (horizontal axis) and battlefield damage κ (vertical axis). We see that the government always prefers providing high-powered incentives to militarily weak rebels, $\kappa \leq \kappa'$, see Proposition 2. The equilibrium features minimal amnesty rather than high-powered incentives when the cost of a credible contract is high relative to efficiency losses, exceeding the bound $\kappa \geq \kappa''$ in (7). The corresponding rebel payoffs are plotted in Figure A4, with lighter colors indicating greater payoffs and black rebel reservation values. Notice that the low-powered payoff varies only with the government's ability to provide monetary incentives, the horizontal dimension. On the vertical dimensions higher levels of fringe activity γ and unrecoverable payments ψ combine to induce the government to use high-powered incentives, increasing the level and variability of payoffs and tying them to rebel battlefield prowess.

FIGURE 4 HERE

5.3 Perfect Bayesian equilibria

We consider perfect Bayesian Nash equilibria of the game with asymmetric information. Rebels are either of the weak or strong type, fixed over time and respectively given by their survival probabilities $\underline{\theta}$ and $\overline{\theta}$ where we denote the type space of the rebels $\mathcal{T} = \{\underline{\theta}, \overline{\theta} : 0 < \underline{\theta} < \overline{\theta} < 1\}$. The government is initially uncertain of whether they face a weak or strong rebel, as encoded in their prior belief $p \in (0, 1) := \mathcal{P}$ that a rebel is weak. A time-t strategy in the dynamic game under asymmetric information is thus a mapping from the type, prior belief and history to stage game strategies, $a_G : \mathcal{H}^t \times \mathcal{P} \times \mathcal{A}_R \rightarrow \mathcal{G}$ for the government and $r : \mathcal{H}^t \cup \mathcal{X} \times \mathcal{T} \times \mathcal{P} \rightarrow \mathcal{R}$ for the government and rebels. Let rebel payoffs for the game of complete information be $\underline{x}$ for the weak and $\overline{x}$ for the strong, referred to interchangeably as limited and generous amnesty. Recall the rebel reservation values defined in (3) and let $\underline{r} = r(\underline{\theta})$ and $\overline{r} = r(\overline{\theta})$ be the reservation value of the weak and strong.

We now introduce parameter restrictions on the rebel payoff under conflict rents α and battlefield damage κ which are sufficient to ensure $\underline{x} \leq \overline{x}$ and the existence of prior beliefs supporting equilibrium path conflict. To begin, consider SPE without oil theft and the belief p that renders the government indifferent between the strategy of offering the strong rebel's SPE payoff $\overline{x}$ or fighting one period and then offering $\overline{x}$,

$$p^* := p : \underbrace{\pi - \overline{x}}_{\text{cost of generous amnesty}} = (1-\delta) \underbrace{(\pi - \kappa)}_{\text{present cost of fighting}} \dots$$

$$\dots + \delta \underbrace{\left[p(\underline{\theta}(\pi - \overline{x}) + (1-\underline{\theta})\pi) + (1-p)(\overline{\theta}(\pi - \overline{x})) + (1-\overline{\theta})\pi\right]}_{\text{belief-contingent continuation value from fighting}}$$

implying that the government prefers to offer $\overline{x}$ immediately for all beliefs

$$p \leq p^*(\overline{x}) := \frac{(1-\delta)\kappa - (1-\delta\overline{\theta})\overline{x}}{\delta(\overline{\theta} - \underline{\theta})\overline{x}} \tag{9}$$

where the critical belief p^* balances the expected cost of fighting with that of a generous transfer. Notice that the critical belief is a function of the government's payoff when facing a strong type in the full information setting, given by $\pi - \overline{x}$ under low- and high-powered incentives. For minimal amnesty with oil theft the payoff is $\pi - \overline{r} + \rho - (1-\gamma)\mu$ and critical belief is given by $p^*(\overline{r} + \ell)$ as $\overline{r} - \rho + (1-\gamma)\mu = \overline{r} + \ell$. We have that if the expected cost of fighting a weak rebel is less than granting the strong rebels payoff

$$\frac{1-\delta}{1-\delta\underline{\theta}}\kappa \leq \overline{x} \implies p^*(\overline{x}) \leq 1 \tag{10}$$

there exist beliefs $p \in (0,1)$ such that the government prefers to fight at least one period over offering $\overline{x}$. To ensure this, recall that the strong rebel's payoff $\overline{x}$ is bounded below by their reservation value, $\overline{x} \geq \alpha(1-\delta)/(1-\delta\overline{\theta})$ so bounding rebel battlefield damage κ to lie in the interval

Assumption 1. ***Bounded rebel military capability.***

$$\kappa \in \mathcal{K} := \left[\mu, \alpha\frac{1-\delta\underline{\theta}}{1-\delta\overline{\theta}}\right] \implies p^*(\overline{r}) \leq 1$$

is sufficient to ensure, for any period length Δ and subgame perfect equilibrium, that the government weakly prefers fighting any weak rebel group to granting them $\overline{x}$. Given the existing bounds on parameters this interval exists only when $\delta\overline{\theta} \to 1$. The assumption may be relaxed for subgame perfect equilibria in which strong type enjoys payoff in excess of their reservation value. We make the assumption here to allow a joint treatment of all full-information payoffs. Let $\max(\mathcal{K}) = \overline{\kappa}$ and recall from Proposition 3 that a high-powered equilibrium existed if κ was inferior to κ''. Appendix D.5 shows that $\overline{\kappa} \leq \kappa''$ for all types if

Assumption 2. ***Bounded rebel conflict payoffs.***

$$\alpha \leq \frac{1-\delta}{1-\delta\underline{\theta}}\rho$$

and when there is efficiency loss from oil theft $\gamma : \ell > 0$. Because $\overline{\kappa}$ is increasing in α, this restriction indirectly constrains the maximum excess cost of a high-powered transfer.

To summarize, a rebel is parameterized by their battlefield damage $\kappa \in \mathcal{K}$ and type $\theta \in \{\underline{\theta}, \overline{\theta}\}$. We have that the SPE payoff of the strong rebel weakly dominates that of the weak $\underline{x} \leq \overline{x}$ for minimal, low- and high-powered amnesties. Hence we refer to $\underline{x}$ and $\overline{x}$ as limited and generous amnesty, respectively. We now formally prove the existence of a perfect Bayesian pooling equilibrium with fighting when the government offers minimal amnesty.

Proposition 4. ***Perfect Bayesian equilibrium with fighting and minimal amnesty****. Suppose $\ell \leq 0$ such that payoffs of the full information are $\underline{x} = \underline{r}$, $\overline{x} = \overline{r}$ for the strong and weak type.*

Consider a strategy profile where rebels of both types reject and fight for all $x \in (-\infty, \bar{r} - \rho)$ accepts and steal for all $x \in [\bar{r} - \rho, \rho/(1-\psi))$ and accept without theft for all $x \in [\rho/(1-\psi), \infty)$. The government offers $\underline{r} - \rho$ for the first $\omega \in \mathbb{Z}_+$ periods and then $\bar{r} - \rho$ forever. Have the government believe with certainty that a rebel accepting $x < \bar{r} - \rho$ is a weak type. Then for all $p > p^*(\bar{r} + \ell)$ there exists a $1 \leq \omega < \infty$ such that this strategy is an essentially unique perfect Bayesian equilibrium when $\Delta \to 0$.

The proof is in Appendix D.6. The intuition is as follows. The government is initially optimistic, $p > p^*(\bar{r} + \ell)$ in (9) and strictly prefers to fight rather than offer the SPE payoff of the strong type. But weak rebels understand that every period they fight and survive the government will revise upward their belief that they face a strong rebel. Weak rebels therefore anticipate that they need to fight only for a finite number of periods before the government grants them a strong-type payoff. If faced with a sufficiently long anticipated conflict there exists a separating offer in the first period which is accepted by the weak and rejected by the strong. In this case uncertainty is resolved after at most one period of fighting. However, if the anticipated conflict is short a separating offer may not exist and both types will strictly prefer to reject offers less than $\bar{x}$.[17]

We demonstrate that there exists an essentially PBE with fighting when incentive provision is low-powered under full information (Proposition 2) in Appendix D.7. The low-powered PBE differs from the minimal amnesty equilibrium above in only two respects. First, strong rebels now may receive a payoff in excess of their reservation value, increasing the scope for separating offers, so the existence of a fighting equilibrium demands lower values for Δ, everything else equal. Second, payoffs do not necessarily vary with type, in which case there is no incentive to fight.

We now turn to high-powered equilibria. Recall from Proposition 3 that the transfer c^h in (8) rendered the government indifferent between fighting and just halting transfers in the event of rebel oil theft, so the contract value was proportional to rebel military capability. But under asymmetric information the cost of fighting is endogenous with respect to beliefs p over rebel type. Given a belief p the transfer x that equates the government's payoff from just halting payments with fighting is

$$\underbrace{(\pi - (1-\gamma)\mu - \psi x)}_{\text{payoff, cease payments}} + \underbrace{\delta(\pi - x)}_{\text{continuation value no fighting}} = (1-\delta)\underbrace{(\pi - \kappa - \psi x)}_{\text{cease payments, fight}} + \dots$$

$$\dots \delta \underbrace{p(\underline{\theta}(\pi - x) + \delta(1-\underline{\theta})\pi) + (1-p)(\bar{\theta}(\pi - x) + \delta(1-\bar{\theta})\pi)}_{\text{belief-contingent continuation value with fighting}}$$

[17] Notice that accepting an inferior offer is never a best response for any rebel type and therefore does not occur on the equilibrium path. Because Bayes' rule is not defined for zero-probability events we specify a government belief such that the continuation payoff is well-defined. The equilibrium is not sensitive to this choice, however.

where solving with respect to x yields the efficient belief-contingent pooling offer

$$x(p) = \max\left\{\frac{(1-\delta)}{1-\delta\overline{\theta}-(1-\delta)\psi}\alpha, \frac{(1-\delta)(\kappa-(1-\gamma)\mu)}{\delta[(1-\overline{\theta})+p(\overline{\theta}-\underline{\theta})]}\right\} \tag{11}$$

bounded below by the reservation value of the strong and strictly decreasing in $p \in [0, 1]$. It follows that

$$\hat{p} := p : p = \frac{(1-\delta)\kappa - (1-\delta\overline{\theta})x(p)}{x(p)\delta(\overline{\theta}-\underline{\theta})} \tag{12}$$

the belief at which the government is indifferent between the making pooling offer (11) forever or fighting one period and then making the pooling offer forever. Note that the solution to (12) implies a unique $\hat{p} \in (0, 1)$ because $x(p) \geq \overline{r}$, the left-hand-side p grows linearly from 0 and the right-hand-side falls convexly from a strictly positive value, crossing once. We can now formally state the existence of an essentially unique PBE when incentive provision is high-powered:

Proposition 5. ***Perfect Bayesian equilibrium with fighting and high-powered incentives****. Let $\ell > 0$ and p_t be the government's posterior belief of rebel type after t periods of fighting. Consider the following strategy profile: The government offers $\underline{r} - \rho$ for the first $\omega \in \mathbb{Z}_+$ periods and then $x(p_\omega)$ for all $t \geq \omega$. Rebels of both reject and fight for all $x \in (-\infty, \overline{r} - \rho)$ accept and steal for all $x \in [\overline{r} - \rho, x(p_t))$ and accept without theft for all $x \in [x(p_t), \infty)$. Have the government believe with certainty that a rebel accepting $x < \overline{r} - \rho$ is a weak type. For all $p > \hat{p}$ there exists a $1 \leq \omega < \infty$ such that this strategy profile is an essentially unique perfect Bayesian equilibrium when $\psi \to 1$ and $\Delta \to 0$.*

The proof is in Appendix D.8. The intuition follows that of Proposition 4 above.

In the perfect Bayesian equilibria discussed so far weak rebels either receive generous offers or are defeated. That no rebel ever accepts a separating offer and a less generous payoff is arguably inconsistent with the heterogeneity in observed amnesty outcomes discussed in Section 2. However, equilibria in which rebels accept limited amnesty arise naturally under two-sided uncertainty. In Appendix D.9 we construct the simplest possible PBE in which weak rebels first fight and then accept a separating offer, revealing their type. The government's prior p is drawn from $\{\underline{p}, \overline{p}\}$ with $p^* < \underline{p} < \overline{p} < 1$ low or high. Because the equilibrium-path conflict duration is increasing in p the government prior beliefs have a natural interpretation as resolve, with $\overline{p}$ a determined type. The intuition is that weak rebels may accept limited amnesty when they are surprised to learn, on the equilibrium path, that the government is determined and that conflict will last longer than first anticipated.

5.4 Predictions

Our theoretical framework captures the idea that rebels fight to credibly signal their strength and extract concessions from the governing elites. In locations where militarily powerful rebels operate alongside active black markets this concession may take the form of tacitly accepting oil theft from the rebel group. We now states these hypotheses in terms of predictions on the sign and functional form of the empirical relationship between military capability, local cost conditions in the black market, amnesty receipt, and aggregate oil theft.

Prediction 1. ***Military capability and amnesty.*** *The probability of receiving amnesty is increasing in military capability κ.*

The proof is in Appendix D.10. We say that amnesty is more probable the less periods of fighting are necessary before the government grants a generous payoff (denoted $\overline{x}$ in 5.3 above). It is optimal for the government to fight so long as the dividends from rebel defeat less the expected cost of fighting exceed the generous amnesty terms the rebels demand. But the expected cost of fighting is a function of rebel survival probability and battlefield damage, so for given amnesty terms the government optimally offers generous terms earlier to militarily powerful rebels. But recall that when the government provides high-powered incentives (Proposition 5) the value of a generous amnesty is itself increasing in rebel battlefield damage, introducing a countervailing incentive for the government to keep fighting. We show that the net effect of an increased in rebel battlefield damage is nevertheless to decrease conflict duration, because under asymmetric information the generous amnesty (11) is offered at a nondegenerate posterior belief. The government therefore discounts the offer relative to the strong type's full-information subgame perfect equilibrium offer (see Proposition 3). As such, Prediction 1 is not contingent on the form of post-amnesty incentives.

Prediction 2. ***Cost conditions and oil theft.*** *Post-amnesty oil theft decreases in the local cost of oil theft.*

This prediction follows from the dominant firm, competitive fringe structure of the black market. An increase in local theft costs reduces competitive entry directly and indirectly by increasing the government's value of providing incentives to the dominant rebel. The direct effect follows from price-taking behavior, see the microfoundation for fringe parameter γ in Appendix D.1. When dominant rebels abstain from oil theft the increase in oil theft from the competitive fringe is greater when the local cost environment is favorable. Everything else equal, a more responsive fringe reduces the government's efficiency loss ℓ from the rebel oil theft, lowering the government's return to providing costly incentives to that rebel, and ultimately rendering it optimal to tacitly allow oil theft. Thus both direct- and indirect effects imply that local oil theft is weakly decreasing in the local cost of oil theft.

Combining Predictions 1 and 2, we have that:

Corollary 1. ***Cost conditions, amnesty and incentive contracts.*** *As the black market becomes increasingly active ($\gamma \to 1$) variation in rebel battlefield damage should predict amnesty offers but not receipt of incentive contracts.*

The following prediction states that the indirect marginal effect of local costs γ on oil theft levels is increasing in rebel military capability κ when mediated by the government's incentive provision decision.

Prediction 3. ***The interaction of military capability and cost conditions on oil theft.*** *The marginal effect of local black market costs on oil theft is greater in locations with militarily powerful rebels.*

Predictions 1 and 2 state that unconditionally more militarily powerful rebels receive better amnesty terms and locations with high costs of oil theft feature greater oil theft. Recall from Proposition 3 that if the government's ability to cease payments is low ($\psi \to 1$) they must credibly threaten to fight the rebel in order to deter oil theft. The cost of an incentive contract is then increasing in rebel military capability and decreasing in cost conditions. Conditional on receiving amnesty there exists a level of military strength κ'' above which the credibility cost of providing high-powered incentive contracts is no longer worth bearing, see the vertical dimension of Figure 4, panel (b). However, notice that, all else equal, an increase in competitive fringe activity γ directly increases the cost of a high-powered contract for militarily powerful rebels, see Figure A4, panel (a).[18] For the government to credibly threaten fighting in response to oil theft the high-powered transfer increases proportionally with the reduction in net losses of rebel oil theft: The greater is fringe activity γ, the lower must rebel battlefield damage κ be for the government to optimally tacitly allow oil theft.[19] Military capability and black market fringe activity must therefore both exceed critical levels before triggering a shift to minimal amnesty. As such, the marginal effect of γ on oil theft depends on κ through a discrete shift in incentive provision. Thus we expect post-amnesty oil theft to concentrate among militarily strong rebels in locations with favorable black market conditions.

The discrete change in government incentive provision implies our final prediction:

Corollary 2. ***Non-linear effects.*** *Aggregate oil theft respectively decreases- and increases non-linearly in local cost conditions and military capability.*

A favorable environment for the black market has two effects on oil theft: a direct effect, increasing activity among the competitive fringe, and indirectly, by inducing the government to tacitly allow oil theft from rebel groups. The indirect effect is given by a discrete shift in

[18] To verify, consider the identity (11) defining the high-powered transfer c^h in the perfect Bayesian equilibrium. We then have that $\partial c^h / \partial \gamma > 0$ when c^h exceeds $\alpha(1-\delta)/(1-\delta\overline{\theta}-(1-\delta)\psi)$.

[19] This can be seen in equation (7), where $\frac{\partial \kappa''}{\partial \gamma} = -\mu \frac{1-\delta\theta}{1-\delta} < 0$.

incentive provision and is strengthened under high-powered incentives, in which increased fringe activity both increases the cost of, and decreases the return to, incentive contracts to militarily powerful rebel groups. Fixing the government's ability to recover amnesty funds ψ there exists a pair of critical values for military capability κ and γ which trigger a shift from incentive provision to minimal amnesty. As such, we expect non-linearity in the effects of these parameters on post-amnesty oil theft.

6 Testing equilibrium mechanisms

We test the implications of our bargaining model presented in Section 5.4. We begin in Section 6.1 by validating a measure of rebel military strength based on local alliance density and then relate heterogeneity in amnesty outcomes to rebel military strength in a cross-section of militant camps. We find that militarily powerful rebels are more likely to receive amnesty, be offered valuable surveillance contracts, and suffer battlefield defeat at lower rates. We also find evidence that the correlation between rebel strength and amnesty outcomes depends on the local black market cost structure. Finally, we provide empirical support for the assumption that rebels retained their military capability post-amnesty. These results are consistent with the core conflict-bargaining mechanism (Prediction 1) and differentiated incentive provision (Corollary 1).

We then move to our village-level panel data, where we employ triple- and quadruple-difference specifications to show how the effect of amnesty on oil theft varies by local cost conditions and rebel military strength. In Section 6.2 we first show that post-amnesty oil theft grows more in areas with low black market costs (Prediction 2), as measured by median wages of young men in the local labor market. Consistent with our overriding hypothesis that post-conflict oil theft is driven in part by government's decision to selectively provide incentive contracts (see Corollary 1), we show in Section 6.3 that the growth in oil theft is increasing in local density of allied camps and therefore concentrated among the stronger rebel groups for whom high-powered incentives are too costly. In Section 6.3, we also find that that the role of black market entry costs is greater for stronger rebel groups, consistent with Prediction 3. Finally, in Section 6.4 we find that the marginal effects of local costs and military strength are highly nonlinear, a pattern suggestive of switching from incentive provision to minimal amnesty equilibria (Corollary 2).

6.1 Rebel military capability and amnesty

The model predicts that militarily stronger rebels are more likely to be offered amnesty for any post-amnesty subgame perfect equilibrium. In this section, we validate our measure of militant strength, and test whether it is correlated with amnesty receipt (Prediction 1) and differentiated incentive contracts (Corollary 1). Using the sample of 69 militant camps, we estimate cross-sectional regressions of the form

$$y_i = \alpha + \varphi k_i^{\delta} + X_i'\beta + \xi_s + v_i \tag{13}$$

where y_i is an outcome, k_i^{δ} is a measure of camp strength based on alliance density, X_i are camp-level controls, and ξ_s is a state fixed effect capturing geographic heterogeneity.

We now validate that alliance density measure is meaningfully capturing the capability of a militant camp to attack and inflict damage on the oil sector. As a first step we rerun the RDiT analysis (Section 4.1) but split the sample of amnestied villages by those in the orbit of low and high military strength rebels, as measured by the numbers of allies (described in Section 3). We then re-estimate the village-level monthly RDiT separately for each subsample. The plots indicate a large and significant reduction in conflict the high strength sample, and a much smaller (though still significant) effect in the low strength sample. To confirm that higher frequency of attacks reflect greater damage we take as outcome variable the percent change in aggregate oil production occurring in all oilfields within 20 km of camp *i* over the period 2005-2009, the height of the Niger Delta conflict. Figure A13 Panel A demonstrates a negative correlation between alliance density and output loss, conditional on controls and state fixed effects. Estimates in Table 3 confirm that this relationship is significant at 5 or 1% level and ranges from 4.7-7.9 percentage points of output loss for every additional local ally. Furthermore, local alliance density accounts for nearly 20% of the variation in oil output change. Thus we establish that more allied groups destroyed more oil output, allowing us to use alliance density to capture heterogeneous rebel fighting capabilities. This correlation is robust to numerous sensitivity checks, including different periods over which conflict is defined, distances over which alliance density is defined, and model specification. See Figure A14 for a complete summary of these tests.

TABLE 3 HERE

Next, we look at whether amnesty outcomes also vary by militant strength, establishing three key results: stronger rebel groups are less likely to be defeated in battle (Figure A15), more likely to receive any amnesty (Figure A16), and more likely to receive incentive contracts (Figure A17). In addition, camps with high density of non-allies are also significantly less likely to receive any amnesty. Estimates in Table 3 verify that the relationships are statistically significant and robust to fixed effects and controls which include a dummy for membership in MEND, the largest umbrella militant group.

To test whether the relationship between amnesty outcomes and militant strength varies by local const conditions, as predicted in Corollary 1, we estimate a nonparametric kernel regression of amnesty outcomes on the interaction between cost conditions, partialling out camp-level controls and state fixed effects. To measure costs, we use the median wages for young men (under age 40) in the local labor market as a proxy for labor outlays. We plot the nonlinear marginal effect of military strength on amnesty and surveillance contracts at different levels of the wage distribution in Figure A18. For very low wage levels ($\gamma \to 1$), we find that there is no relationship between security contracts and military strength. In these markets, the fringe is so responsive that offering high-powered incentive contracts is never optimal, so payoffs are constant in κ. Only when γ is sufficiently low (that is, wages are sufficiently high) to allow for high-powered incentives does amnesty payoff generosity significantly increase in κ. In contrast, we find that amnesty receipt is increasing in rebel military strength throughout the wage distribution, since the Bayesian updating process that leads government to offer amnesty always depends on κ for any equilibrium payoff.

Lastly, we provide empirical support for the assumption that rebel military capability κ does not change. Recall that for the bargain to be self-enforcing, rebels must be able to re-mobilize militarily. This assumption is plausibly violated if the amnesty was successful at demobilization and reintegration of ex-combatants into civilian life. In Appendix E.4, we argue that the 2015 election of president Muhammadu Buhari, who held a hawkish stance toward the amnesty, constituted a shock to the equilibrium bargain. Exploiting another RDiT, we show that Buhari's announcement that amnesty payments would be slashed by 70% led to an immediate jump in oil sector militancy. This response is remarkably similar in magnitude to the reduction in conflict achieved by amnesty itself.

In general, the results indicate that government differentiates amnesty offers across groups, targeting stronger rebels with better deals, while weaker groups received limited amnesty, nothing at all, or were defeated in battle. This is consistent with the key mechanism in our theoretical model, where battle strength is revealed through conflict over time and retained post-conflict, inducing the government to separate militant types with targeted peace offers.

6.2 Local cost conditions and oil theft

Post-amnesty oil theft will be greatest in low-cost (high γ) markets with a responsive fringe via two mechanisms.[20] First, for a given post-amnesty subgame perfect equilibrium, a responsive fringe steals a greater share of residual output, all else equal. Second, γ determines the efficiency losses from rebel theft, and therefore the optimality of restraining rebels. Since the value government of contracts that dissuade rebel oil theft are falling in γ, a decrease in bunkering costs may trigger a shift from low or high-powered contracts to minimal amnesty, accompanied by a jump in oil theft.

[20] Recall that the fringe responsiveness parameter γ is defined as the price-cost ratio for fringe entrepreneurs in the black market (see Appendix D.1).

To measure entry costs, we use three variables: *i*) distance from village to the Atlantic coast, *ii*) distance to the Niger River, and *iii*) median wages for young men (under age 40) in the local labor market.[21] Distance to rivers and the coast provides a measure of the transport costs faced by oil theft gangs to move the oil and access export markets, respectively. Local wages determine the opportunity costs faced by young men on the margin of joining organized crime, and thus the prevailing black market wage.

To test for heterogeneity, we estimate a triple-differences model that interacts the treatment indicator T_{it}^{d} with the above-mentioned cost factors.[22] Results are given in Table 4. Column (1) reprints the main results for reference, column (2) interacts with distance to coast, column (3) with distance to the Niger River, and column (4) with median wages. Columns (5)-(6) include all cost factors interacted simultaneously.

TABLE 4 HERE

The results suggest that there is robust heterogeneity in post-amnesty bunkering growth by cost factors. In each case, the amnesty effect is largest for low-cost villages, and falls to zero as costs increase. The estimates for the interactions with distance to coast (export costs) and local wages (labor costs) are negative and significant in every specification. In contrast, distance to the Niger River (transport costs) is only negative and significant when all cost variables and controls are included in column (6). Post-amnesty bunkering is concentrated primarily in low-cost markets that are attractive targets for fringe entry. In Figure A24, we split the sample by whether a village is above or below the median of the market-level wage distribution and estimate dynamic event-study coefficients for each subsample. While parallel trends obtain in both subsamples, post-amnesty oil theft effects are only observed in the low-wage subsample.

The triple-difference interaction coefficient on young mens' median wages is robustly negative in Table 4. But wages may simply be proxying for overall economic conditions, and effects may be greater in poor communities for a variety of reasons unrelated to the fringe entry mechanism. In Table B11, we use wages of other demographic groups as a falsification test. For each triple-difference specification, we include interactions with median wages for young men, old men, or young women, with and without controls. As long as wages of these other groups are also correlated with overall development, they should produce similar effects if the result is spurious. However, the results for young mens' wages are negative and significant, while those for old men and young women are insignificant when controls are included, and much smaller than the young male estimates, as visualized in

[21] Labor markets are defined at the level of the local government area to minimize spillovers, and wages are measured in 2009. These cost factors are derived from the qualitative literature, which suggests that labor costs dominate for most small-scale bunkering operations, followed by transport and bribery costs.

[22] The cost factors are also interacted with year fixed effects.

Figure A25. Only wages in our demographic group of interest display evidence of significant heterogeneity, ruling out generalized income effects.

6.3 Local cost conditions, military strength, and oil theft

In the high-powered equilibrium, contract values are increasing in κ. At the margin, an increase in military strength may lead to a violation of the government's participation constraint, triggering a shift to a minimal amnesty equilibrium in which theft is implicitly tolerated and generating a positive correlation between military strength and theft. Furthermore, a lower γ shifts up the critical κ below which high-powered contracts are viable. As such, the range in which theft is increasing in κ may shrink, or disappear altogether. Therefore, as black market costs rise and γ falls, the response of oil theft to κ should fall. We test these predictions in the following quadruple-difference specification:

$$\begin{aligned} y_{it} =& \alpha_0 + \alpha_1 Post_t \times T_i^d + \alpha_2 Post_t \times T_i^d \times k_i^\delta + \alpha_3 Post_t \times T_i^d \times w_i \\ &+ \alpha_4 Post_t \times T_i^d \times k_i^\delta \times w_i + \xi_t(k_i^\delta + w_i + k_i^\delta w_i) + \xi_i + X_{it}'\beta + u_{it} \end{aligned} \tag{14}$$

Where k_i^δ measures rebel military strength as the number of local allies along a pipeline within δ kilometers for village i's nearest militant camp and w_i measures costs as young male wages. All specifications use $\delta = 10$ kilometers. Here, the object of interest is the quadruple-difference coefficient α_4, which tests the null hypothesis that the heterogeneous effect of local alliance density k_i^δ does not vary with costs w_i. The triple-interaction terms α_2 and α_3 test heterogeneity of the main effect with respect to k_i^δ and w_i, respectively, when the other variable is fixed at zero. Given the theoretical prediction, we also expect $\alpha_2 > 0$, since greater military strength increases the likelihood of a shift from high-powered contracts to minimal amnesty, particularly so when costs are low and the fringe is responsive. Finally, α_1 gives the main treatment effect for low-cost, militarily weak areas. 2- and 3-way interactions with the year fixed effects complete the quadruple-difference specification.

The results are given in Table 5. We begin by re-estimating the main results in column (1) on the sample of villages for which control variables and wage data are available. In columns (2)-(5) we consider the triple-difference model that interacts the time-varying amnesty treatment variable with our measure of militant strength. Columns (6)-(7) complete the quadruple-difference model by adding the interaction with young male wages in the local labor market. The results of the triple-difference-only model in columns (2)-(5) indicate that post-amnesty oil theft increases in military strength: an additional local ally increases the effect of amnesty on oil theft by 0.02-0.07 incidents annually. These estimates are significant at 5% or lower in all but column (2). They are also robust to including interacted controls (column 3), accounting for membership in MEND (column 4) and controlling for overall militant camp density (column 5).

Figure A26 illustrates parallel trends using a split-sample event-study: in the both low- and high-κ samples, pre-trends are flat, and large post-amnesty effects only emerge in the sample of villages controlled by militants with at least one local ally. In Figure A27, we plot the triple-difference coefficients for varying levels of δ from 10 to 100 km. These effects effects fall to zero as distance along the pipeline grows, implying that fighting capacity is a fundamentally local property dependent on a camp's position in the alliance and physical infrastructure networks.

TABLE 5 HERE

Columns (6) and (7) show that while low-wage areas experience oil theft that increases in the number of local allies, this effect reverses as wages rise. The quadruple-difference effects are statistically significant at 5% in column (6) and 10% in column (7), though similar in size under both specifications. In magnitudes, the heterogeneous effect of an additional ally is roughly 3 times greater when labor costs are zero than in the triple-difference results of column (2). The partial derivative $\frac{\partial y_{it}}{\partial k_i^{\delta}}$ then falls to zero when $w_i = 0.22$ (220 Naira/hr), corresponding to roughly the 95^{th} percentile of median wages across villages. For the highest-cost bunkering markets, then, the effect of military strength disappears entirely.[23]

6.4 Non-linearity

Heterogeneity in κ and γ induce equilibrium shifts by increasing the costs and reducing the benefits of no-bunkering contracts. As such, we should expect non-linear dynamics in the interaction effects tested in Sections and , as these equilibrium shifts induce discrete jumps in post-amnesty oil theft.

We test for nonlinearity in local wages by non-parametrically estimating the marginal difference-in-difference effect of amnesty across the market-level wage distribution, after partialing out year and village fixed effects. Figure 7 plots the marginal effects, estimated by kernel regression, as well as a histogram showing the density of the data across the wage distribution. The results reveal that the linear heterogeneity estimates in Table 4 are driven almost exclusively by very low-wage (high γ) markets. Quantitatively, the effect of amnesty on oil theft is 5 times larger at the 5th percentile of the wage distribution than at the median, falling to nearly zero shortly thereafter. The results are consistent with a critical point at which fringe activity falls enough that it becomes optimal for government to restrain oil theft, indicating shift from minimal amnesty to either a high or low-powered equilibrium.

We similarly test for nonlinearity in local alliance density in Figure 8. However, since the

[23] In additional tests, available upon request, we do not find similarly significant quadruple-difference interaction coefficients for our export cost measures, which is consistent with the results of Table 4, and with qualitative reporting that labor costs dominate the cost structure of illegal firms.

local allies variable takes only discrete values from 0 to 7 and is therefore essentially categorical, we use a semi-parametric binned estimation instead of a full nonparametric regression. In particular, we estimate the main differences-in-differences regression on subsamples with levels degrees of local alliance density.[24] The results again reveal nonlinear dynamics that are suggestive of an equilibrium shift. The effect of amnesty on oil theft is a small positive or indistinguishable from zero for all of the subsamples with 5 or fewer allies. However, for the 6-7 allies subsample, the result is 4.7 times larger than the full-sample effect in Table 2 column (1).

Note that nonlinearity in wages only suggests an equilibrium shift from restraining to minimal amnesty, but cannot distinguish between high and low-powered initial equilibria. However, nonlinearity in military strength only occurs in the shift from high-powered to minimal amnesty, since low-powered contracts are independent of κ. This pattern therefore, together with the results of Section 6.3, points to high-powered and minimal amnesty as the primary equilibria in the data.

7 Conclusion

We study how the ruling elite and rebels bargain over resource rents in the context of the Niger Delta insurgency, a decades-long conflict between militant groups and the Nigerian state that ended in amnesty for combatants. We develop a no-commitment dynamic model of bargaining over resource rents in the presence of a black market.

Leveraging rich data on the location, alliances, black market activity, amnesty participation, and violent attacks of militant commanders and the villages under their control, we test the theoretical predictions of the model. We show that amnesty was effective at reducing violence, but also led to increased oil theft in amnestied territories, driven by an equilibrium switching effect where the ruling elites optimally allowed oil theft from militarily powerful rebels in locations with active black markets.

Our analysis provides empirical support for the conflict bargaining framework and the amnesty policy itself. We find that payments to Niger Delta rebels enforced by threats of fighting are effective deterrents of bad behavior. Amnesty restored order to the Niger Delta, while pipeline surveillance contracts to militant leaders reduced oil theft. Despite this, inefficient outcomes can persist in a second-best world where government may not have the incentives to fully eliminate illegal activity. Thus, our analysis illustrates the importance of treating conflict and illegal activity jointly when studying the political economy of resource-rich and institutionally weak states.

An important implication of our work is that socially efficient outcomes are more likely

[24] We divide the data into 0-1, 2-3, 4-5, and 6-7 allies

where government more fully internalizes the costs of theft. Greater state ownership of oil revenue, reduced agency frictions, and internalization by the state of the local costs of criminality – e.g. environmental externalities – all increase government's willingness to make deals. A countervailing force, however, is the possibility of dynamic entry of new combatants in response to amnesty, which, like entry in the black market, reduces the space for mutually beneficial bargains. While we are unable to explicitly address these issues, we view these as fruitful opportunities for future research on the economics and politics of black markets in conflict.

Nigeria has long been seen as the prototypical example of the resource curse. As such, it provides an ideal setting for studying the broader phenomena of bargaining, conflict, and black markets for natural resources. The lessons from Nigeria can be applied to contexts in which violent conflict over resource rents and black markets are salient features of the environment.

References

Acemoglu, Daron, Thierry Verdier, and James A Robinson (2004). "Kleptocracy and divide-and-rule: A model of personal rule". In: *Journal of the European Economic Association* 2.2-3, pp. 162–192.

Amnesty International (2018). *Negligence in the Niger Delta: Decoding Shell and Eni's Poor Record on Oil Spills*. Tech. rep. Amnesty International.

Angrist, Joshua D. and Adriana D. Kugler (2008). "Rural Windfall or a New Resource Curse? Coca, Income, and Civil Conflict in Colombia". In: *The Review of Economics and Statistics* 90.2, pp. 191–215.

Arkhangelsky, Dmitry et al. (2019). "Synthetic Differences in Differences". In: *NBER Working Paper 25532*.

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

Bannon, Ian and Paul Collier (2003). *Natural resources and violent conflict: Options and actions*. The World Bank.

Bazzi, Samuel and Christopher Blattman (2014). "Economic Shocks and Conflict: Evidence from Commodity Prices". In: *American Economic Journal: Macroeconomics* 6.4.

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

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

Chassang, Sylvain and Gerard Padro-i-Miquel (2009). "Economic shocks and civil war". In: *Quarterly Journal of Political Science* 4.3, pp. 211–228.

Collier, Paul and Anke Hoeffler (2004). "Greed and grievance in civil war". In: *Oxford economic papers* 56.4, pp. 563–595.

— (2005). "Resource rents, governance, and conflict". In: *Journal of conflict resolution* 49.4, pp. 625–633.

Collier, Paul, Anke Hoeffler, and Dominic Rohner (2009). "Beyond greed and grievance: feasibility and civil war". In: *oxford Economic papers* 61.1, pp. 1–27.

Collier, Paul, Anke Hoeffler, and Mans Soderbom (2008). "Post-Conflict Risks". In: *Journal of Peace Research* 45.4, pp. 461–478.

Dal Bó, Ernesto and Robert Powell (2009). "A model of spoils politics". In: *American Journal of Political Science* 53.1, pp. 207–222.

Dancy, Geoff (2018). "Deals with the Devil? Conflict Amnesties, Civil War, and Sustainable Peace". In: *International Organization* 72.Spring, pp. 387–421.

Dell, Melissa (2015). "Trafficking Networks and the Mexican Drug War". In: *American Economic Review* 105.6, pp. 1738–1779.

Dube, Oeindrila and Juan F. Vargas (2013). "Commodity Price Shocks and Civil Conflict: Evidence from Colombia". In: *The Review of Economic Studies* 40.4.

Ebiede, Tarila Marclint, Arnim Langer, and Jale Tosun (2020). "Disarmament, Demobilisation, and Reintegration: Analysing the Outcomes of Nigeria's Post-Amnesty Programme". In: *Stability: International Journal of Security and Development* 9.1, pp. 1–17.

Economist (2016). *The Niger Delta Avengers: Danegeld in the Delta*. Tech. rep.

Ellingsen, Tore and Topi Miettinen (2008). "Commitment and conflict in bilateral bargaining". In: *American Economic Review* 98.4, pp. 1629–35.

Fearon, James (2007). "Fighting rather than Bargaining". In: *Annual Meetings of the American Political Science Association*.

— (1995). "Rationalist Explanations for War". In: *International Organization* 49.3, pp. 379–414.

Fearon, James, Macartan Humphreys, and Jeremy M. Weinstein (2009). "Can Development Aid Contribute to Social Cohesion after Civil War? Evidence from a Field Experiment in Post-Conflict Liberia". In: *American Economic Review: Papers and Proceedings* 99.2, pp. 287–291.

Fearon, James and David D Laitin (2003). "Ethnicity, insurgency, and civil war". In: *American political science review* 97.1, pp. 75–90.

Fetzer, Thiemo and Stephan Kyburz (2018). "Cohesive Institutions and Political Violence". In: *CAGE Working Paper, University of Warwick*.

Francois, Patrick, Ilia Rainer, and Francesco Trebbi (2014). "How is Power Shared in Africa". In: *Econometrica* 83.2, pp. 465–503.

Freixas, Xavier, Roger Guesnerie, and Jean Tirole (1985). "Planning under incomplete information and the ratchet effect". In: *The Review of Economic Studies* 52.2, pp. 173–191.

Gehrin, Kai, Sarah Langlotz, and Stefan Kienberger (2018). "Stimulant or depressant? Resource-related income shocks and conflict". In: *University of Heidelberg Economics Discussion Paper Series* 652.

Ghosh, Sambuddha, Gabriele Gratton, and Caixia Shen (2019). "Intimidation: Linking negotiation and conflict". In: *International Economic Review*.

Grossman, Herschel I (1999). "Kleptocracy and revolutions". In: *Oxford Economic Papers* 51.2, pp. 267–283.

Hausman, Catherine and David S Rapson (2018). "Regression Discontinuity in Time: Considerations for Empirical Applications". In: *Annual Review of Resource Economics* 10, pp. 533–552.

Hazen, Jennifer M and Jonas Horner (2007). "Small Arms, Armed Violence, and Insecurity in Nigeria: The Niger Delta in Perspective". In: *Small Arms Survey Occasional Paper*.

Hegre, Håvard and Nicholas Sambanis (2006). "Sensitivity analysis of empirical results on civil war onset". In: *Journal of conflict resolution* 50.4, pp. 508–535.

Hodler, Roland (2006). "The curse of natural resources in fractionalized countries". In: *European Economic Review* 50.6, pp. 1367–1386.

Holmstrom, Bengt (1982). "Design of incentive schemes and the new Soviet incentive model". In: *European Economic Review* 17.2, pp. 127–148.

Katsouris, Christina and Aaron Sayne (2013). *Nigeria's criminal crude: international options to combat the export of stolen oil*. Chatham House London.

Konig, Michael D et al. (2017). "Networks in Conflict: Theory and Evidence from the Great War of Africa". In: *Econometrica* 85.4, pp. 1093–1132.

Laffont, Jean-Jacques and Jean Tirole (1990). "Adverse selection and renegotiation in procurement". In: *The Review of Economic Studies* 57.4, pp. 597–625.

— (1988). "The dynamics of incentive contracts". In: *Econometrica: Journal of the Econometric Society*, pp. 1153–1175.

Lütkepohl, Helmut (2005). *New Introduction to Multiple Time Series Analysis*. Vol. 1. New York: Springer.

Mehlum, Halvor, Karl Moene, and Ragnar Torvik (2006). "Institutions and the resource curse". In: *The economic journal* 116.508, pp. 1–20.

Miceli, Andrea Di (2017). "Chasing the Key Player: A Network Approach to the Myanmar Civil War". In: *Working Paper*.

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

Nielsen, Bent (2006). "Order determination in general vector autoregressions". In: *Institute of Mathematical Statistics – Monograph Series* 52, pp. 93–112.

Nillesen, Eleonora and Erwin Bulte (2014). "Natural Resources and Violent Conflict". In: *Annual Review of Resource Economics* 6, pp. 69–83.

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

Ojakorotu, Victor and Lysias Dodd Gilbert (2010). *Oil Violence in Nigeria: Checkmating its Resurgence in the Niger delta*. Vol. 1. Lambert Academic Publishing.

Perron, Pierre (2006). "Dealing with structural breaks". In: *Palgrave Handbook of Econometrics: Econometric Theory, Vol I*. Ed. by Terence C. Mills and Kerry Patterson. Basingstoke, UK: Palgrave, pp. 278–352.

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

Ralby, Ian and David Soud (2018). *Oil on the Water: Illicit Hydrocarbons Activity in the Maritime Domain*. Tech. rep. The Atlantic Council.

Robinson, James A, Ragnar Torvik, and Thierry Verdier (2006). "Political foundations of the resource curse". In: *Journal of development Economics* 79.2, pp. 447–468.

Romero, Mauricio and Santiago Saavedra (2019). "Local Incentives and National Tax Evasion: Unintended Effects of a Mining Royalties Reform in Colombia". In: *Working Paper*.

Ross, Michael L. (1999). "The Political Economy of the Resource Curse". In: *World Politics* 51.2, pp. 297–322.

— (2015). "What Have We Learned about the Resource Curse?" In: *Annual Review of Political Science* 18, pp. 239–259.

Rustad, Siri Aas and Helga Malmin Binningsbø (2012). "A price worth fighting for? Natural resources and conflict recurrence". In: *Journal of Peace Research* 49.4, pp. 531–546.

Sala-i-Martin, Xavier and Arvind Subramanian (2013). "Addressing the natural resource curse: An illustration from Nigeria". In: *Journal of African Economies* 22.4, pp. 570–615.

Shirkey, Zachary C (2016). "Uncertainty and war duration". In: *International Studies Review* 18.2, pp. 244–267.

Sierra, Raul Sanchez de la (2019). "On the Origins of the State: Stationary Bandits and Taxation in Eastern Congo". In: *Journal of Political Economy* Forthcoming.

Simonsohn, Uri, Joseph P. Simmons, and Leif D. Nelson (2019). "Specification Curve: Descriptive and Inferential Statistics on All Reasonable Specifications". In: *Working Paper*.

Spaniel, William and Peter Bils (2018). "Slow to learn: Bargaining, uncertainty, and the calculus of conquest". In: *Journal of Conflict Resolution* 62.4, pp. 774–796.

Stakeholder Democracy (2019a). *Communities not Criminals: Illegal Oil Refining in the Niger Delta*. Tech. rep. Stakeholder Democracy.

— (2019b). *More money, more problems: Economic dynamics of the artisanal oil industry in the Niger delta over five years*. Tech. rep. Stakeholder Democracy.

— (2019c). *Pipeline surveillance contracts in the Niger Delta*. Tech. rep. Stakeholder Democracy.

Thomas, Jakana L, William Reed, and Scott Wolford (2016). "The rebels' credibility dilemma". In: *International Organization* 70.3, pp. 477–511.

Ugwu, Chukwuka E and A N Oben (2011). "An Assessment Of The Amnesty Programme Implementation In The Niger Delta Of Nigeria (2009-2011)". In: *Working Paper*.

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

Weitzman, Martin L (1976). "The new Soviet incentive model". In: *The Bell Journal of Economics*, pp. 251–257.

Wolford, Scott, Dan Reiter, and Clifford J Carrubba (2011). "Information, commitment, and war". In: *Journal of Conflict Resolution* 55.4, pp. 556–579.

World Peace Foundation (2018). "CAS-Global Ltd. and the Private Nigerian Coast Guard Fleet". In: *Compendium of Arms Trade Corruption*. Accessed 19 November 2019.

Figure 1: Crime and violence, quarterly event-study

(a) Militant attacks (b) Theft

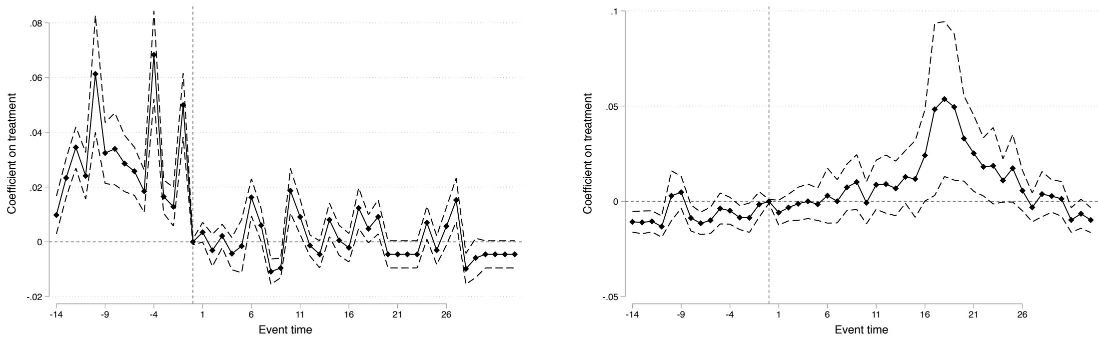

Note: This figure shows trends in militant activity and oil theft before and after amnesty. Coefficients and dashed 95% confidence intervals are from a quarterly event-study regression of the outcome on dummies for periods before and after amnesty, interacted with an indicator for within 30 kilometers from a militant camp, with the initial period of the data as the omitted reference group. Vertical line indicates date of amnesty program.

Figure 2: Pipeline network map

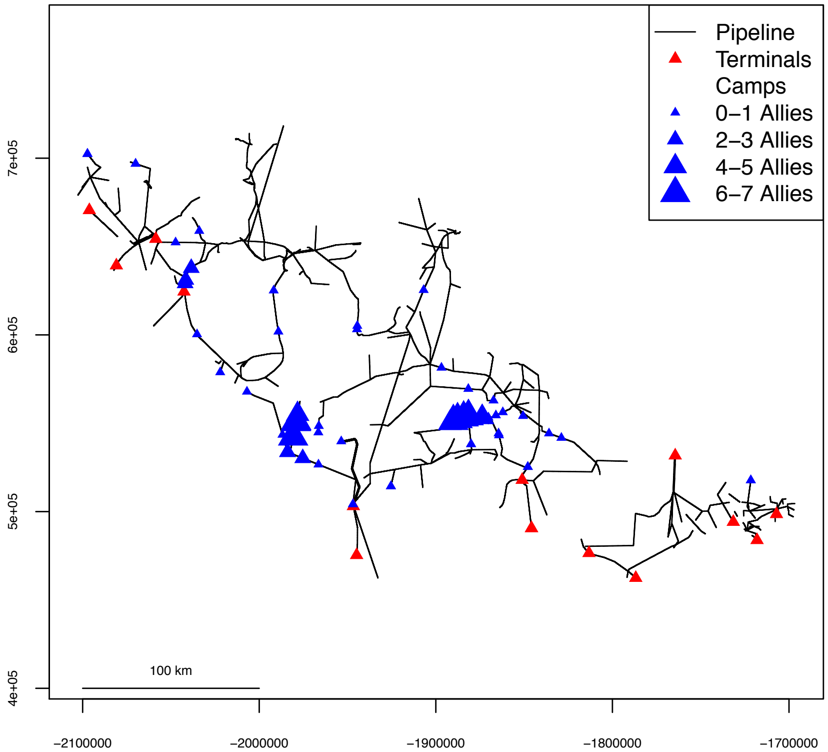

Note: This figure shows the network of oil pipelines in the coastal Niger Delta in black, overlaid with the location of the 69 militant camps in our sample in blue, and the location of oil export terminals (pipeline endpoints) in red. Each militant camp is "snapped" to its nearest pipeline to determine its location in the pipeline network. The size of the camp indicates κ_i^{10}, the number of local allies each camp is connected to within 10 kilometers along the pipeline.

Figure 3: Stage game timing

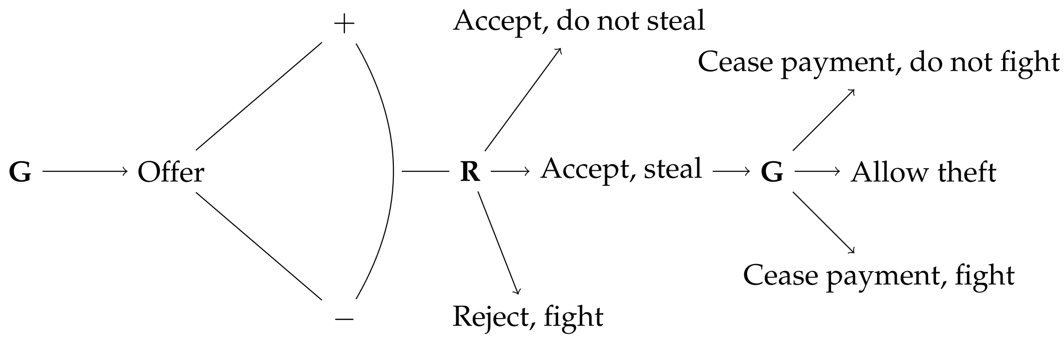

Note: This figure represents stage game timing. **G**, **R** denote government and rebel.

Figure 4: Equilibrium selection

(a) Recovery ψ vs. fringe activity γ (b) Recovery ψ vs. rebel battlefield damage κ

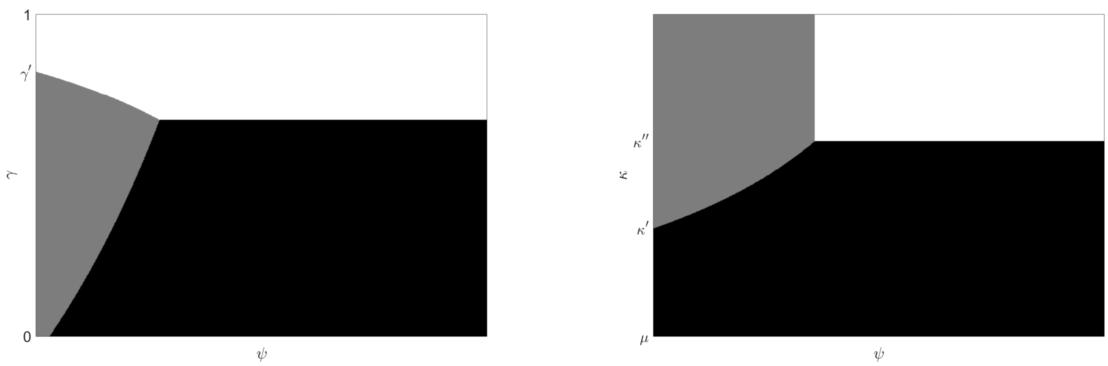

Note: This figure shows subgame perfect equilibrium behavior as a function of parameter values. White, gray- and black correspond to minimal, low- and high-powered incentive provision, see Propositions 1, 2, and 3. The government's ability to recover amnesty payments is increasing on the vertical axis, fringe activity increases along the horizontal axis. Baseline parameters are $\alpha = 1$, $\rho = 3$, $\mu = 15$, $\kappa = 30$, $\psi = .3$, $\gamma = .3$, $\delta = .99$, and $\theta = .95$.

Figure 5: Militant violence, RDiT event-study

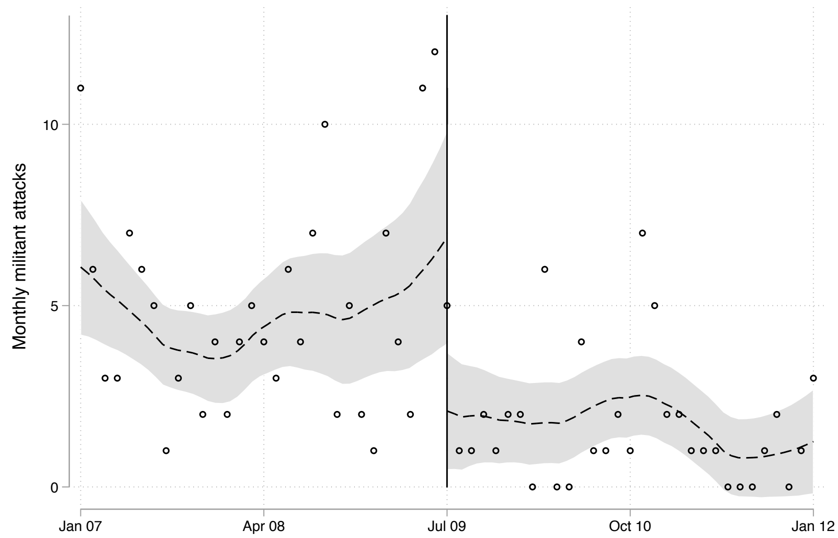

Note: This figure shows the results of a nonparametric RDiT event-study for the July 2009 amnesty date, which is indicated by the vertical solid line. The scatterplot shows the monthly total of oil-related militant attacks over time from ACLED. Dotted lines indicate local polynomial time trends fitted on either side of the event-date cutoff, with shaded area indicating 95% confidence intervals.

Figure 6: Sabotage and malfunction, event-study coefficients

(a) Oil theft (b) Malfunction

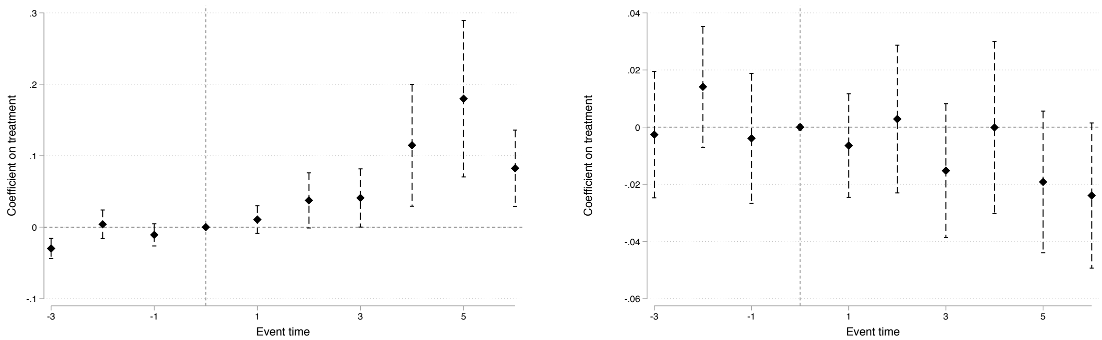

Note: This figure shows the coefficients and 95% confidence intervals from a dynamic differences-in-differences event-study regression of the outcome on dummies for years pre-and-post amnesty, interacted with the treatment indicator, which equals one for villages within 30 km of an amnestied militant camp (see equation 2). 2009 is the omitted reference group. Panel A displays the results with the oil theft outcome, measured as the total annual sabotage-related oil spills. Panel B displays the results for the malfunctions placebo outcome measured as the total annual non-sabotage oil spills.

Figure 7: Nonlinear marginal effects by local wages

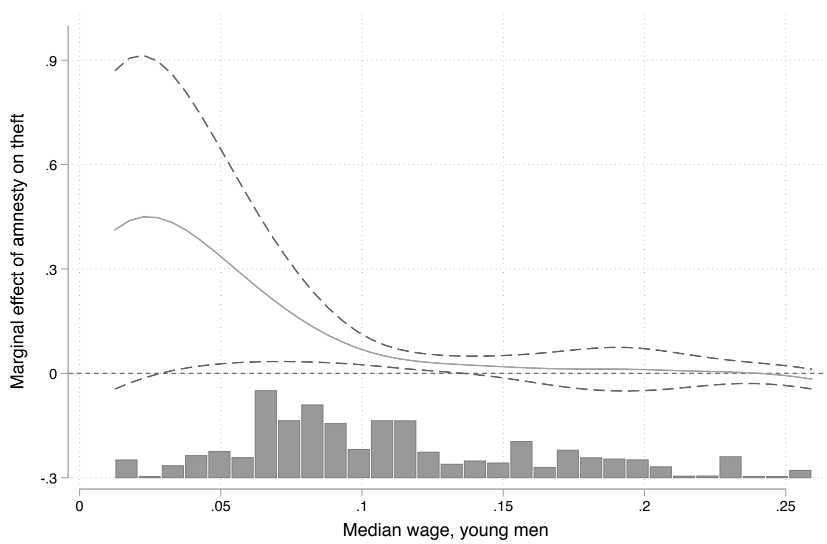

Note: This figure shows the marginal effect of amnesty on sabotage spills by γ, measured as hourly wages for young men in the local labor market. Wages are measured in thousands of Naira per hour. Marginal effects are estimated using a nonparametric kernel regression with a bandwidth of 30, after residualizing year and village fixed effects. Histogram bars indicate the relative size of bins of the wage distribution. Plot excludes outlier wages, defined as above zero and below 300 Naira per hour.

Figure 8: Binned amnesty effects by local alliance density

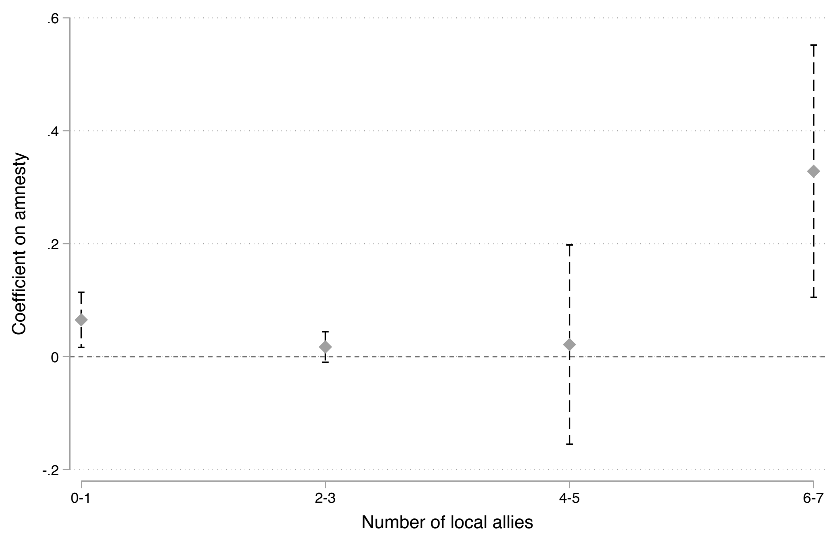

Note: This figure shows the differential effect of amnesty on sabotage spills by κ, measured as the number of allied groups within 10 kilometers along the pipeline. Coefficients are estimated using the difference-in-differences regression in (1) on subsamples for the number of allies, as indicated on the x-axis.

Table 1: The effect of Amnesty on militant events

Polynomial	Linear		Quadratic		Quartic		AR(1)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Full sample								
Post-amnesty	-4.898***	-3.351***	-4.291**	-3.064***	-5.592***	-5.687***	-4.600***	-2.121***
	(1.587)	(0.672)	(1.725)	(0.980)	(1.818)	(1.235)	(1.637)	(0.619)
m_{t-1}							0.078	0.393***
							(0.113)	(0.095)
Month FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	240	240	240	240	240	240	239	239
R^2	0.515	0.276	0.521	0.433	0.536	0.474	0.517	0.388
$\Delta = 30$								
Post-amnesty	-4.701**	-3.274**	-7.791***	-6.935***	-6.589**	-6.869***	-4.224**	-2.766*
	(1.773)	(1.417)	(2.281)	(2.059)	(2.589)	(2.410)	(1.855)	(1.445)
m_{t-1}							0.138	0.205
							(0.162)	(0.151)
Month FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	61	61	61	61	61	61	61	61
R^2	0.372	0.412	0.468	0.499	0.513	0.514	0.385	0.438
$\Delta = 20$								
Post-amnesty	-6.076***	-5.605***	-6.197*	-6.675**	-6.828**	-9.072**	-6.375***	-5.331**
	(1.848)	(1.780)	(3.067)	(2.468)	(2.917)	(4.277)	(2.047)	(2.040)
m_{t-1}							-0.058	0.056
							(0.212)	(0.187)
Month FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	41	41	41	41	41	41	41	41
R^2	0.357	0.531	0.379	0.536	0.568	0.610	0.360	0.533
$\Delta = 10$								
Post-amnesty	-5.982*		-9.699**		-5.918**		-6.598	
	(3.070)		(4.420)		(2.523)		(3.784)	
m_{t-1}							-0.103	
							(0.307)	
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	21	21	21	21	21	21	21	21
R^2	0.373	0.000	0.485	0.000	0.741	0.000	0.379	0.000

Robust standard errors. Outcome variable is the number of monthly oil-related militancy events. Treatment is defined as an indicator for after July 2009. Window refers to the number of months included in the estimation before and after the event date. All windows apart from the full sample are symmetric. AR(1) specifications include a lagged dependent variable and a linear polynomial of event time. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 2: The effect of amnesty on oil theft

Dependent variable	Oil theft							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Within 30 km × Post-amnesty	0.069***	0.067***	0.058***	0.075***	0.041***	0.045***	0.047***	0.064***
	(0.021)	(0.022)	(0.021)	(0.023)	(0.012)	(0.013)	(0.018)	(0.024)
Distance to oil field (00s km) × Post-amnesty		-0.007	0.057*					
		(0.014)	(0.034)					
Distance to pipeline (00s km) × Post-amnesty			-0.072***					
			(0.025)					
Within 30 km × Oil price (USD/barrel)				0.001**		0.001**		
				(0.000)		(0.000)		
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	No	No	No	Yes	Yes	No	No
Village FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE × year	No	No	No	No	No	No	Yes	No
Village FE × year	No	No	No	No	No	No	No	Yes
Observations	139572	139572	139572	139572	135372	135372	139572	139572
R^2	0.416	0.416	0.416	0.416	0.418	0.418	0.419	0.520

Standard errors clustered at the village level. Outcome variable is the number of annual oil bunkering events. Treatment is defined as all villages within 30 km of an amnestied militant camp. Control group is all villages outside of this range. Controls include distance to nearest oilfield, distance to nearest pipeline, distance to state capital, distance to Niger River, distance to the coast, and population density. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: Militant strength and differentiated amnesty

	(1)	(2)	(3)	(4)	(5)
			Panel A: Δ oil production, 2005-2009		
Allied, 10 km	-7.905***	-7.902***	-5.983***	-3.487**	-4.658***
	(1.671)	(1.688)	(1.578)	(1.405)	(1.424)
Non-allied, 10 km		0.162	0.634	0.240	1.355
		(1.553)	(1.812)	(1.365)	(1.367)
MEND				-38.007***	
				(11.419)	
R^2	0.188	0.188	0.471	0.558	0.637
			Panel B: Defeated in battle		
Allied, 10 km	-0.053***	-0.053***	-0.065***	-0.048***	-0.067***
	(0.015)	(0.015)	(0.020)	(0.017)	(0.023)
Non-allied, 10 km		0.034	0.040	0.037	0.045
		(0.025)	(0.032)	(0.032)	(0.033)
MEND				-0.256**	
				(0.125)	
R^2	0.111	0.145	0.319	0.370	0.332
			Panel C: Any amnesty		
Allied, 10 km	0.080***	0.079***	0.090***	0.061**	0.095***
	(0.017)	(0.017)	(0.022)	(0.023)	(0.024)
Non-allied, 10 km		-0.043***	-0.048**	-0.044*	-0.050*
		(0.015)	(0.023)	(0.025)	(0.025)
MEND				0.442***	
				(0.123)	
R^2	0.184	0.223	0.372	0.485	0.423
			Panel D: Surveillance contract		
Allied, 10 km	0.074**	0.074**	0.064**	0.045	0.062**
	(0.033)	(0.033)	(0.028)	(0.032)	(0.029)
Non-allied, 10 km		-0.005	-0.018	-0.015	-0.024
		(0.024)	(0.028)	(0.029)	(0.033)
MEND				0.280	
				(0.209)	
R^2	0.149	0.149	0.337	0.379	0.408
Controls	No	No	Yes	Yes	Yes
State FE	No	No	No	No	Yes
Observations	69	69	69	69	69

Standard errors clustered at the militant commander level. Outcome variable is indicated in panel header. The independent variable of interest is k_i^{10}, the number of allies of camp i within 10 kilometers along the pipeline network. $MEND$ is a dummy variable indicating that the group belongs to the Movement for Emancipation of the Niger Delta. Controls are distance from pipeline, distance from state capital, distance from coast, camp latitude, annual precipitation, monthly average temperature, altitude, and slope. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: The effect of amnesty on oil theft by cost factors

Dependent variable	Oil theft					
	(1)	(2)	(3)	(4)	(5)	(6)
Within 30 km × Post-amnesty	0.063**	0.119***	0.106**	0.195***	0.243**	0.306***
	(0.025)	(0.034)	(0.051)	(0.075)	(0.095)	(0.100)
Within 30 km × Post-amnesty × Distance to coast (km)		-0.002***			-0.002**	-0.003***
		(0.001)			(0.001)	(0.001)
Within 30 km × Post-amnesty × Distance to Niger River (km)			-0.001		-0.001	-0.001**
			(0.000)		(0.000)	(0.000)
Within 30 km × Post-amnesty × Median wage, young men				-1.008***	-0.756**	-0.875**
				(0.389)	(0.330)	(0.348)
Village FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	Yes	Yes	Yes	Yes	No	Yes
Observations	135120	135120	135120	134304	138468	134304
R^2	0.417	0.417	0.417	0.417	0.416	0.418

Standard errors clustered at the village level. Outcome variable is the number of annual oil bunkering events. Treatment is defined as all villages within 30 km of an amnestied militant camp. Control group is all villages outside of this range. Median wage of young men is the median wage of men aged 10-40 in the local government area, measured in 2009. Controls include distance to nearest oilfield, distance to nearest pipeline, distance to state capital, and population density. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: The effect of amnesty on oil theft by alliance density and cost factors

Dependent variable	Oil theft						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Within 30 km × Post-amnesty	0.071***	0.048**	0.021*	0.012	0.062***	0.119	0.098
	(0.023)	(0.022)	(0.012)	(0.025)	(0.021)	(0.078)	(0.065)
Within 30 km × Post-amnesty × Allied camps, 10 km		0.024*	0.032**	0.031**	0.070***	0.074**	0.075**
		(0.013)	(0.014)	(0.014)	(0.026)	(0.033)	(0.033)
Within 30 km × Post-amnesty × Median wage, young men						-0.638	-0.583
						(0.441)	(0.418)
Within 30 km × Post-amnesty × Allied camps, 10 km × Median wage, young men						-0.342**	-0.323*
						(0.172)	(0.170)
Village FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	No	Yes	Yes	Yes	No	Yes
MEND controls	No	No	No	Yes	No	No	No
Density controls	No	No	No	No	Yes	No	No
Observations	134304	134304	134304	134304	134304	134304	134304
R^2	0.416	0.417	0.418	0.419	0.419	0.419	0.420

Standard errors clustered at the village level. Outcome variable is the number of annual oil bunkering events. Treatment is defined as all villages within 30 km of an amnestied militant camp. Control group is all villages outside of this range. Controls include distance to nearest oilfield, distance to nearest pipeline, distance to state capital, distance to Niger River, distance to the coast, and population density. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A Graphs

Figure A1: Oil spills by cause over time

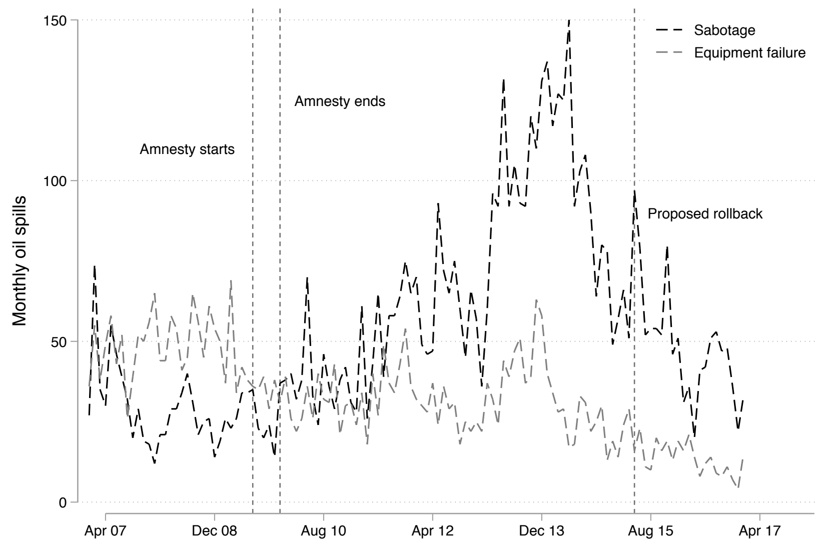

Note: This figure shows total monthly sabotage and non-sabotage malfunction spills from the NOSDRA data over time. Important dates in the amnesty process are indicated by labeled vertical lines.

Figure A2: Map of the Niger Delta

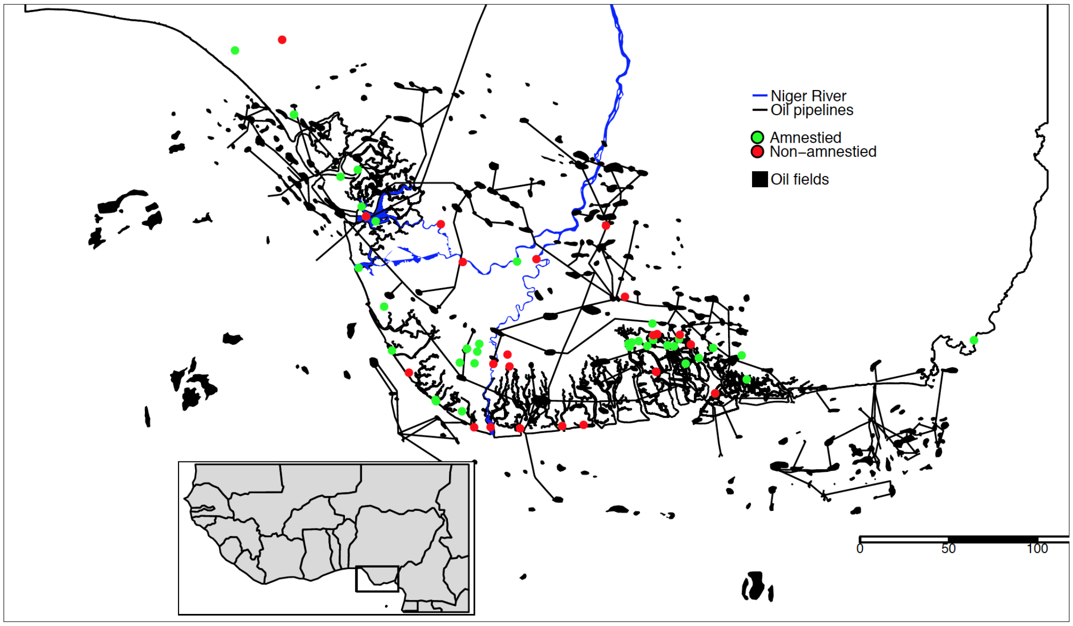

Note: This figure shows geography of amnesty, oil production, and militant activity in the Niger Delta. The southern coastline of the Niger Delta is outlined in black, overlaid with the locations of oil fields, pipelines, and the Niger River. Points indicate the locations of militant camps in our sample, color-coded by their amnesty status. The map inset indicates the location of the Niger Delta on the coast of West Africa.

Figure A3: Crime and violence over time

(a) Militant attacks (b) Oil production

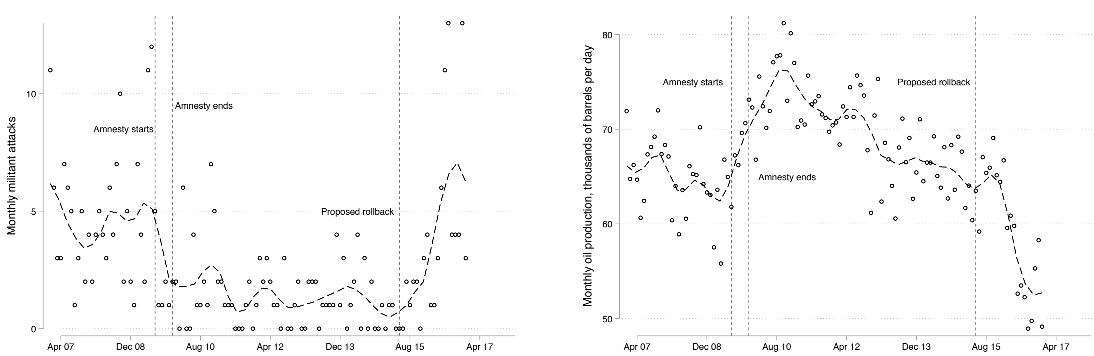

Note: This figure shows total monthly militant attacks from ACLED in Panel A and oil production from the DPR in Panel B over time. Important dates in the amnesty process are indicated by labeled vertical lines.

Figure A4: Rebel payoffs

(a) Recovery ψ vs. fringe activity γ (b) Recovery ψ vs. rebel battlefield damage κ

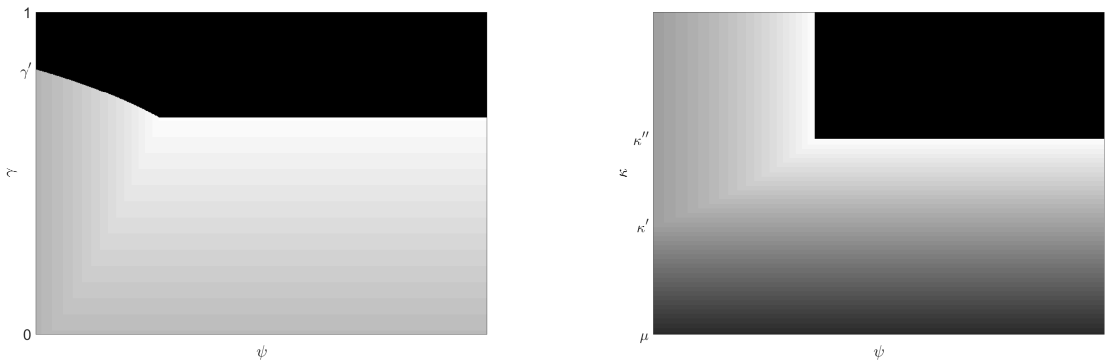

Note: This figure shows the subgame perfect equilibrium payoffs of the rebels as a function of parameter values. Lighter colors indicate greater payoffs, with the rebels reservation payoff (3) denoted in black. The government's ability to recover amnesty payments is increasing on the vertical axis, fringe activity increases along the horizontal axis. Baseline parameters are $\alpha = 1$, $\rho = 3$, $\mu = 15$, $\kappa = 30$, $\psi = .3$, $\gamma = .3$, $\delta = .99$, and $\theta = .95$.

Figure A5: RDiT event-study by treatment status

(a) < 30 km from amnestied militant camp (b) > 30 km from amnestied militant camp

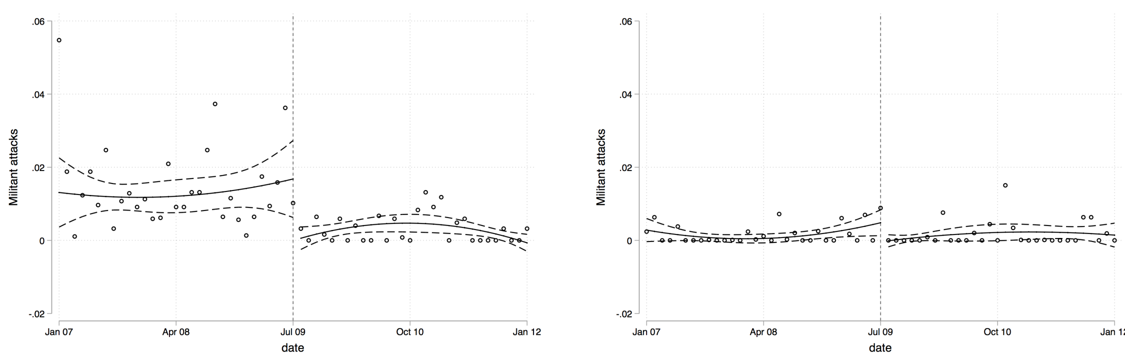

Note: This figure shows the RDit event-study on militant activity separately for treatment and control groups. Each panel estimates a quadratic RDiT event-study for an event-window of 30 months before and after the amnesty date, with parametric trends and dashed 95% confidence intervals estimated separately on either side of the event. These trends are overlaid with a scatterplot of mean monthly village-level oil-related militant attacks over time. Panel A estimates the RDiT only on the sample of treated villages within 30 km of an amnestied camp, while Panel B does the same for control villages.

Figure A6: RDiT placebo test

(a) Regression coefficient β (b) t-statistic

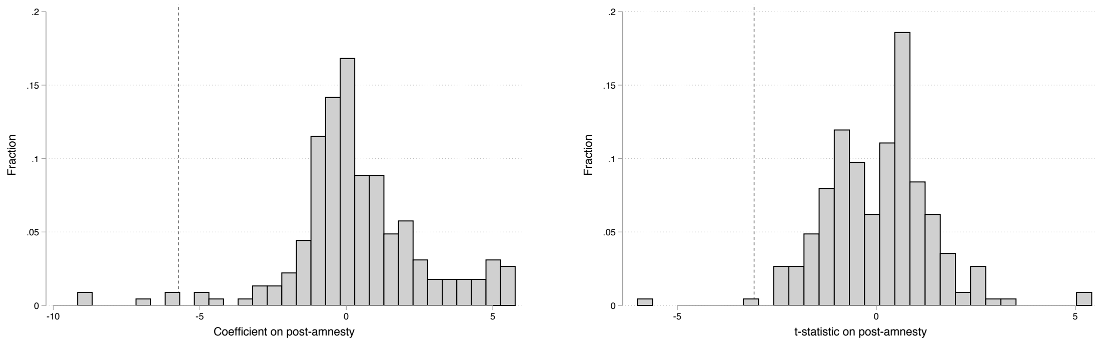

Note: This figure plots a histogram of the RDiT event-study placebo test for militant activity. We re-estimate the monthly RDiT equation in Section 4.1 with quadratic time trends and year fixed effects for each of 225 possible event-dates between July 1997 and April 2016. Panel A then plots the distribution of treatment coefficients from these regressions while Panel B plots the associated t-statistics. The dotted vertical line indicates the estimate for July 2009, the true amnesty date.

Figure A7: Event-study influential observations test

(a) Regression coefficient β (b) t-statistic

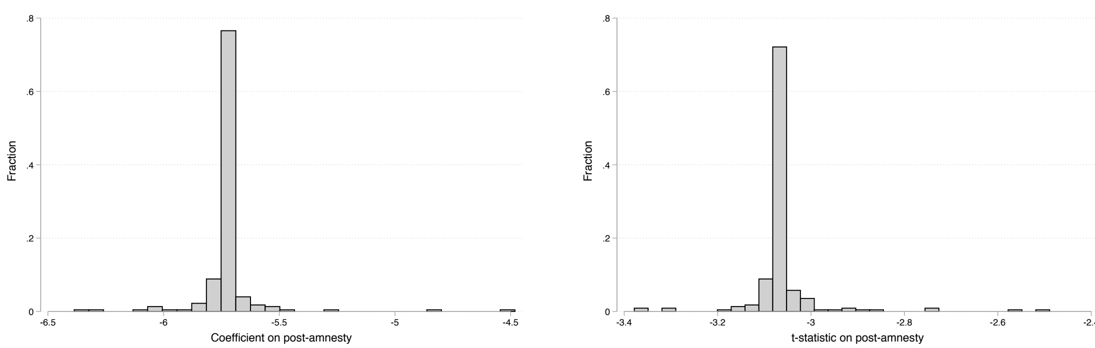

Note: This figure plots RDiT event-study estimates from a leave-one-out influential observations test. We re-estimate the monthly RDiT equation in Section 4.1 with quadratic time trends and year fixed effects, dropping each of 225 possible event-dates between July 1997 and April 2016. Panel A plots the distribution of treatment coefficients from these regressions while Panel B plots the associated t-statistics.

Figure A8: Histograms of Wald statistics for a structural break test at varying bandwidths

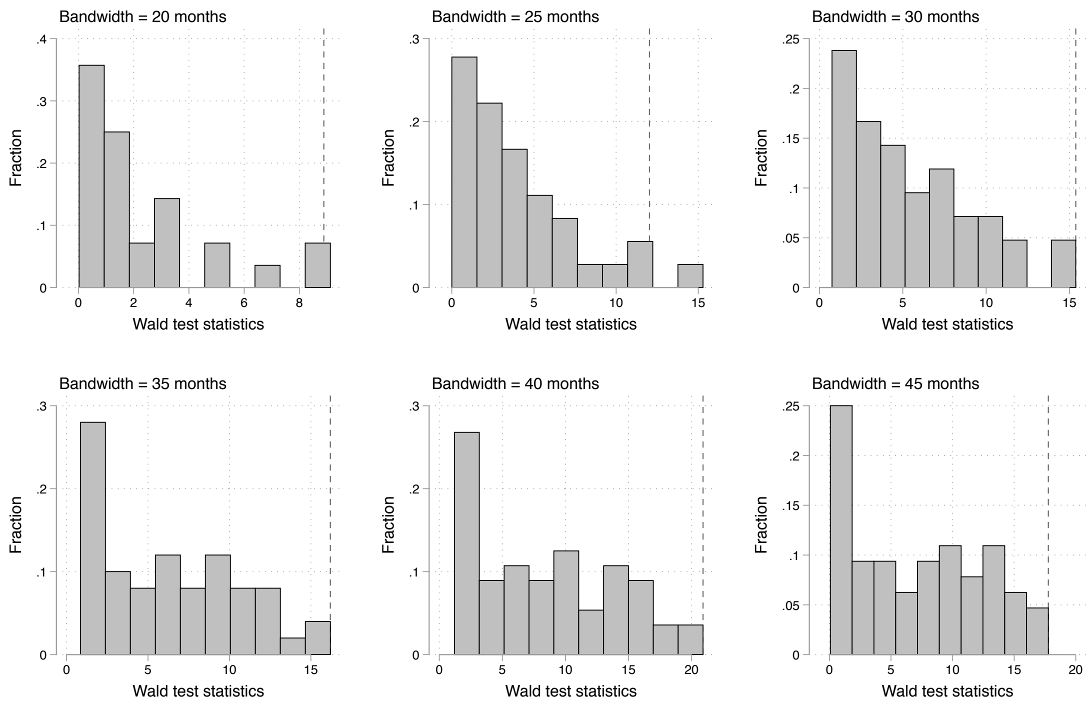

Note: This figure displays Wald statistics for structural break tests by month for varying event-windows, indicated in the figure sub-header. Vertical line indicates the Wald statistic on a structural break test of July 2009, the true amnesty date. The outcome variable is oil-related militancy. See E.1 for greater detail on the structural break test.

Figure A9: RDiT event-study by local alliance density

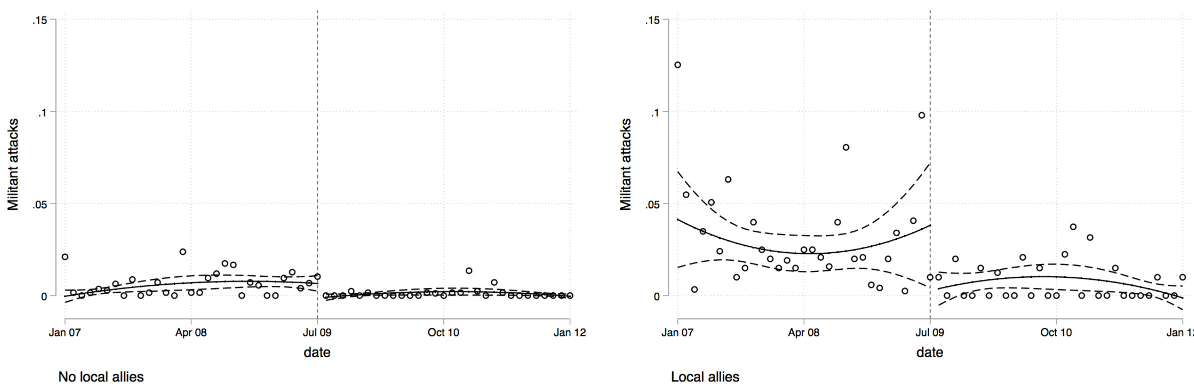

Note: This figure shows the RDiT event-study on militant activity separately for villages with militant groups of low and high military strength. Each panel estimates a quadratic RDiT event-study for an event-window of 30 months before and after the amnesty date, with parametric trends and dashed 95% confidence intervals estimated separately on either side of the event. These trends are overlaid with a scatterplot of mean monthly village-level oil-related militant attacks over time. The righthand panel estimates the RDiT only on the sample of villages whose nearest militant camp has no local allies, while the leftmost panel does the same for militants with at least one local ally.

Figure A10: Heterogeneity by surveillance contracts

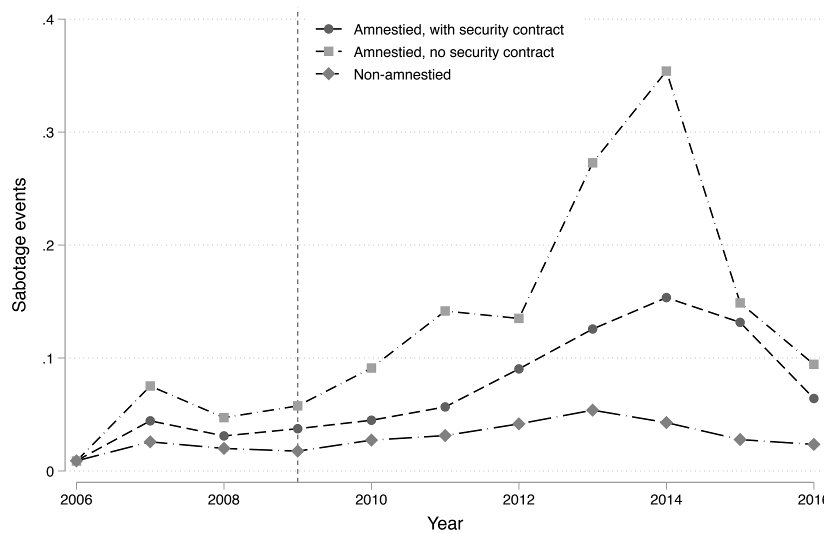

Note: This figure shows estimated mean annual village-level oil theft sabotage events by amnesty and surveillance contract status. Amnestied without contract are those within 30 kilometers of an amnestied militant camp, in which the nearest militant commander did not receive a pipeline surveillance contact, while amnestied with contract are villages in the treatment group in which the nearest militant commander did receive a pipeline surveillance contact. Non-amnestied are villages more than 30 kilometers from an amnestied militant camp. Dashed vertical line indicates 2009, the year of the amnesty.

Figure A11: RDiT event-study on Buhari announcement of amnesty funding cuts

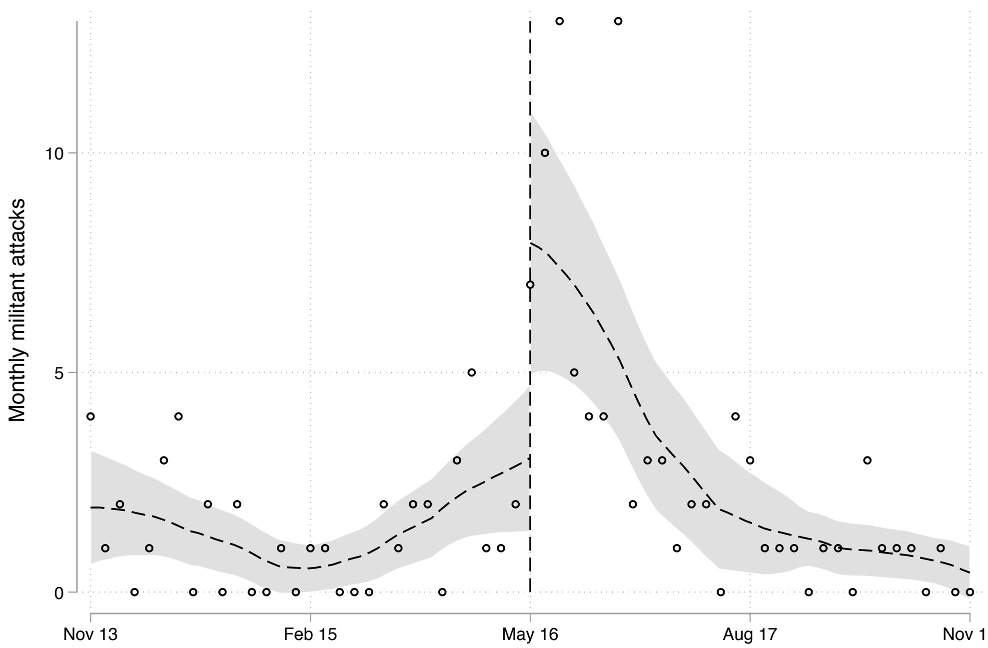

Note: This figure shows the results of a nonparametric RDiT event-study for the May 2016 date on which President Buhari announced a budget with a 70% cut in amnesty funding, which is indicated by the vertical solid line. The scatterplot shows the monthly total of oil-related militant attacks over time from ACLED. Dotted lines indicate local polynomial time trends fitted on either side of the event-date cutoff, with shaded area indicating 95% confidence intervals.

Figure A12: RDiT placebo test on Buhari announcement of amnesty funding cuts

(a) Regression coefficient β (b) t-statistic

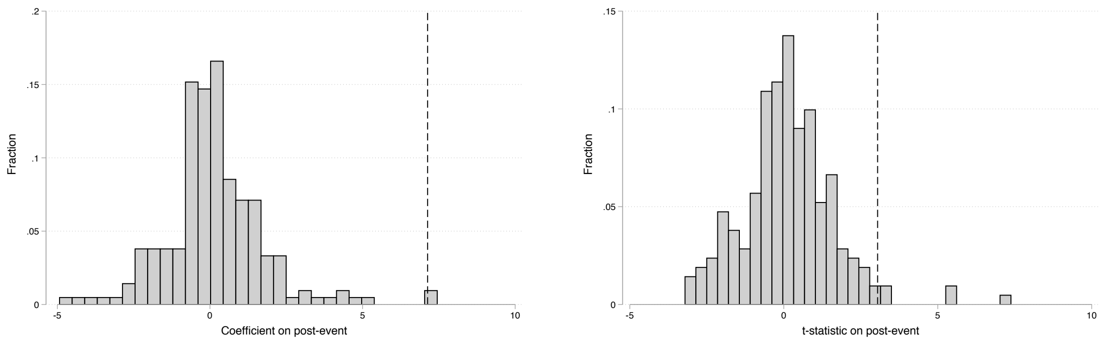

Note: This figure plots a histogram of the RDiT event-study placebo test for militant activity. We re-estimate the monthly RDiT equation in Section 4.1 with quadratic time trends and year fixed effects for each of 211 possible event-dates between September 2001 and March 2019. Panel A then plots the distribution of treatment coefficients from these regressions while Panel B plots the associated t-statistics. The dotted vertical line indicates the estimate for May 2016, the date on which President Buhari announced a budget with a 70% cut in amnesty funding.

Figure A13: Local alliance density and damage

(a) Allied camps (b) Non-allied camps

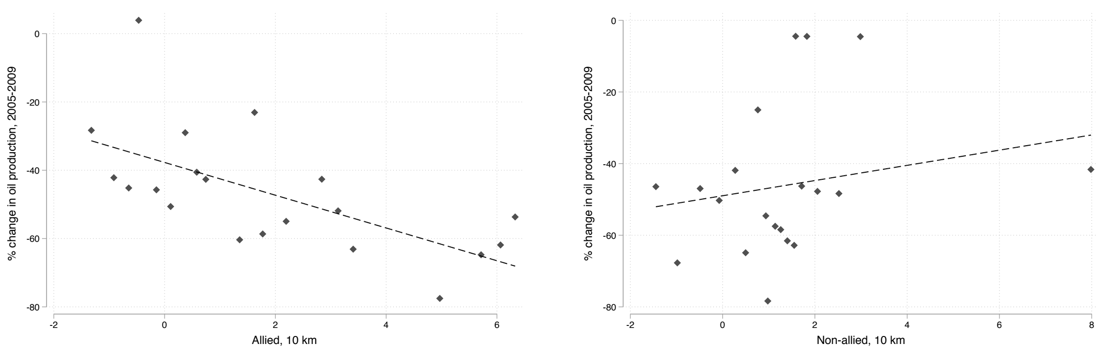

Note: This figure shows a binned scatterplot of the militant-camp-level relationship between the number of allied connections along the pipeline and damage inflicted during the height of the Niger Delta Crisis. Damage is measured as the % change in onshore oil production within 20 km of the militant camp location between 2005 and 2009. Panel A uses the number of allied camps within 10 km along the pipeline as the independent variable, while Panel B uses the number of non-allied camps. Scatterplots include state fixed-effects and camp-level controls for slope, altitude, average temperature, average precipitation, latitude, and distance to the nearest pipeline, state capital, and Atlantic coast.

Figure A14: Camp-level output change specification plot

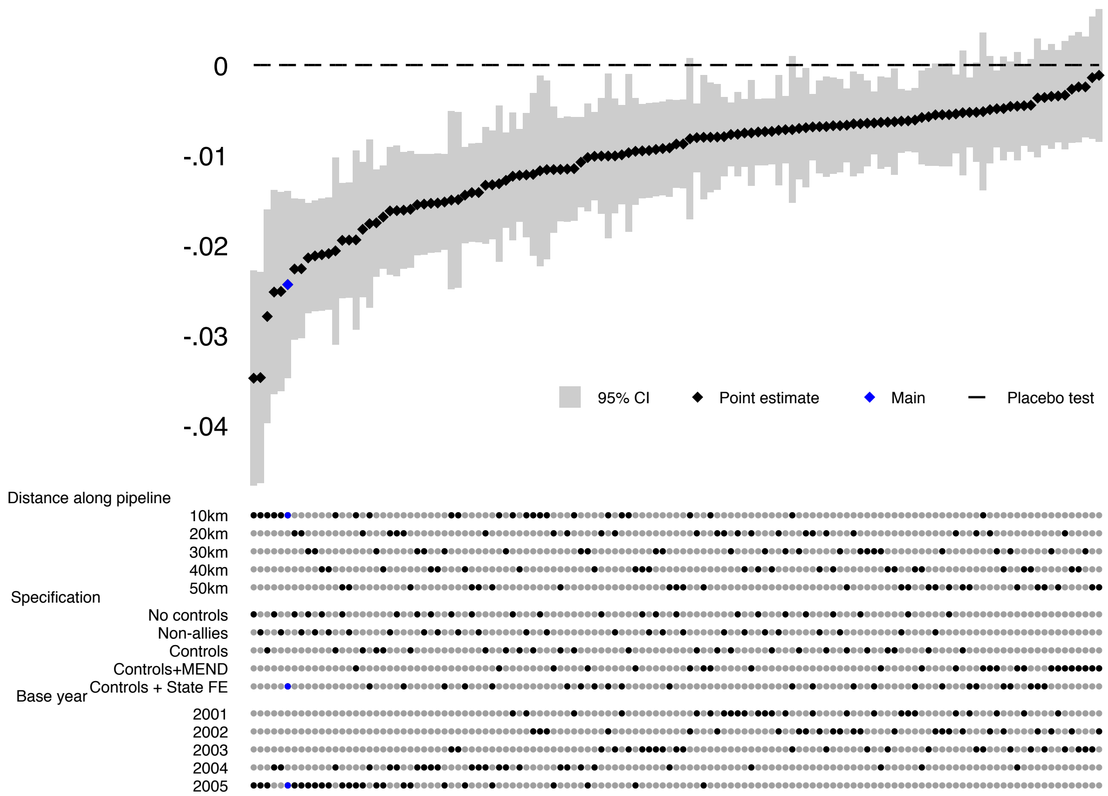

Note: This figure shows robustness across 125 specifications for a regression of camp-level percent output change on the number of allied camps within 10 km along the pipeline, as estimated in equation 13. Each point represents the estimated coefficient of an individual specification surrounded by 95% confidence bars. Specification set is all combinations of: *i*) varying the distance δ along the pipeline over which k_i^δ is defined from 10 to 50 km, *ii*) inclusion of controls, MEND dummy, state FE, and non-allies, *iii*) varying the base-year over which the output change outcome variable is defined from 2001 to 2005. Specification type is indicated in the figure footer. Main estimate is indicated in blue.

Figure A15: Local alliance density and military defeat

(a) Allied (b) Non-allied

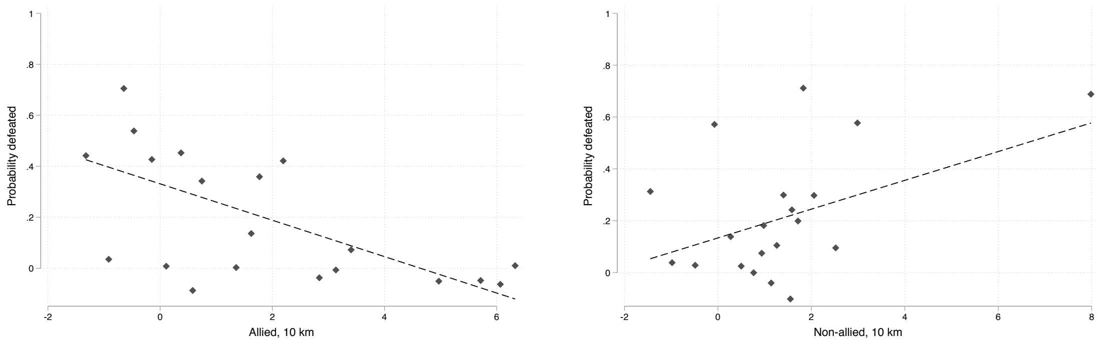

Note: This figure shows a binned scatterplot of the militant-camp-level relationship between the number of allied connections along the pipeline and defeat probability. Defeat is measured as a dummy variable that equals one if the camp is controlled by a commander who was arrested or killed during the conflict. Panel A uses the number of allied camps within 10 km along the pipeline as the independent variable, while Panel B uses the number of non-allied camps. Scatterplots include state fixed-effects and camp-level controls for slope, altitude, average temperature, average precipitation, latitude, and distance to the nearest pipeline, state capital, and Atlantic coast.

Figure A16: Local alliance density and amnesty

(a) Allied (b) Non-allied

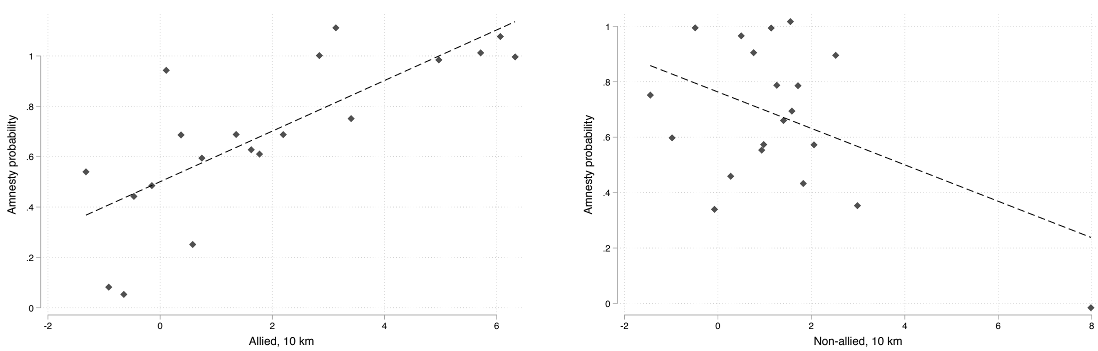

Note: This figure shows a binned scatterplot of the militant-camp-level relationship between the number of allied connections along the pipeline and the probability of receiving any amnesty deal. Amnesty is measured as a dummy variable that equals one if the camp is controlled by a commander who was received amnesty. Panel A uses the number of allied camps within 10 km along the pipeline as the independent variable, while Panel B uses the number of non-allied camps. Scatterplots include state fixed-effects and camp-level controls for slope, altitude, average temperature, average precipitation, latitude, and distance to the nearest pipeline, state capital, and Atlantic coast.

Figure A17: Local alliance density and surveillance contracts

(a) Allied (b) Non-allied

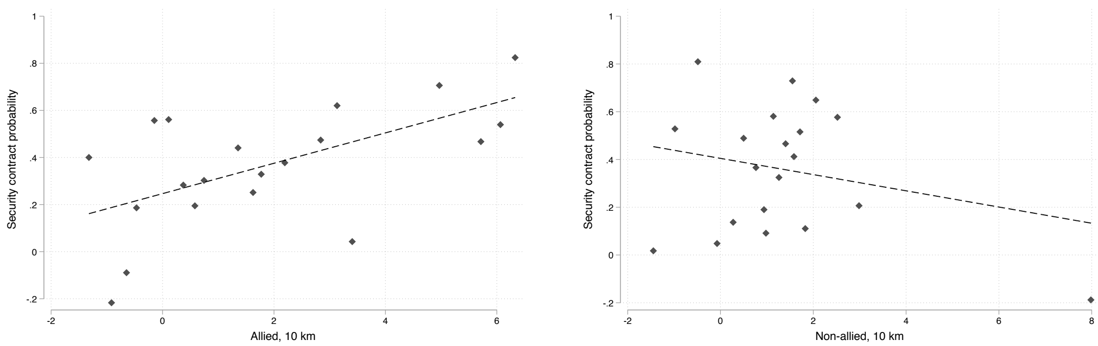

Note: This figure shows a binned scatterplot of the militant-camp-level relationship between the number of allied connections along the pipeline and the probability of receiving a surveillance contract. This is measured as a dummy variable that equals one if the camp is controlled by a commander who was received a pipeline security contract. Panel A uses the number of allied camps within 10 km along the pipeline as the independent variable, while Panel B uses the number of non-allied camps. Scatterplots include state fixed-effects and camp-level controls for slope, altitude, average temperature, average precipitation, latitude, and distance to the nearest pipeline, state capital, and Atlantic coast.

Figure A18: Marginal effect of κ on contract receipt by levels of γ

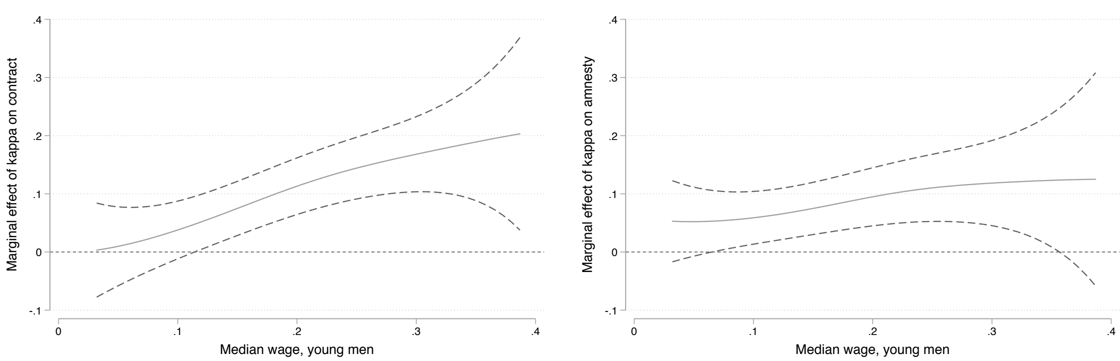

Note: This figure shows the marginal effect of the number of allied connections along the pipeline on the probability of receiving a surveillance contract (left panel) or any amnesty (right panel) at different points along the local young mens' wage distribution. Marginal effects are estimated using a kernel regression of contract outcomes on the interaction between the number of local allies and local wages. All regressions include state fixed-effects and camp-level controls for slope, altitude, average temperature, average precipitation, latitude, and distance to the nearest pipeline, state capital, and Atlantic coast. Wages are measured as the median nominal wages in thousands of Nigerian Naira for men aged 10-40 in the local government area in which a camp is located.

Figure A19: Parallel trends plot

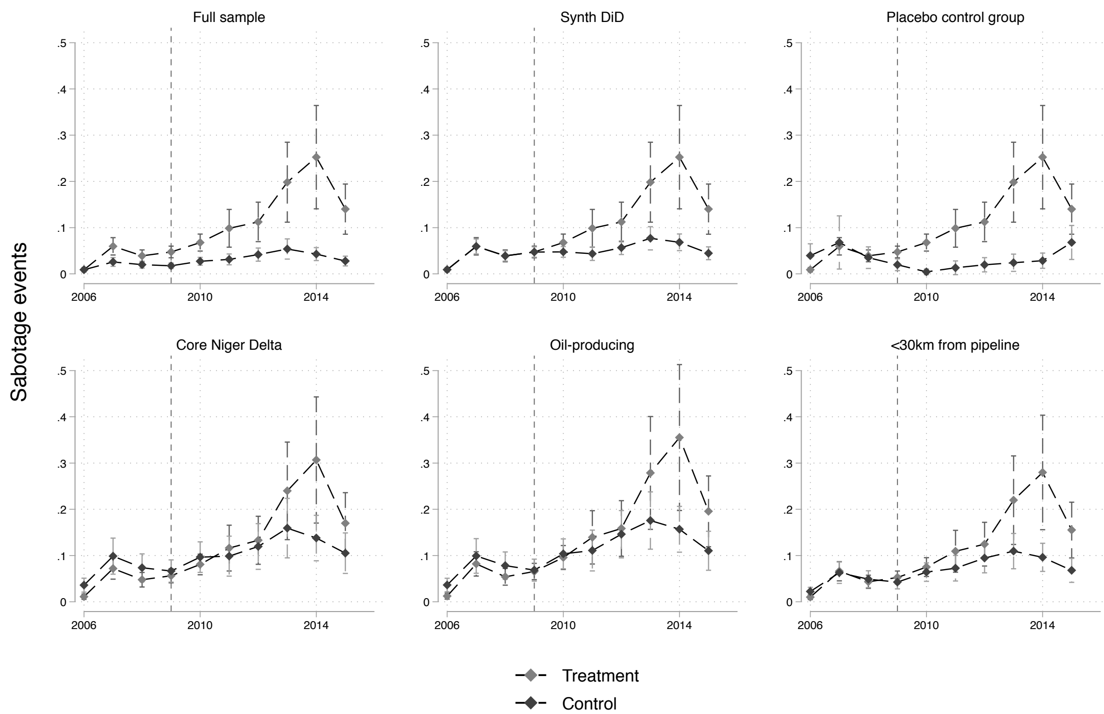

Note: This figure shows estimated mean annual village-level theft incidents with 95% confidence intervals by treatment status across various permutations of control group. Full sample uses all villages and defines the treatment group as those within 30 km of an amnestied camp. Synth DD uses the same treatment group but re-weights the control group to match the pre-trends in oil theft, following Arkhangelsky et al. (2019). Placebo control group restricts the control group to only villages within 30 km of a non-amnestied militant camp. Core Niger Delta maintains the same treatment definition, but restricts the sample to Bayelsa, Rivers, and Delta states. Oil-producing restricts the sample to villages within 10 km of oil infrastructure. Dashed vertical line indicates 2009, the year of the amnesty.

Figure A20: Sabotage parallel trends, pipeline-level

(a) Sabotage (b) Normalized sabotage

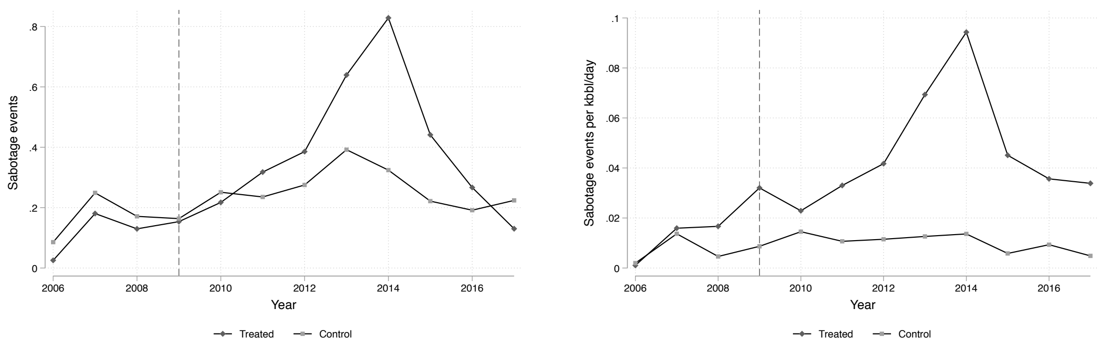

Note: This figure shows estimated mean annual pipeline-segment-level sabotage events (Panel A) and normalized sabotage events (Panel B) by treatment and control group. Normalized sabotage events are defined as total sabotage events divided by total oil flow in pipeline segment i in year t, as detailed in Section E.2. Oil flow is calculated as the average daily barrels flowing through a pipeline segment from all upstream fields in the given year. Treated pipeline-segments are those within 30 kilometers of the nearest amnestied militant camp. Dashed vertical line indicates 2009, the year of the amnesty.

Figure A21: Difference-in-differences estimation, robustness to treatment radius

(a) Sabotage (b) Malfunction

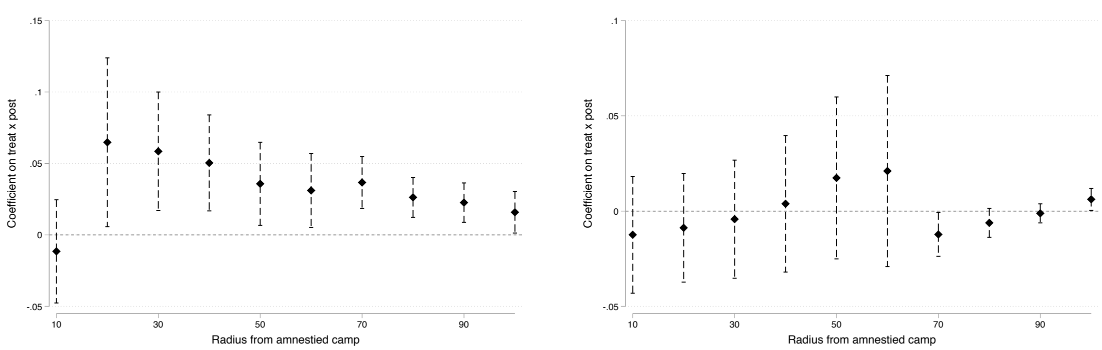

Note: This figure shows estimated difference-in-differences coefficients for treatment definitions of varying distance, for both oil theft (Panel A), measured as village-level annual sabotage events, and the placebo outcome equipment malfunctions (Panel B). The treatment indicator is a dummy equaling one if the village is within d kilometers from the nearest amnestied camp; d is varied between 10 and 100 kilometers, as indicated.

Figure A22: Difference-in-differences oil theft specification plot

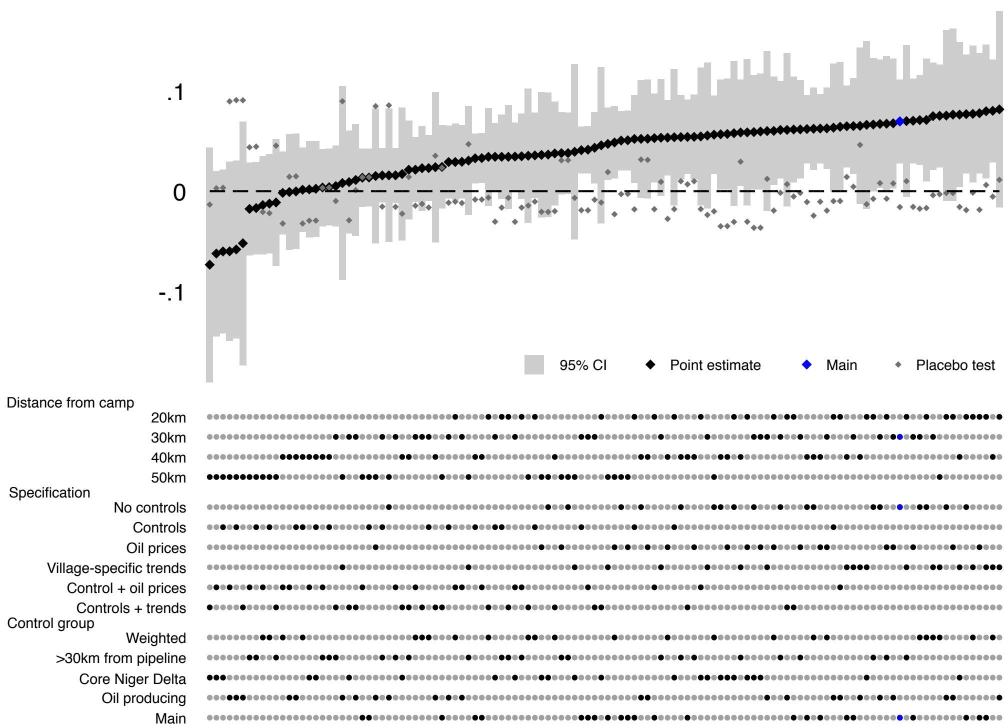

Note: This figure shows robustness across 120 specifications for a difference-in-differences regression of village-level oil theft on the amnesty treatment, as estimated in equation 1. Each black point represents the estimated coefficient of an individual specification surrounded by 95% confidence bars. Points in grey represent the estimates from the corresponding placebo test in which equipment malfunctions is the outcome variable. The preferred estimate is indicated in blue. Specification set is all combinations of: *i*) varying the distance *d* from an amnestied militant camp within which the treatment is defined from 20 to 50, *ii*) inclusion of controls, oil price interactions, village-specific linear time trends, and combinations of these *iii*) varying the control group between the main specification, the re-weighted synth-DD, the oil producing sample, the core Niger Delta sample, and the sample within 30 km of a pipeline. Specification type is indicated in the figure footer.

Figure A23: Amnesty impacts by oil infrastructure type

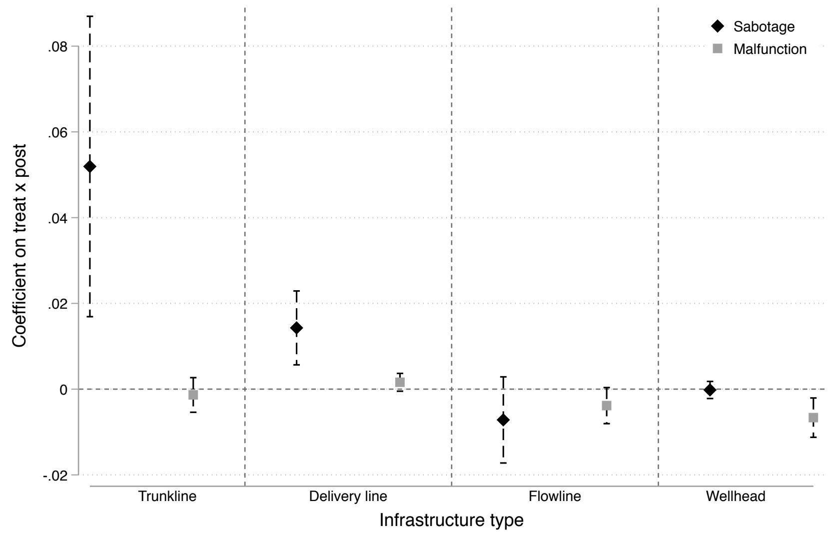

Note: This figure shows the effect of amnesty on sabotage spills by infrastructure type. The plot contains coefficients and 95% confidence intervals from a differences-in-differences regression of the outcome on the interaction between treatment and post-amnesty. The outcome variables are, respectively, total theft or malfunction spills that arise specifically on trunklines, delivery lines, flowlines, or wellheads. The infrastructure type is indicated in the x-axis, while the marker shape indicates whether theft or the malfunctions placebo is the outcome.

Figure A24: Event-study heterogeneity by local wages

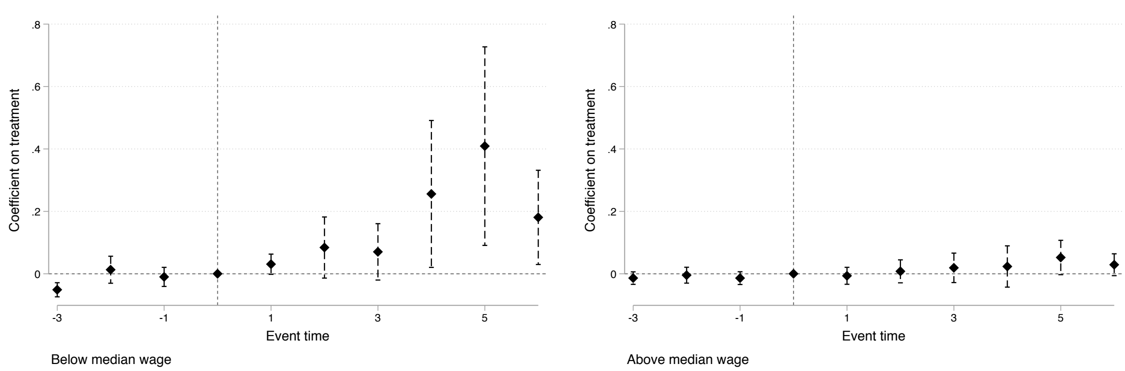

Note: This figure shows the coefficients and 95% confidence intervals from a dynamic differences-in-differences event-study regression of the outcome on dummies for years pre-and-post amnesty, interacted with the treatment indicator, which equals one for villages within 30 km of an amnestied militant camp (see equation 2). 2009 is the omitted reference group and the outcome variable is oil theft, measured as the total annual sabotage-related oil spills. The sample is split by local labor markets below and above the median of the distribution of market-level wages, in the left and righthand subfigures, respectively.

Figure A25: Heterogeneity by local wages by demographic groups

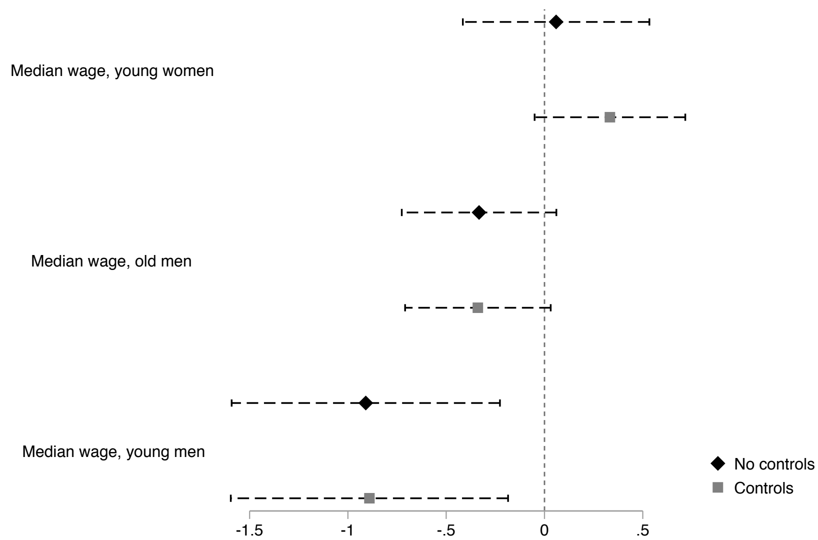

Note: This figure shows village-level triple-difference estimates and 95% confidence intervals of the effect of amnesty on oil theft (sabotage spills) across dimensions of labor market heterogeneity. The coefficient indicates the differential effect of amnesty on oil theft for a one-unit change in the heterogeneous variable. Labor market factors are the median nominal wages, in thousands of Nigerian Naira, of young men, young women (both ages 10-40), and old men (ages 40+) in village i's local government area in 2007. Each model is estimated with and without controls, as indicated. Controls are population density, distance to the nearest oil field, distance to the nearest oil pipeline, distance to the state capital, distance to the Niger River, and distance to the Atlantic coast.

Figure A26: Event-study heterogeneity by local alliance density

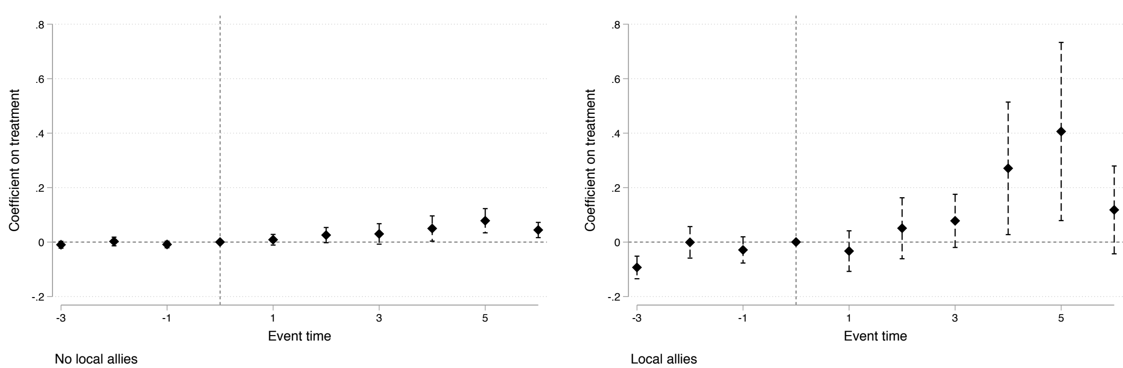

Note: This figure shows the coefficients and 95% confidence intervals from a dynamic differences-in-differences event-study regression of the outcome on dummies for years pre-and-post amnesty, interacted with the treatment indicator, which equals one for villages within 30 km of an amnestied militant camp (see equation 2). 2009 is the omitted reference group and the outcome variable is oil theft, measured as the total annual sabotage-related oil spills. The sample is split by villages for which the nearest militant camp has zero or greater than zero local allies within 10 kilometers along the pipeline, in the left and righthand subfigures, respectively.

Figure A27: Heterogeneity by alliance network

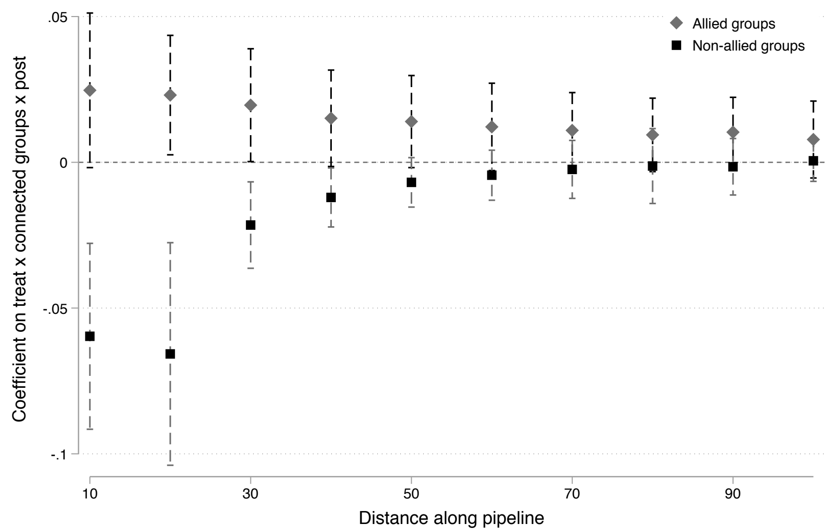

Note: This figure shows the differential effect of amnesty on sabotage spills by density of allied and non-allied groups for varying distances along the pipeline network. The figure plots estimates and 95% confidence intervals from triple difference regressions of the form in (**??**), in which the model is re-estimated for varying δ_i from 10 to 100 kilometers. Marker shape indicates whether the interaction is on the number of local allies or non-allies.

B Tables

Table B1: Summary statistics

Group	T (1)	C (2)	S-C (3)	P-C (4)	T-OP (5)	C-OP (6)	All (7)
Panel A: Pre-treatment outcome variables							
Sabotage spills	0.04	0.02	0.04	0.04	0.05	0.07	0.02
	(0.41)	(0.31)	(0.51)	(0.37)	(0.48)	(0.62)	(0.35)
Equipment failure	0.06	0.03	0.04	0.01	0.09	0.12	0.04
spills	(0.60)	(1.05)	(0.93)	(0.11)	(0.71)	(2.10)	(0.93)
Oil-related militant	0.06	0.01	0.02	0.06	0.09	0.02	0.03
attacks	(0.73)	(0.11)	(0.16)	(0.35)	(0.86)	(0.19)	(0.42)
Total conflict	0.21	0.06	0.07	0.42	0.29	0.13	0.11
events	(1.89)	(0.63)	(0.52)	(2.24)	(2.24)	(1.12)	(1.19)
Panel B: Cluster-level covariates							
Distance to oil	11.67	48.72	48.47	8.80	7.47	13.67	36.93
field (km)	(9.09)	(48.86)	(48.85)	(4.91)	(4.31)	(24.09)	(44.18)
Distance to pipeline	10.98	55.58	55.29	5.23	4.82	9.05	41.39
(km)	(12.83)	(55.70)	(55.69)	(4.19)	(4.18)	(14.97)	(50.99)
Distance to state	70.89	60.47	60.53	106.18	65.76	77.34	63.79
capital (km)	(49.82)	(46.63)	(46.53)	(40.81)	(50.63)	(35.64)	(47.91)
Distance to Niger	58.59	98.70	98.22	21.00	49.16	48.95	85.94
River (km)	(54.47)	(59.15)	(59.35)	(25.06)	(42.42)	(49.01)	(60.65)
Distance to coast	19.52	78.89	78.92	23.61	17.24	51.09	60.01
(km)	(21.18)	(49.33)	(49.22)	(13.30)	(20.02)	(35.97)	(50.66)
Population density	432.28	445.55	444.64	363.91	529.10	384.88	441.57
	(1116.37)	(754.01)	(751.95)	(489.21)	(1311.13)	(778.04)	(878.42)
Cassava suitability	9.94	34.94	34.80	13.98	9.58	27.41	26.99
	(11.88)	(16.20)	(16.27)	(14.15)	(11.90)	(16.63)	(18.96)
Mean annual	26.83	26.72	26.72	26.86	26.80	26.86	26.75
temperature (deg. C)	(0.17)	(0.38)	(0.38)	(0.11)	(0.12)	(0.23)	(0.33)
Total annual	2358.31	2177.62	2178.13	2439.50	2400.29	2297.00	2235.09
rainfall (mm)	(208.41)	(310.55)	(309.88)	(36.26)	(92.41)	(204.37)	(294.38)
Altitude (m)	13.78	89.21	88.86	13.22	12.28	40.12	65.22
	(8.14)	(93.77)	(93.65)	(3.27)	(6.69)	(61.95)	(85.15)
Slope	95.53	90.68	90.70	95.41	95.72	94.90	92.22
	(1.84)	(7.93)	(7.92)	(0.93)	(1.54)	(3.98)	(7.01)
Surveillance	0.51	0.36	0.36	0.84	0.58	0.73	0.41
contract recipient	(0.50)	(0.48)	(0.48)	(0.36)	(0.49)	(0.44)	(0.49)
Number of clusters	3692	7915	7915	456	2613	1958	11607

Standard deviations in parentheses. T stands for treatment and C for control. "S-C" is a synthetic-DD weighted control group, as in Arkhangelsky et al. (2019). "P-C" refers to placebo control group of untreated areas within 30 kilometers of a non-amnestied militant camp. "OP" indicates the sample is restricted to oil producing communities.

Table B2: The effect of Amnesty on gang activity

Polynomial	Linear		Quadratic		Quartic		AR(1)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Full sample								
Post-amnesty	-2.618***	-2.573***	-2.674**	-2.525***	-2.944**	-1.869*	-2.315**	-1.654**
	(0.753)	(0.606)	(1.080)	(0.809)	(1.131)	(1.028)	(0.914)	(0.651)
m_{t-1}							0.098	0.345***
							(0.119)	(0.109)
Month FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	146	132	146	132	132	132	145	131
R^2	0.431	0.350	0.459	0.364	0.542	0.458	0.436	0.423
$\Delta = 30$								
Post-amnesty	-3.174***	-3.821***	-2.913**	-2.498**	-1.593	-1.657	-3.205***	-2.280**
	(0.848)	(0.724)	(1.285)	(1.220)	(1.285)	(1.446)	(1.036)	(0.904)
m_{t-1}							-0.010	0.394**
							(0.162)	(0.169)
Month FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	61	61	61	61	61	61	61	61
R^2	0.509	0.446	0.510	0.508	0.596	0.562	0.509	0.529
$\Delta = 20$								
Post-amnesty	-3.206***	-3.329***	-1.639	-2.987**	-0.815	-0.699	-3.365**	-1.806
	(1.037)	(0.964)	(1.609)	(1.360)	(1.065)	(1.742)	(1.259)	(1.282)
m_{t-1}							-0.048	0.407*
							(0.185)	(0.219)
Month FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	41	41	41	41	41	41	41	41
R^2	0.513	0.560	0.530	0.616	0.680	0.830	0.515	0.632
$\Delta = 10$								
Post-amnesty	-2.535		-2.185		1.113		-4.283***	
	(1.541)		(1.778)		(1.985)		(1.087)	
m_{t-1}							-0.634***	
							(0.180)	
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	21	21	21	21	21	21	21	21
R^2	0.586	0.000	0.589	0.000	0.746	0.000	0.752	0.000

Robust standard errors. Outcome variable is the number of monthly gang-related events. Treatment is defined as an indicator for after July 2009. Window refers to the number of months included in the estimation before and after the event date. All windows apart from the full sample are symmetric. AR(1) specifications include a lagged dependent variable and a linear polynomial of event time. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B3: The effect of Amnesty on non-oil conflict

Polynomial	Linear		Quadratic		Quartic		AR(1)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Full sample								
Post-amnesty	2.807	5.825	-0.371	-9.643	19.387*	34.043***	1.173	3.071
	(4.269)	(5.492)	(5.745)	(9.116)	(10.013)	(12.343)	(4.931)	(4.705)
m_{t-1}							0.376***	0.637***
							(0.088)	(0.086)
Month FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	240	240	240	240	240	240	239	239
R^2	0.898	0.827	0.899	0.840	0.915	0.885	0.915	0.896
$\Delta = 30$								
Post-amnesty	7.146	16.940*	15.681	21.263	-2.156	-7.637	8.151	10.209
	(5.588)	(8.890)	(10.357)	(14.488)	(14.101)	(15.661)	(6.050)	(9.241)
m_{t-1}							-0.100	0.313*
							(0.188)	(0.163)
Month FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	61	61	61	61	61	61	61	61
R^2	0.677	0.537	0.725	0.539	0.749	0.676	0.681	0.576
$\Delta = 20$								
Post-amnesty	11.441	22.690*	2.523	-5.224	-16.109	-42.449*	14.068**	11.073
	(7.213)	(12.809)	(14.561)	(12.810)	(10.938)	(23.785)	(6.722)	(13.246)
m_{t-1}							-0.191	0.482**
							(0.167)	(0.192)
Month FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	41	41	41	41	41	41	41	41
R^2	0.768	0.506	0.774	0.620	0.810	0.700	0.782	0.615
$\Delta = 10$								
Post-amnesty	1.815		-20.869**		-7.328		-3.959	
	(13.504)		(7.696)		(31.138)		(8.904)	
m_{t-1}							-0.461***	
							(0.111)	
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	21	21	21	21	21	21	21	21
R^2	0.848	0.000	0.891	0.000	0.901	0.000	0.910	0.000

Robust standard errors. Outcome variable is the number of monthly non-oil-related conflict events. Treatment is defined as an indicator for after July 2009. Window refers to the number of months included in the estimation before and after the event date. All windows apart from the full sample are symmetric. AR(1) specifications include a lagged dependent variable and a linear polynomial of event time. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B4: Structural break tests

Bandwidth	Break month	Wald statistic	p-value
20	1	9.12	0.040
25	18	15.29	0.002
30	0	15.38	0.002
35	0	16.19	0.001
40	0	20.91	0.000
45	0	17.77	0.001

Table gives results for a structural break test of oil militancy. The null hypothesis for the test is that the constant term, $\hat{\alpha} = E[y_t|y_{t-1}] - \hat{\rho}y_{t-1}$, does not vary over time. The estimated break month refers to the month with the highest Wald statistic for this test, measured in event-time. The test-statistic is the supremum of these χ^2 values for each possible date, while the p-value refers to the rejection probability associated with a given test statistic. We vary the event-window around the amnesty date from 20 to 45 months.

Table B5: Optimal lag selection for an AR(p) process

Criterion	LR	FPE	AIC	HQIC	SBIC
	(1)	(2)	(3)	(4)	(5)
Full sample					
Post-amnesty	-2.117***	-2.117***	-2.117***	-2.035***	-2.035***
	(0.612)	(0.612)	(0.612)	(0.596)	(0.596)
p	9	9	9	5	5
Observations	231	231	231	235	235
R^2	0.448	0.448	0.448	0.419	0.419
$\Delta = 30$					
Post-amnesty	-3.149**	-3.149**	-3.149**	-3.034**	-3.034**
	(1.363)	(1.363)	(1.363)	(1.502)	(1.502)
p	5	5	5	1	1
Observations	61	61	61	61	61
R^2	0.455	0.455	0.455	0.339	0.339
$\Delta = 20$					
Post-amnesty	-5.580***	-5.791***	-5.791***	-5.791***	-5.580***
	(1.680)	(1.922)	(1.922)	(1.922)	(1.680)
p	0	1	1	1	0
Observations	41	41	41	41	41
R^2	0.330	0.332	0.332	0.332	0.330
$\Delta = 10$					
Post-amnesty	-5.715**	-5.839*	-5.839*	-5.715**	-5.715**
	(2.544)	(3.042)	(3.042)	(2.544)	(2.544)
p	0	1	1	0	0
Observations	21	21	21	21	21
R^2	0.363	0.363	0.363	0.363	0.363

Robust standard errors. Outcome variable is the number of monthly oil-related militancy events. Treatment is defined as an indicator for after July 2009. Window refers to the number of months included in the estimation before and after the event date. All windows apart from the full sample are symmetric. p lags are included in each specification, with p indicated in the table footer. Each AR(p) specification includes a linear polynomial of event time. Optimal lags are selected by the likelihood ratio test of p vs. $p-1$ (1), or minimizing the final prediction error (2), the Akaike information criterion (3), the Hannan-Quinn information criterion (4), and Schwarz's Bayesian information criterion (5). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B6: The effect of Buhari's announcement on militancy

Polynomial	Linear		Quadratic		Quartic		AR(1)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Full sample								
Post-announcement	4.708**	3.293***	5.739***	7.258***	7.131***	10.485***	5.359**	2.098*
	(1.963)	(1.136)	(1.938)	(1.562)	(2.340)	(1.815)	(2.076)	(1.070)
m_{t-1}							-0.110	0.321***
							(0.110)	(0.092)
Month FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	276	276	276	276	276	276	275	275
R^2	0.474	0.212	0.481	0.301	0.494	0.341	0.479	0.291
$\Delta = 30$								
Post-announcement	4.894**	4.770***	5.417**	5.929***	6.998**	7.948***	6.468**	3.659**
	(2.007)	(1.470)	(2.172)	(1.985)	(3.418)	(2.060)	(2.670)	(1.672)
m_{t-1}							-0.265	0.197
							(0.253)	(0.191)
Month FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	61	61	61	61	61	61	61	61
R^2	0.550	0.521	0.583	0.663	0.605	0.698	0.575	0.538
$\Delta = 20$								
Post-announcement	4.908**	5.044**	6.453*	6.026***	8.164*	7.408*	6.628**	5.575**
	(2.033)	(1.817)	(3.202)	(2.148)	(3.998)	(3.632)	(2.953)	(2.584)
m_{t-1}							-0.288	-0.082
							(0.307)	(0.230)
Month FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	41	41	41	41	41	41	41	41
R^2	0.574	0.647	0.608	0.711	0.625	0.718	0.601	0.649
$\Delta = 10$								
Post-announcement	7.249**		6.946		9.975		9.236**	
	(3.073)		(3.930)		(5.612)		(3.196)	
m_{t-1}							-0.317	
							(0.372)	
Year FE	Yes	No	Yes	No	Yes	No	Yes	No
Observations	21	21	21	21	21	21	21	21
R^2	0.512	0.000	0.528	0.000	0.567	0.000	0.551	0.000

Robust standard errors. Outcome variable is the number of monthly oil-related militant attacks. Treatment is defined as an indicator for after the May 2016 announcement of funding cuts for the Niger Delta amnesty. Window refers to the number of months included in the estimation before and after the event date. All windows apart from the full sample are symmetric. AR(1) specifications include a lagged dependent variable and a linear polynomial of event time. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B7: The effect of amnesty on oil bunkering: control groups

Control group	All untreated (1)	Synth-DD (2)	Pipeline dist<30km (3)	Core ND (4)	Oil-producing (5)	Placebo (6)
Within 30 km × Post-amnesty	0.069***	0.070***	0.061**	0.059**	0.067**	0.097***
	(0.021)	(0.023)	(0.025)	(0.029)	(0.033)	(0.024)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Village FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	139572	139572	78444	58656	54888	49992
R^2	0.416	0.361	0.421	0.423	0.421	0.411

Standard errors clustered at the village level. Outcome variable is the number of annual oil bunkering events. Treatment is defined as all villages within 30 km of an amnestied militant camp. Column (1) includes full sample of control observations. Column (2) reweights the control sample with synthetic DD weights (Arkhangelsky et al., 2019). Column (3) uses a sample of villages located < 30 km from a pipeline. Column (4) uses the villages in only the core Niger Delta states (Delta, Bayelsa, and Rivers). Column (5) uses only oil-producing villages. Column (6) uses all treated observations and defines the control group as all militant-affected villages where the amnesty was not implemented. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B8: The effect of amnesty on malfunctions: specifications

Dependent variable	Equipment failure and malfunction spills							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Within 30 km × Post-amnesty	-0.016	-0.005	-0.005	-0.014	-0.020**	-0.020**	-0.000	0.005
	(0.010)	(0.015)	(0.016)	(0.010)	(0.008)	(0.008)	(0.009)	(0.012)
Distance to oil field (00s km) × Post-amnesty		0.031**	0.031**					
		(0.016)	(0.013)					
Distance to pipeline (00s km) × Post-amnesty			0.001					
			(0.010)					
Within 30 km × Oil price (USD/barrel)				0.000**		-0.000		
				(0.000)		(0.000)		
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Year FE	No	No	No	No	Yes	Yes	No	No
Village FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Municipality FE × year	No	No	No	No	No	No	Yes	No
Village FE × year	No	No	No	No	No	No	No	Yes
Observations	139572	139572	139572	139572	135372	135372	139572	139572
R^2	0.506	0.506	0.506	0.506	0.331	0.331	0.527	0.664

Standard errors clustered at the village level. Outcome variable is the number of annual oil spills due to equipment failure and malfunction and other non-sabotage causes. Treatment is defined as all villages within 30 km of an amnestied militant camp. Control group is all villages outside of this range. Controls include distance to nearest oilfield, distance to nearest pipeline, distance to state capital, distance to Niger River, distance to the coast, and population density. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B9: The effect of amnesty on sabotage: robustness to oil production

Dependent variable	Sabotage		Malfunction		Normalized sabotage
	(1)	(2)	(3)	(4)	(5)
Within 30 × Post-Amnesty	0.175***	0.173***	0.039	0.048	0.023***
	(0.059)	(0.057)	(0.059)	(0.060)	(0.005)
Oil flow (kbbl/day)		-0.000		0.001***	
		(0.000)		(0.000)	
Pipeline-segment FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	21256	21256	21256	21256	21256
R^2	0.421	0.421	0.402	0.402	0.326

Standard errors clustered at the pipeline-segment level. Outcome variable is given in the table header. Treatment is defined as all pipeline-segments within 30 km of an amnestied militant camp. Control group is all pipeline segments outside of this range. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B10: The effect of amnesty on oil bunkering, maximum likelihood estimation

Dependent variable	Oil bunkering					
Control group	Full sample			Core ND	OP	Placebo
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Fixed-effects Poisson regression						
Within 30 km × Post-amnesty	0.529**	0.531**	0.461**	0.490***	0.416**	1.447***
	(0.215)	(0.217)	(0.182)	(0.183)	(0.182)	(0.317)
Within 30 km × Oil price (USD/barrel)		0.001	0.004	0.005	0.005	0.028***
		(0.005)	(0.004)	(0.004)	(0.004)	(0.007)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	No	Yes	Yes	Yes	Yes
Observations	10056	10056	10056	8712	9372	6180
Log-likelihood	-7371.667	-7371.322	-7190.343	-6436.517	-6895.349	-4403.466
Panel B: Zero-inflated Poisson regression						
Within 30 km × Post-amnesty	0.690***	0.691***	0.806***	0.700***	0.793***	1.557***
	(0.224)	(0.225)	(0.230)	(0.216)	(0.237)	(0.406)
Within 30 km × Oil price (USD/barrel)		0.003	0.006	0.005	0.006	0.021**
		(0.005)	(0.005)	(0.005)	(0.005)	(0.008)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	No	Yes	Yes	Yes	Yes
Observations	139284	139284	139284	58656	54852	49776
Log-likelihood	-1.72e+04	-1.72e+04	-1.70e+04	-1.40e+04	-1.56e+04	-9.69e+03

Standard errors clustered at the village level. Outcome variable is the number of annual oil spills due to sabotage. Treatment is defined as all villages within 30 km of an amnestied militant camp. Control group is all villages outside of this range. Top panel uses a fixed-effects poisson regression for estimation. Bottom panel uses a zero-inflated poisson regression, with a first-stage logit that models selection into oil theft as a function of distance to nearest oilfield and pipeline. Controls include distance to nearest oilfield, distance to nearest pipeline, distance to state capital, distance to Niger River, distance to the coast, and population density. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B11: The effect of amnesty on oil bunkering by median wages

Dependent variable	Oil bunkering					
	(1)	(2)	(3)	(4)	(5)	(6)
Amnestied × Post-amnesty	0.177***	0.165***	0.121***	0.093**	0.038	0.034
	(0.066)	(0.058)	(0.046)	(0.038)	(0.030)	(0.026)
Amnestied × Post-amnesty × Median wage, young men	-0.890**	-0.909***				
	(0.360)	(0.348)				
Amnestied × Post-amnesty × Median wage, old men			-0.339*	-0.333*		
			(0.189)	(0.201)		
Amnestied × Post-amnesty × Median wage, young women					0.333*	0.059
					(0.195)	(0.242)
Village 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	138468	134304	138468	134304	138468	134304
R^2	0.416	0.418	0.416	0.418	0.416	0.418

Standard errors clustered at the village level. Outcome variable is the number of annual oil spills due to sabotage. Treatment is defined as all villages within 30 km of an amnestied militant camp. Control group is all villages outside of this range. Controls include distance to nearest oilfield, distance to nearest pipeline, distance to state capital, distance to Niger River, distance to the coast, and population density. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table B12: Price responsiveness by amnesty status and period, triple difference

Dependent variable	Oil bunkering			
	(1)	(2)	(3)	(4)
Within 30 km × Oil price (USD/barrel)	-0.000	-0.000	-0.000	-0.000
	(0.001)	(0.001)	(0.001)	(0.001)
Within 30 km × Oil price (USD/barrel) × Post-amnesty	0.002**	0.001	0.003***	0.002**
	(0.001)	(0.001)	(0.001)	(0.001)
Field FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Controls × Year FE	No	Yes	No	Yes
Village linear time trends	No	No	Yes	Yes
Observations	139284	135120	139284	135120
R^2	0.416	0.418	0.520	0.522

Standard errors clustered at the village level. Outcome variable is the number of annual oil bunkering events. Treatment is defined as all villages within 30 km of an amnestied militant camp. Control group is all villages outside of this range. Oil price is the WTI crude price in dollars per barrel. Controls include distance to nearest oilfield, distance to nearest pipeline, distance to state capital, distance to Niger River, distance to the coast, and population density. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C Data and measurement

C.1 Data sources

Militant events data: Data on oil-related conflict events comes from several sources. We use the Armed Conflict Location Event Dataset (ACLED) to measure high-profile attacks on the oil sector perpetrated by militant groups. This primarily consists of bombings of major oil infrastructure, kidnappings and killings of oil workers, and battles against the Nigerian military. To identify all such events from 1997-2017, we conduct keyword searches,[25] and further subset events containing the keywords to include only those attacks that were perpetrated by political militias or rebel groups.[26] We also include all additional events perpetrated by MEND, the largest Niger Delta militant group, that may not include the specified strings.[27] This yields at total of 419 events, or 4.2% of Nigeria's total conflict events over the period.

Oil bunkering data: Information on the time, location, and details of 11,327 georeferenced oil spills comes from the administrative records of the the National Oil Spill Detection and Response Agency (NOSDRA), made publicly available on their Oil Spill Monitor. For every oil spill discovered in Nigeria, a Joint Investigative Report (JIV), is filed by the relevant oil companies, the Department of Petroleum Resources (DPR), NOSDRA, and the affected community. This JIV consists of a field visit to the spill site, in which the catchment area and size of the spill is assessed, as well as other basic data: the cause of the spill (e.g. sabotage or operational failure), a description of the environment (e.g. marsh, wetland, creek, or offshore), the site,[28] and of the incident itself.[29] 9,783 spills occur onshore, while 1,544 are offshore. Sabotage accounts for 61.7% of all spills, a proportion that rises over time (see Figure A1). Using these reports we generate a panel of quarterly and annual sabotage and malfunction spills with their characteristics at the village-level from 2006-2016[30] by designating affected villages as all those within 5 kilometers of a given spill. Controversy abounds over the extent of underreporting and misattribution of spill causes, either because of uneven government access to conflict zones or incentives to misreport on the part of oil companies (Amnesty International 2018). We test the robustness of our main results to measurement error correlated with the amnesty.

Militancy and amnesty data: Data on militant commanders was collected by the authors from several sources. In 2018, we visited Warri, Delta State, one of the epicenters of the Niger Delta crisis. We first collated a list of militant commanders from previous qualitative work

[25]These key words include oil, petroleum, drilling, rig, pipeline, flow, and the names of all major multinational and state-owned oil companies operating in Nigeria.

[26]This excludes communal conflict over oil revenues that is not directly related to militant groups.

[27]Results are not sensitive to this choice.

[28]For example "18 inch Assa-Rumuekpe Trans Niger Pipeline at Emohua."

[29]One reads "a 6 inch ball valve was installed on the line by suspected oil bunkers. A 4 inch galvanized pipe ran from the valve to water front where the bunkers load from."

[30]JIV reports from oil spills prior to 2006 are not available.

on the program, including Ugwu and Oben (2011) and Ojakorotu and Dodd Gilbert (2010), as well as consultation with AA Peaceworks (AAPW), a local non-profit organization with more than two decades of experience in conflict mitigation and peace-building in the region. For each militant commander, AAPW provided the following information: *i*) the group that this commander was affiliated with, *ii*) the location of their camp(s), usually denoted by the exact creek or a nearby village, and *iii*) whether they accepted amnesty. Gaps in the data and verification of accuracy were addressed by consulting Nigerian newspapers. We supplement this dataset with a list of pre-amnesty militant camps collected in a similar exercise by Blair and Imai (2013). Using Google searches of Nigerian newspapers we determine which militant commanders on our list received government contracts to perform security services in the oil sector.

This data collection procedure yielded 69 militant camps led by 41 unique commanders belonging to 24 different groups. Of the geo-referenced camps, 47 were controlled by commanders who accepted amnesty, while 22 were controlled by those who did not.[31]

Oil production data: Data on oil production comes from two sources: the administrative records of the Department of Petroleum Resources (DPR), and the Monthly Petroleum Information (MPI) of the Nigerian National Petroleum Corporation (NNPC), both made available to the author by these respective government agencies. These data contain field-level monthly production quantities for all oil fields in Nigeria from 1999-2015, excluding 2009, where data is unavailable.[32] The data were digitized from PDF files, resulting in an unbalanced panel 3,773 field-year observations covering 319 unique fields. We consider "birth" year of a field to be the first year in which it enters either of the government records. Since we are interested in capturing both extensive and intensive margins of field production, we include a field for all months after it is born, even if it does not appear in the data in a given month.[33] We then match each of the oilfields and their production history to detailed geographic data on oil and gas infrastructure as described below.

Oil and gas infrastructure: Data on the geographic location of oil and gas infrastructure comes from the DPR and Google Maps. Taken together, out of 337 ever-producing oil fields in the production data, we are able to georeference 321, or 95.3%. The data also feature a pipeline network of 4,284 kilometers. This latter number underestimates the official length of pipeline network of roughly 5,000 km, per NNPC, by 15%. This is expected, as this data primarily measure the larger flow and trunklines that connect fields to other fields and ex-

[31] Determining how to define amnesty take-up required some judgment calls. For example, Eziz Ogunboss, a MEND commander, formally accepted amnesty but was rumored to remain involved in the conflict in various clandestine ways. We code such cases as accepting amnesty. In contrast, John Togo of the Niger Delta Liberation Front (NDLF) took amnesty, and then publicly renounced it shortly thereafter and returned to the creeks. In this case, we consider him as not having accepted amnesty, since he did not actually participate in the program.

[32] Where possible, we also crosscheck this data with aggregated monthly totals provided in the Annual Reports of the Nigerian Extractive Industries Transparency Initiative (NEITI).

[33] In such cases, I assume the field did not produce anything in that month.

port terminals, and are unlikely to capture smaller delivery lines that flow within fields from wellheads themselves. A map of the oil infrastructure, along with the locations of militant camps and their amnesty status, can be found in Figure A2. We use infrastructure data in order to identify oil-producing villages by both their distance to the nearest field and/or pipeline, allowing us to restrict the sample to only oil-producing communities where appropriate. This data also allows us to control flexibly for differential trends in outcomes for oil-producing and/or pipeline-rich areas around the time of the amnesty.

C.2 Measurement

To test the comparative statics of the model, we measure the key equilibrium quantities and exogenous variables: oil theft, violence, extent of amnesty, and rebel capacity to inflict damage.

Bunkering: Throughout the paper, we measure oil bunkering as all spills that occurred as an act of sabotage. To allay concerns that this definition is overly broad – capturing non-bunkering acts of sabotage such as political protest – we categorize sabotage spills by infrastructure type and confirm that only economically significant pipelines are affected by Amnesty. However, several sources of measurement error in spills may contaminate our empirical analysis. Firstly, the manner in which these causes are recorded is a matter of substantial controversy in Nigeria. Since oil companies are not liable for compensation in the event of oil spills caused by third party tampering, the designation of spills in the JIVs is often hotly contested between communities, environmental rights groups, oil companies, and regulators. In 2018, Amnesty International, a rights group, investigated 1,830 JIVs filed by Agip and Shell, two of the major oil companies in the region, and found 89 were misclassified as sabotage.[34] While this number is lamentably nonzero, it represents a relatively modest 4.8% share of the total and is thus unlikely to affect the results.

More importantly, spills could be under-reported in general, arising from difficulty accessing militant-affected host communities, or limited capacity and funding constraints of Nigerian regulators. To affect the results, it must be that the physical ability to collect data changed at the time of the Amnesty in amnestied areas. Prior to the Amnesty, it is possible that security concerns prevented regulators and companies from accessing spill sites in militant-controlled communities. Since such measurement error would be differentially worse in treated areas, we use operational failure spills – affected by the same measurement error bias but not the Amnesty itself – as a placebo check to provide evidence that this is unlikely to be driving the results.

Amnesty: For each camp, mapped in Figure A2, we calculate the distance between that

[34] www.amnesty.org/en/latest/news/2018/03/niger-delta-oil-spills-decoders/

camp and all 11,607 villages in the sample.[35] All villages that fall within a 30 kilometer radius[36] of any amnestied camp are considered affected by amnesty for the purposes of the difference-in-differences analysis. This is based on the assumption, validated by review of the qualitative literature (Asuni 2009) and interviews with AA Peaceworks, that a given militant controls oil bunkering primarily in territories immediately surrounding his camp. We test robustness of all results to numerous different definitions of the non-amnestied control group.

Differentiated amnesty is a critical aspect of the theoretical and empirical setting. To measure this, we define three outcomes at the militant camp level: *i*) if the camp received any amnesty at all, *ii*) if the camp received a "generous" amnesty, which includes pipeline security contracts, and *iii*) if the camp was defeated in battle. This allows us to correlate differential amnesty outcomes with the strength of militants to test key predictions of the model.

Militant strength: The extent of output damage that a militant group is able to inflict – a key parameter in the model – is unobserved. To proxy for the underlying strength of a militant camp, we begin with the assumption that militant commanders coordinate with nearby allied groups in order to carry out attacks on oil infrastructure. We therefore measure military strength at the camp-level by the number of allied camps within a given area. To construct our measure, we leverage two networks. The first is the network of alliances between groups at the time of the amnesty. The Niger Delta conflict is known for a complex web of alliances between militant commanders, most prominently in the form of the umbrella group MEND. Let $\mathbf{\Pi}$ be the alliance network matrix, where $\pi_{ij} = 1$ if camps i and j belong to the same larger umbrella group (e.g. Movement for the Emancipation of the Niger Delta, Niger Delta People's Volunteer Force, et cetera) or if they are identified as allies in a comprehensive 2007 Small Arms Survey review of the state of the Niger Delta conflict (Hazen and Horner, 2007).[37] Since rebel strength is inherently localized by its dependence on access to physical infrastructure targets, we overlay the alliance structure with the network of physical oil pipeline infrastructure, $\boldsymbol{P}^d$, where $p^d_{ij} = 1$ if i and j are located on connected pipelines and $d(i, j) \leq d$, where $d(\cdot)$ is the distance function that measures the shortest path between two points along the pipeline network. This gives us our measure of strength, $k^d_i = \sum_{j \in Z} \pi_{ij}$ where $Z = \{j \neq i | P^d_{ij} = 1\}$, or the number of allies for camp i within d kilometers along the pipeline. In Section 6.1, we validate this measure by showing its correlation with output

[35]The village sample is defined as all villages in the Niger Delta geopolitical region, comprising the states of Akwa Ibom, Abia, Bayelsa, Cross River, Delta, Edo, Ondo, and Rivers. Where appropriate, we restrict the sample to the "core" onshore oil-producing states of Delta, Bayelsa, and Rivers.

[36]We consider robustness of results to ten-kilometers intervals from 10 to 100.

[37]We take no position on the strength of connections, since some within-group alliances may well be weaker than cross-group alliances, given substantial variance in organizational capabilities of umbrella groups (Asuni 2009).

losses.

FIGURE 2 HERE

The network of alliances and pipelines is visualized on Figure 2. This map plots the physical pipeline infrastructure, with each militant camp attached to its nearest segment of pipeline. Red triangles indicate oil export terminals, the geographic endpoints of a pipeline. Blue triangles indicate the location of militant camps along pipelines. The map also overlays the geographic distribution of k_i^d for $d = 10$ kilometers. The size of the militant camps markers correspond to a number of allies within 10 kilometers. Variation in k_i^d is driven by variation in the location of camp i, the location of camps $-i$, and the number and type of alliances camp i possesses. High k_i^{10} camps are concentrated along the eastern pipelines connecting Port Harcourt to Soku field in Rivers state, and the southern pipelines connecting Ogbainbiri to the coast in Bayelsa state. Finally, we measure damage inflicted at the camp-level as the percent change in oil production within 20 km of that camp location between 2005 and 2009, the peak years of the conflict.[38] The following quantities are measured at the village or village-year level: attacks, bunkering, distance to amnestied camps, and geographic control variables relating to access to oil infrastructure. In contrast, amnesty generosity, group power, and damage inflicted are measured at the camp-level. Lastly, outcomes are also aggregated to the month-level for our event-study models.

[38]We test robustness to this choice in Section 6.1.

D Theoretical framework

D.1 Fringe bunkering, motivation and microfoundation

Based on the reports discussed in Section 2.2 we assume that the bunkering has a dominant firm, competitive fringe structure. We derive a micro-foundation of γ, our parameter of competitive oil theft. There is no strategic interaction between government and fringe. Let P the price of crude oil and Q output, both exogenously fixed. A rebel chooses oil theft $\tilde{q} \in \{0, q\}$, where $0 < q < Q$. The corresponding pay-offs are $\{0, \rho\}$. Let $\rho < \mu$, so the rebel income from stealing oil is strictly less than the market value of that oil. The representative fringe maximizes profits, stealing $z \in \mathbb{R}_+$ at quadratic cost $cz^2(2[Q - \tilde{q}])^{-1}$, capturing the notion of decreasing returns to scale. The first-order condition for profit maximization yields equilibrium quantities

$$z = \frac{p}{c}(Q - \tilde{q}) = \arg\max Pz - \frac{cz^2}{2(Q - \tilde{q})}$$

so fringe theft increases linearly in residual output $z = {}^{p}/_{c}(Q - \tilde{q})$ by a factor $\gamma := {}^{p}/_{c}$ in residual quantities. A reasonable and tractable assumption is that theft occurs irrespective of conflict. Thus, government income is given by $\pi = (1 - \gamma)pQ$ when the rebel does not steal and $\pi - (1 - \gamma)\mu$ when there is theft.

D.2 Proposition 1: Minimal amnesty equilibrium.

Proof. We verify that this strategy profile constitutes a Nash equilibrium for any history and then argue that it is unique.

Begin with the rebel's one-shot deviations. Consider first offers in $x \in (-\infty, m)$, which together with income from oil theft ρ yield a payoff less than r. For these offers the rebel prefers to reject and fight. Next consider $x \in [m, \rho/(1 - \psi)]$. The rebel strictly prefers to steal after having accepted and is at most indifferent between fighting and the equilibrium action. Suppose the rebel is offered $x \in [\rho/(1 - \psi), \infty) > 0$. If stealing the government's best response to oil theft is to cease payments and perhaps to fight. The payoff from stealing is a most $(1 - \delta)(\psi x + \rho) + \delta x$ and no greater than x if $x \geq \rho/(1 - \psi)$. For any offer x the rebels strategy is a best response.

Now consider the government's one-shot deviations. An offer $x < m$ is rejected by rebels and not preferred if $m + (1 - \gamma)\mu = r + \ell < g$ where r and g are the rebel- and government payoffs from fighting (3) and (4). But $g > r$ since $\kappa > \alpha$ and $\ell \leq 0$ by $\gamma \geq \gamma'$ so offers $x < m$ are not preferred to m. The government prefers an offer $\psi/(1 - \gamma)$ to $x > \psi/(1 - \gamma)$ but does not prefer $\psi/(1 - \gamma)$ to m: We have that $\pi - \rho/(1 - \psi) < \pi - \ell - r$ implies $r + \ell < \rho/(1 - \psi)$. But again $\ell < 0$ and $r < \alpha \leq \rho \leq \rho/(1 - \psi))$. Once the offer m has been accepted ceasing the transfer is never a best response because the equilibrium transfers surplus to the government, $m < 0$ by $\alpha < \rho$. Finally, as $\pi - m > \pi - \psi m - g$ for all ψ ceasing payments and fighting is not preferred.

We now argue that the proposed equilibrium is essentially unique. Facing an offer m, accepting and stealing is a weak best response for the rebel. By subgame perfection we may exclude all strategies that prescribe any other action. Next, notice that the stated strategy profile maximizes government payoffs: Without efficiency loss $\gamma : \ell < 0$ the net available surplus is maximized by allowing oil theft, $\pi - \ell > \pi$. The government then extracts all but the rebel reservation value r from this surplus. But then for any other strategy profile the government always weakly prefers deviating to an offer of m and thus it cannot be a subgame perfect equilibrium. □

D.3 Proposition 2: Low-powered equilibrium.

Proof. We verify that the proposed strategy constitutes a Nash equilibrium for any history and then argue that it is unique.

Begin with the rebel's one-shot deviations. Rebels prefer to reject and fight if $(1-\delta)\alpha + \delta\theta c^\ell > c^\ell$, implying $c^\ell < r$, which does not hold as c^ℓ is bounded below by $\rho \geq \alpha > r$. If the rebel accepts c^ℓ and then steals their payoff is $(1-\delta)(\psi c^\ell + \rho) + (1-\delta)c^\ell = c^\ell$, and thus rebels weakly prefer the equilibrium action of not stealing.

Now consider the government's one-shot deviations, beginning with their decision to cease payments and not fight. Suppose they have made an offer $x > 0$ and that the rebels have stolen oil. Ceasing the transfer is trivially preferred to allowing oil theft. A one-shot deviation to payment cessation and fighting is not preferred if

$$\lim_{\psi \to 0}(1-\delta)(\pi - \kappa - \psi c^\ell) + \delta\theta(\pi - c^\ell) + \delta(1-\theta)\pi \leq (\pi - \psi c^\ell - (1-\gamma)\mu)(1-\delta) + \delta(\pi - c^\ell) \tag{15}$$

implying $\kappa \geq \rho\delta(1-\theta)/(1-\delta\theta) + (1-\gamma)\mu$ under full recovery as we have $\lim_{\psi\to 0} c^\ell = \rho$. The right-hand-side is decreasing in Δ so taking $\Delta \to 0$ yields the proposed bound $\kappa' = \sigma/(\sigma+\nu)\rho + (1-\gamma)\mu$. by L'Hôpital's rule.

We now examine government deviations in terms of offers. They prefer to offer $x < m$ and fight one period to offering c^ℓ if $\pi - c^\ell \leq (1-\delta)(\pi - \kappa) + \delta\theta(\pi - c^\ell) + \delta(1-\theta)\pi$, implying $\rho > (1-\delta)/(1-\delta\theta)\kappa$ under full recovery $\psi \to 0$, which does not hold by assumption. The government strictly prefers c^ℓ to all $x > c^\ell$ and m to any $x \in (m, c^\ell)$. A one-shot deviation to m from c^ℓ is preferred if $\pi - m - (1-\gamma)\mu \geq \pi - c^\ell$, or $c^\ell \leq r + (1-\gamma)\mu - \rho = r + \ell$, where $\ell > \rho$ is the efficiency loss of rebel oil theft. But this is not the case under full recovery as $\lim_{\psi\to 0} c^\ell = \rho < r + \ell$. Hence it is optimal for the government to offer c^ℓ for any history and for the rebel to accept it and not steal.

The argument above establishes that c^ℓ is the least-cost subgame perfect offer that provides incentives not to steal, and that the government strictly prefers providing these incentives to allowing theft. Any other strategy profile either allows theft or transfers unnecessary surplus to rebels, and are not subgame perfect as the government will weakly prefers deviating to an offer of c^ℓ. The subgame perfect equilibrium in low-powered incentives is therefore essentially unique under the stated conditions.

□

D.4 Proposition 3: High-powered equilibrium.

Proof. We verify that the proposed strategy constitutes a Nash equilibrium for any history and then argue that it is unique.

Begin with the rebel's one-shot deviations. The rebels response to offers $x \in (-\infty, r - \rho)$ is given in the proof of Proposition 1 in Appendix D.2. Accepting and stealing in response to offers $x \in (r - \rho, c^h)$ is a best-response if the government does not fight, with payoff $\lim_{\psi \to 1} x + \psi(1 - \delta)\rho > x$. We confirm below that fighting is not a best response from the government. To accept and steal is not preferred to accepting without theft when the offer $x \geq (1 - \delta)(\alpha + \psi x) + \delta\theta x$, implying that $x \geq (1 - \delta)/(1 - \delta\theta - (1 - \delta)\psi)\alpha$. But this is the lower bound on the high-powered offer and thus holds by construction. The payoff of from rejecting and fighting is inferior to accepting and stealing by $(1 - \delta)\psi x \geq 0$.

Now consider the government's one-shot deviations, beginning with their decision to fight in the case of oil theft. Suppose they have made an offer $x > 0$ and that the rebels have stolen oil. Ceasing the transfer is trivially preferred to allowing oil theft. Cessation of payments and fighting is weakly preferred to just ceasing the transfer for all offers x such that

$$(\pi - (1 - \gamma)\mu - \psi x)(1 - \delta) + \delta(\pi - x) \leq (1 - \delta)(\pi - \kappa - \psi x) + \delta\theta(\pi - x) + \delta(1 - \theta)\pi$$
$$\Rightarrow x \geq \frac{1 - \delta}{\delta(1 - \theta)}(\kappa - (1 - \gamma)\mu)$$

where the lower bound, given that it exceeds the rebel recovery-adjusted reservation value, yields the proposed offer (8). Thus, given an offer of at least c^h, ceasing payments and fighting is always weakly preferred by the government to just halting the payment.

We now examine government deviations in terms of offers. An offer $x > c^h$ is not preferred to c^h. Consider next an offer $r - \rho \leq x < c^h$. By the argument above, fighting is not a best-response to oil theft for offers $x < c^h$. However, for a positive offer it is still optimal for the government to cease the transfer. If this threat is sufficient to deter rebel oil theft then this deviation is strictly preferred. Anticipating that payments will cease, the least value that gives rebels incentives not to steal solves

$$x : (1-\delta)(\psi x + \rho) + (1-\delta)x = x \implies x = \rho/(1-\psi)$$

which diverges to infinity when $\psi \to 1$. Hence, there does not exist an offer $x < c^h$ which provides rebels with an incentive not to steal. Consider next a one-time offer of $m \leq x < c^h$ that induces rebels to engage in oil theft. By Proposition 1 the payoff from this deviation is bounded above by offering minimal amnesty m and not ceasing payments, which extracts maximum surplus from rebels. The equilibrium offer c^h is weakly preferred to this deviation if

$$\pi - c^h \geq \pi - (1-\gamma)\mu - m = \pi - r - \ell$$

where we recall that $\ell = (1-\gamma)\mu - \rho$ the efficiency loss from rebel oil theft. If the government transfers reservation values $c^h = r$ the high-powered contract is strictly preferred as $\pi - r > \pi - r - \ell$. To evaluate the general case, insert the identities for interior high-powered- and limited amnesty payments c^h, m from (8), (5) and consolidate terms to yield

$$\begin{aligned} \pi - \frac{1-\delta}{\delta(1-\theta)}\left(\kappa - (1-\gamma)\mu\right) &\geq \pi - \frac{1-\delta}{1-\delta\theta}\alpha - \ell \\ \kappa &\leq \frac{\delta(1-\theta)}{1-\delta\theta}\alpha + (1-\gamma)\mu + \frac{\delta(1-\theta)}{1-\delta}\ell \end{aligned} \tag{16}$$

as proposed. Finally we verify that the equilibrium is essentially unique. The argument above establishes that c^h is the least-cost subgame perfect offer that provides incentives not to steal, and that this is preferred to allowing theft. Any other strategy profile either allows theft, violates subgame perfection, or transfers unnecessary surplus to rebels. Hence under any other strategy profile the government always weakly prefers deviating to an offer of c^h. The subgame perfect equilibrium in high-powered incentives is therefore essentially unique under the stated conditions. □

D.5 Existence of high-powered equilibrium for strong types.

Recall from Proposition 3 that when $\psi \to 1$ the deviation to a low-powered transfer c^ℓ in (6) is dominated by the high-powered transfer c^h in (8) and that the government weakly prefers c^h to the minimal transfer m in (5) when

$$\kappa \leq \kappa^* = \frac{\delta(1-\theta)}{1-\delta\theta}\alpha + (1-\gamma)\mu + \frac{\delta(1-\theta)}{1-\delta}\ell \tag{17}$$

as derived in (16). Inserting $\kappa = \overline{\kappa}$ from Assumption 1 and solving for α yields

$$\alpha(\gamma) \leq \frac{1-\delta\overline{\theta}}{1-\delta+\delta(\overline{\theta}-\underline{\theta})}\left((1-\gamma)\mu + \frac{\delta(1-\overline{\theta})}{1-\delta}\ell\right) \tag{18}$$

which is decreasing in γ. Recall from Section 5.1 that the upper bound $\gamma \leq (\mu-\rho)/\mu$ on γ implied net efficiency loss from rebel oil theft $\ell \geq 0$. Evaluating, we have that

$$\alpha\left(\frac{\mu-\rho}{\mu}\right) \leq \frac{1-\delta\overline{\theta}}{1-\delta+\delta(\overline{\theta}-\underline{\theta})}\rho \tag{19}$$

where the right-hand-side is bounded below by $\rho(1-\delta)/(1-\delta\underline{\theta})$, which holds by Assumption 2.

D.6 Proposition 4: Perfect Bayesian equilibrium with fighting and minimal amnesty.

Proof. We verify that for sufficiently optimistic government beliefs the strategy profile constitutes a perfect Bayesian equilibrium unique up to the value of rejected offers.

Notice first that rebels of both type play the full information strategy of the strong, see Proposition 1. It therefore follows directly that payoffs are bounded above by the reservation value of strong rebels. We show that there exist priors p such that weak rebels weakly prefer to mimic the strong and no separating offer exists. It then follows from the assumption $p > p^*(\overline{r}+\ell)$ that the government prefers to make pooling offers that are rejected by both types.

To begin we may exclude $(-\infty, \underline{r}-\rho)$ and $(\overline{r}-\rho, \infty)$ from the set of potential separating offers. Proposition 1 shows that in the unique subgame perfect equilibrium of the full information game the weak and strong rebels receive their reservation values $\underline{x} = \underline{r} < \overline{r} = \overline{x}$. We may therefore exclude $x > \overline{r}-\rho$ because the government strictly prefers $x = \overline{r}-\rho$ which is accepted by both rebel types. Next, it is without loss of generality to consider only $x \geq \underline{r}-\rho$ as all offers $x < \underline{r}-\rho$ are rejected with fighting by both rebel types and are therefore equivalent in terms of actions. An offer $x \in (\underline{r}-\rho, \overline{r}-\rho)$ is separating if the best response of the strong rebel is to reject it and fight but the weak optimally accepts it, rendering rebel type common knowledge after at most one period of conflict.

We now argue that there exist priors p such that no separating offers exist. The argument proceeds in two steps. We first derive the maximal number of periods $\omega : (0,1) \to \mathbb{Z}$ a government with prior p is willing to fight before preferring to make offer of $\overline{r}-\rho$, given beliefs. We then show that the maximal number of fighting periods $\tau \in \mathbb{Z}$ a weak rebel must expect to accept a separating offer exceeds ω when $\Delta \to 0$.

Recall that $p^*(\ell+\overline{r})$ in (9), hereafter p^*, is defined as the belief for which the government is indifferent between fighting one more period and offering $\overline{r}-\rho$ or making the offer $\overline{r}-\rho$ immediately. Fix a prior belief $p > p^*$ that the rebel is weak. We derive the length of time the government must fight the rebel before the posterior is weakly less than p^*, and the offer $\bar{r} - \rho$ is preferred to fighting. By Bayes' rule, at the beginning of period t a rebel that has survived $t-1$ periods of fighting is believed weak with probability

$$\Pr\left(\text{rebel is weak}|\text{rebel survived } t-1 \text{ fighting periods}\right) := p_t = \frac{\underline{\theta} p_{t-1}}{\overline{\theta}(1-p_{t-1}) + \underline{\theta} p_{t-1}} \tag{20}$$

for prior p_{t-1}. The posterior is recursively defined as

$$\frac{p_t}{1-p_t} = \left(\underline{\theta}/\overline{\theta}\right)\frac{p_{t-1}}{1-p_{t-1}} = \left(\underline{\theta}/\overline{\theta}\right)^2 p_{t-2}/1-p_{t-2} = \cdots = \left(\underline{\theta}/\overline{\theta}\right)^t \frac{p}{1-p} \tag{21}$$

so $p_t < p_{t-1}$. The required periods of fighting before the government holds posterior belief in q is therefore

$$\omega(p, p^*) := \min_{z \in \mathbb{Z}} z : \left(\underline{\theta}/\overline{\theta}\right)^z \frac{p}{1-p} \leq \frac{p^*}{1-p^*} = -\frac{\ln(p^*/1-p^*)\ln(1-p/p)}{\Delta(\underline{\sigma} - \overline{\sigma})} + c \tag{22}$$

where $c \in (0, 1)$ rounds ω up to the nearest nonnegative integer.

Let the time-averaged rebel pay-off from fighting ω periods and then receiving a net payoff of $\bar{r}$ be

$$v(\theta, \omega) := (1-\delta)\frac{1-(\delta\theta)^\omega}{1-\delta\theta}\alpha + (\delta\theta)^\omega \bar{r} = r(\theta)(1-(\delta\theta)^\omega) + (\delta\theta)^\omega \bar{r} \tag{23}$$

where we have replaced $r(\theta) = \alpha(1-\delta)/(1-\delta\theta)$ the rebel reservation value. Denote $v(\underline{\theta}, \omega) = \underline{v}$ and $v(\overline{\theta}, \omega) = \bar{r}$ the value to weak and strong types. A weak rebel accepts any offer

$$x : \underline{v} \leq (1-\delta)x + \delta\underline{x} \tag{24}$$

when anticipating that they will receive $\underline{r}$ forever after accepting. Solve (24) with equality

$$x = \underline{r} + \frac{(\delta\underline{\theta})^\omega}{1-\delta}(\bar{r} - \underline{r})$$

to obtain the efficient offer. We have that (24) separates if the strong have no incentive to accept it

$$x \leq \bar{r} \implies (\delta\underline{\theta})^\omega \leq (1-\delta) \Rightarrow \omega \geq -\frac{\ln(1-\exp(-\Delta\nu))}{\Delta(\nu + \underline{\sigma})} + c =: \tau \tag{25}$$

where $c \in (0, 1)$ is an implicitly defined constant rounding τ up to the nearest integer. Call τ the minimum separating threat. If $\tau > \omega \geq 1$ no separation by offers exists, weak rebels mimic the strong, and the government strictly prefers fighting ω periods to offering $\bar{r}$. Thus we have a unique pooling equilibrium with fighting when $\tau > \omega$ which implies

$$\frac{\ln(1-\exp(-\Delta\nu))}{\nu+\underline{\sigma}} < \frac{\ln(p^*/1-p^*)\ln(1-p/p)}{\underline{\sigma}+\overline{\sigma}} \tag{26}$$

where the parameter Δ has been canceled in the denominators. But since $\lim_{\Delta\to 0} p^* \in (0,1)$ we have that (26) holds for all $p > p^*$ when $\Delta \to 0$.

$\square$

D.7 Perfect Bayesian equilibrium with fighting and low-powered incentives

Recall from Proposition 2 that $c^\ell = \rho/(1-\psi)$ is the value of a low-powered incentive contract.

Proposition 6. ***Low-powered perfect Bayesian equilibrium****. Let $\kappa < \underline{\sigma}/(\nu+\underline{\sigma})\rho + (1-\gamma)\mu$ and suppose rebel income from oil theft is exceeded by the government's average cost of fighting $\rho < \kappa(1-\delta)/(1-\delta\theta)$ and efficiency loss from oil theft $\rho < \ell$. Consider the following strategy profile. Rebels of both types reject and fight for all $x \in (-\infty, r-\rho)$ accepts and steals for all $x \in [r-\rho, c^\ell)$ and accepts without theft for all $x \in [c^\ell, \infty)$. The government offers $\underline{r}-\rho$ for the first $\omega \in \mathbb{Z}_+$ periods and then $\rho/(1-\psi)$ forever. For all $p > p^*(\rho/(1-\psi))$ this strategy profile is an essentially unique perfect Bayesian equilibrium when $\Delta \to 0$ and $\psi \to 0$.*

Proof. We verify that the strategy profile constitutes a perfect Bayesian equilibrium unique up to the value of rejected offers. The stated parametric restrictions imply that the low-powered payoff in Proposition 2 is not granted to the weak type because the government prefers to provide them with high-powered incentives. The subgame perfect equilibrium payoffs therefore satisfy $\lim_{\psi\to 0}\underline{x} = \underline{r} < \overline{x} \geq \overline{r}$ under full recovery. Notice again that rebels of both type play the full information strategy of the strong, see Proposition 2. We show that there exist priors p such that no separating offer exists, the weak rebels have no profitable deviations, and that the government prefers to make pooling offers rejected by both types. The remaining argument now follows the proof of Proposition 4 in Appendix D.6 above, up to the functional form of the minimum separating threat τ.

Following the argument in Appendix D.6 above, it is without loss of generality to consider separating offers $x \in (\underline{r}-\rho, c^\ell)$ and verify that they do not exist when $\Delta \to 0$. The weak- and strong rebel's payoffs from the equilibrium strategy are $\underline{v} = \underline{r} + (\delta\underline{\theta})^\omega(\overline{x}-\underline{r})$ and $\overline{v} = \overline{r} + (\delta\overline{\theta})^\omega(\overline{x}-\overline{r})$. The efficient offer is thus

$$x = \frac{\underline{r} + (\delta\underline{\theta})^\omega(\overline{x}-\underline{r}) - \delta\underline{x}}{1-\delta}$$

which strong rebels reject in the full-recovery limit if $\lim_{\psi\to 0}\overline{v} \geq (1-\delta)x + \delta\overline{v}$ which by (22) implies that

$$\omega(p,p^*) > \tau := \min_{z\in\mathbb{Z}} (\underline{\theta}/\overline{\theta})^z \leq (1-\delta)\left(\frac{\overline{r}-\underline{r}}{(\overline{x}-\underline{r})(\delta\overline{\theta})^q} + \frac{\overline{x}-\underline{r}}{\overline{x}-\underline{r}}\right)$$

so we have an essentially unique pooling equilibrium if $\tau > \omega \geq 1$ and

$$\Delta\tau < \frac{\ln(p^*/1-p^*)\ln(1-p/p)}{\underline{\sigma}+\overline{\sigma}} \tag{27}$$

which holds when $\Delta \to 0$ as $\lim_{\Delta\to 0} p^* \in (0,1)$.

□

D.8 Proposition 5: Perfect Bayesian equilibrium with fighting and high-powered incentives

Proof. We verify that the strategy profile constitutes a perfect Bayesian equilibrium unique up to the value of rejected offers. We first verify that the payoffs are increasing in type $\underline{x} < x(p_\omega) = \overline{x}$. Note that by Assumption 1 and 2 the high-powered equilibrium in Proposition 3 is essentially unique under full information. The pooling transfer $x(p)$ is strictly greater the weak type's high-powered payoff in the limit if

$$\lim_{\psi\to 1} \frac{1-\delta}{\delta(1-\overline{\theta})}\alpha > \frac{(1-\delta)(\kappa-(1-\gamma)\mu)}{\delta(1-\underline{\theta})}$$

which implies $\alpha(1-\underline{\theta})/(1-\overline{\theta}) > \kappa - (1-\gamma)\mu$. But by Assumption 1 the left-hand-side is bounded below by κ, so the inequality holds and $c^h|(\theta=\underline{\theta}) = \underline{x} < x(p) \leq c^h|(\theta=\overline{\theta})$. Notice again that rebels of both type play the full information strategy of the strong, see Proposition 3. We show that there exist priors p such that no separating offer exists, the weak rebels have no profitable deviations, and that the government prefers to make pooling offers rejected by both types. The remaining argument now follows the proof of Proposition 4 in Appendix D.6 above, up to the functional form of the minimum separating threat τ.

Following the argument in Appendix D.6 above, it is without loss of generality to consider separating offers $x \in (\underline{r}-\rho, x(p_t))$ and verify that they do not exist when $\Delta \to 0$. The weak- and strong rebel's payoffs from the equilibrium strategy are $\underline{v} = \underline{r} + (\delta\underline{\theta})^\omega(\overline{x}-\underline{r})$ and $\overline{v} = \overline{r} + (\delta\overline{\theta})^\omega(\overline{x}-\overline{r})$. The efficient offer is thus

$$x = \frac{\underline{r} + (\delta\underline{\theta})^\omega(\overline{x}-\underline{r}) - \delta\underline{x}}{1-\delta}$$

which strong rebels reject in the no-recovery limit if $\lim_{\psi\to 0}\overline{v} \geq (1-\delta)x + \delta\overline{v}$. Recall that (12) defined the belief for which the government was indifferent between making a pooling offer and fighting. We can therefore evaluate (22) for $\omega(p,\hat{p})$ to compute the number of periods the government with prior p fights until making the pooling offer. Strong rebels reject the separating offer when

$$\omega(p,\hat{p}) > \min_{z\in\mathbb{Z}}(\underline{\theta}/\overline{\theta})^z \leq (1-\delta)\left(\frac{\overline{r}-\underline{r}}{(\overline{x}-\underline{r})(\delta\overline{\theta})^q}+\frac{\overline{x}-\underline{r}}{\overline{x}-\underline{r}}\right) =: \tau$$

so we have an essentially unique pooling equilibrium if $\tau > \omega \geq 1$ which by (22) implies

$$\Delta\tau < \frac{\ln(\hat{p}/1-\hat{p})\ln(1-p/p)}{\underline{\sigma}+\overline{\sigma}} \tag{28}$$

which holds when $\Delta \to 0$ as $\lim_{\Delta\to 0}\hat{p} \in (0,1)$.

□

D.9 A perfect Bayesian equilibrium with equilibrium-path fighting and separation.

Let nature draw the government prior $p \in \mathcal{P} := \{(\underline{p},\overline{p}) : 0 < \underline{p} < \overline{p} < 1\}$ with probability $\Pr(p = \underline{p}) = b \in \mathcal{B} := \{(b, 1-b) : b \in (0,1)\}$ commonly known. Because the number of periods the government is willing to fight increases in their prior belief p that the rebel is weak, call $\underline{p}$, $\overline{p}$ the wavering and determined government type. The time-t strategies are then of the form $g : \mathcal{H}^t \times \mathcal{P} \times \mathcal{B} \to \mathcal{G}$ and $r : \mathcal{H}^t \times \mathcal{T} \times \mathcal{P} \times \mathcal{B} \to \mathcal{R}$. The following proposition shows that there exists a PBE with government types such that weak rebels first fight and then accept limited amnesty.

Proposition 7. ***Perfect Bayesian equilibrium with fighting and separation*** *Suppose $\gamma : \ell \leq 0$ such that the minimal amnesty equilibrium from Proposition 1 pertains under full information with payoffs $\underline{r}$, $\overline{r}$ for the weak- and strong rebel. Let $1 \leq \underline{\omega} < \overline{\omega} < \infty$ two integer-values and $x(\omega) = \underline{r} + (\delta\underline{\theta})^\omega(\overline{r}-\underline{r})/(1-\delta)$. Consider the following strategy profile. A wavering government offers $x(\overline{\omega}-t)$ for every $0 \leq t \leq \underline{\omega}$ and then offers $\overline{x}$ for all $t > \underline{\omega}$. A determined governments offer $x(\overline{\omega}-t)$ for every $0 \leq t \leq \underline{\omega}+1$. If the final offer is accepted, offer $\underline{x}$ for all $t > \underline{\omega}+1$, if it is rejected, fight and then offer $\overline{x}$ forever. Rebels of both types reject and fight for all $x \in (-\infty, \overline{r}-\rho)$ accepts and steals for all $x \in [\overline{r}-\rho, \rho/(1-\psi))$ and accepts without theft for all $x \in [\rho/(1-\psi), \infty)$. fight up to period $\underline{\omega}$. At $t = \underline{\omega}+1$ the weak type accepts $x(\overline{\omega}-\underline{\omega}-1)$ and the strong rejects. There exist government prior beliefs $\underline{p}$, $\overline{p}$ and a probability $b \in (0,1)$ of $p = \underline{p}$ such that this strategy profile is an essentially unique perfect Bayesian equilibrium.*

Proof. The argument proceeds analogously with the proof of Proposition 4, Appendix D.6. We first show that there always exist exists government types $\underline{p}$, $\overline{p}$ implying any conflict duration $\underline{\omega}$, $\overline{\omega}$. We then verify that no government type has a profitable deviation from the strategy, given beliefs. Next, we show that there exists a probability $b \in (0,1)$ such that there are no profitable deviation for the rebels. Finally, we argue that the strategy profile is unique up to actions off the equilibrium path.

To begin, notice by Proposition 4 that the function $x(\omega)$ defines the efficient separating offer under minimal amnesty for an expected conflict duration of $\omega \in \mathbb{Z}_+$. Fix some $\Delta > 0$, implying a minimum separating threat τ in (25), and integers $\underline{\omega} < \tau$, $\overline{\omega} > \underline{\omega} + \tau + 1$. Let $\underline{g}(\omega) = \pi - (1-\delta)x(\omega) - \delta\underline{x}$ and $\overline{g} = \pi - (1-\delta)\kappa - \delta\overline{\theta}\overline{x}$ the government's value of making a separating offer to a weak and strong rebel type under a threatened conflict duration of $\omega > \tau$. The expected value of a separating offer is $p\underline{g} + (1-p)\overline{g}$. Treating ω as a parameter, consider the beliefs $\tilde{p}(\omega)$ for which the government weakly prefers offering generous amnesty $\pi - \overline{x}$ to the separating offer $x(\omega)$

$$\begin{aligned} \tilde{p}(\omega) &:= p : 1 - \overline{x} \geq p\underline{g}(\omega) + (1-p)\overline{g} \\ \implies \tilde{p}(\omega) &= \frac{(1-\delta)\kappa + \delta\overline{\theta}\overline{x} - \overline{x}}{(1-\delta)\kappa + \delta\overline{\theta}\overline{x} - [(\delta\underline{\theta})^{\omega}\overline{x} + (1-(\delta\underline{\theta})^{\omega})\underline{x})]} \end{aligned} \tag{29}$$

where $0 < \tilde{p}(\omega) < 1$ for all $\omega \geq 1$. We set the determined government type's belief so the fighting time (22) evaluates to $\overline{p} : \omega(\overline{p}, p^*(\overline{r}+\ell)) = \overline{\omega}$, which exists as $\lim_{p\to 1}\omega(p, p^*/r + \ell) = \infty$ and $\lim_{p\to p^*(\overline{r}+\ell)}\omega(p, p^*/r+\ell) = 0$. The wavering government's belief may is then set to satisfy $\underline{p} : \omega(\underline{p}, \tilde{p}(\overline{\omega} - \underline{\omega} - 1)) = \underline{\omega}$, where $\underline{\omega}$ is treated as a fixed number.

We first consider deviations by the wavering government type with belief $\underline{p}$. By construction they strictly prefer to fight rather than offer $\overline{x}$ for the first $\underline{\omega}$ periods and to offer $\overline{x}$ for $t \geq \underline{\omega}$. Next, consider the determined government type with belief $\overline{p}$. Suppose this was common knowledge. Then there exist separating offers that both determined government and rebels prefer to the equilibrium action of fighting. But the wavering type strictly prefers this separating offer and would mimic any signal. The rebel therefore does not infer anything from offers made before $\underline{\omega} + 1$.

Turn to the rebels, where following the reasoning in Proposition 4, the best response of weak types is reject and fight at $t = 0$ if the expected value of rejecting exceeds the bound $\overline{r}$ of a separating offer

$$b\left(\underline{r} + (\delta\underline{\theta})^{\underline{\omega}}(\overline{r} - \underline{r})\right) + (1-b)\left(\underline{r} + (\delta\underline{\theta})^{\overline{\omega}}(\overline{r} - \underline{r})\right) \geq \overline{r} \tag{30}$$

which holds as $b \to 1$ since $\underline{\omega} < \tau$. But if they are willing to fight at $t = 0$ they will certainly fight at $t \geq 1$ as their beliefs are unchanged. At $\underline{\omega} + 1$ they infer that the government is determined for sure and a separating offer exists. □

D.10 Prediction 1 Military capability, amnesty terms

Notice that for subgame perfect equilibria with minimal or low-powered incentives $\overline{x}$ is independent of rebel battlefield damage. It may be verified directly that the critical belief $p^*(\overline{x})$

in (9) is increasing in κ as

$$\frac{\partial p^*}{\partial \kappa} = \frac{1-\delta}{\delta(\overline{\theta}-\underline{\theta})\overline{x}} > 0$$

and hence the duration of fighting periods before generous amnesty receipt is falling. But for high-powered equilibria the cost of generous amnesty, rebel battlefield damage, and the critical belief (12) are jointly determined. If the pooling offer equals the recovery-adjusted reservation value

$$\frac{1-\delta}{1-\delta\theta-(1-\delta)\psi}\alpha$$

the payoff is independent of κ and conflict duration unambiguously decreasing in κ by the argument above. If not, it may be verified that the critical belief $\hat{p}$ is increasing in κ by the following steps

$$\hat{p}(\kappa) = \frac{(1-\delta)\kappa}{x(\hat{p}(\kappa))\delta(\overline{\theta}-\underline{\theta})} - \frac{1-\delta\overline{\theta}}{\delta(\overline{\theta}-\underline{\theta})}$$

$$\Longrightarrow$$

$$\hat{p}(\kappa) = \frac{\kappa}{\kappa-(1-\gamma)\mu}\frac{\delta[(1-\overline{\theta})+\hat{p}(\kappa)(\overline{\theta}-\underline{\theta})]}{\delta(\overline{\theta}-\underline{\theta})} - \frac{1-\delta\overline{\theta}}{\delta(\overline{\theta}-\underline{\theta})}$$

$$\Longrightarrow$$

$$\hat{p}(\kappa) = \left(1-\frac{\kappa}{\kappa-(1-\gamma)\mu}\right)^{-1}\left(\frac{\kappa}{\kappa-(1-\gamma)\mu}\frac{1-\overline{\theta}}{\overline{\theta}-\underline{\theta}} - \frac{1-\delta\underline{\theta}}{\delta(\overline{\theta}-\underline{\theta})}\right) \tag{31}$$

$$\Longrightarrow$$

$$\hat{p}(\kappa) = \frac{\kappa-(1-\gamma)\mu}{\kappa}\frac{1-\delta\underline{\theta}}{\delta(\overline{\theta}-\underline{\theta})} - \frac{\kappa-(1-\gamma)\mu}{(1-\gamma)\mu}\frac{1-\overline{\theta}}{\overline{\theta}-\underline{\theta}}$$

$$\Longrightarrow$$

$$\frac{\partial\hat{p}(\kappa)}{\partial\kappa} = \frac{1}{(1-\gamma)\mu}\left(\frac{1-\delta+\delta(\overline{\theta}-\underline{\theta})}{\delta(\overline{\theta}-\underline{\theta})}\right) > 0$$

and so conflict duration is decreasing in κ.

E Additional empirical results

E.1 Robustness of event-study

In this section, we conduct several additional robustness tests to support the validity of teh event-study estimates in Section 4.1.

Differences-in-discontinuities: To restore the notion of a control group, consider the differences-in-discontinuities estimate, which captures the difference in the structural breaks in the time series' for treated and control areas separately. Figure A5 plots the event studies on subsamples that are treated (Panel A) and untreated (Panel B), with flexible polynomials in time plotted on either side of the amnesty date for both groups. As Figure A5 reveals, the aggregate event-study effect is driven entirely by areas that received amnesty, ruling out spurious contemporaneous shocks that would affect all parts of the Niger Delta equally. This should bolster confidence that Figure 5 captures the short-run effect of the amnesty itself, rather than global trends or correlated shocks. Interestingly, Figure A5 also demonstrates the failure of parallel trends, as the pre-amnesty conflict trend slopes upward in Panel A but remains generally flat in Panel B.

Specifications: Columns (1)-(2), (3)-(4), and (5)-(6) of Table 1 consider linear, quadratic, and quartic specifications of time trends, respectively. For each specification, we include either Year or Month-of-Year effects. Each panel of Table 1 contains a different bandwidth. Lastly, columns (7) and (8) revert to a linear time trend specification but include a lagged dependent variable (AR(1)) term. The estimates on the post-amnesty indicator are significant at the 5 or 1% level in all but 5 out of 36 specifications, though even in these insignificant specifications the magnitudes remain large but standard errors increase as sample size falls. Amnesty robustly reduces violence by between 3-7 monthly militant attacks on average.

Placebo and outlier tests: In Figure A6 we consider whether July 2009 may have simply been a "lucky choice" of event date, and whether the results can be replicated with other placebo event dates. We estimate the event-study for 225 placebo dates and compare the estimate of $\hat{\theta}$ to this distribution of coefficients and t-statistics. We find that the results are not replicated by these placebo dates, and the estimated $\hat{\theta}$ falls in the left tail of this distribution. A similar permutation-type test is conducted in Figure A7 to ensure that influential observations do not drive the results, by using a leave-one-out estimation procedure for each observation in the event window. Again, we plot the histogram of estimates and t-statistics, demonstrating that for any observation that is dropped, the results remain robustly negative and significant, allaying concerns that very high conflict levels immediately before the amnesty are entirely responsible for the observed effect.

Placebo conflict outcomes. We consider robustness to placebo outcomes. In particular, we support the assumption of counterfactual continuity over the event date by showing that non-oil conflict events are not discontinuous at the event date. In table B3 we estimate the same event-study specifications as in Table 1, only using all non-oil conflict as the outcome of interest. The effect magnitudes and signs do not display any consistent pattern and are generally insignificant. In fact, the estimates are only significant at the 5% level in 2 out of 28 specifications, likely due to noise. In general, it appears as though non-oil conflict does not change discontinuously in the month of the amnesty announcement. This suggests that the treatment effect is probably not being driven by monthly time effects and/or correlated aggregate shocks that would affect the likelihood of conflict broadly. Indeed, only conflict specifically targeted by the amnesty actually falls in response to the policy.

Structural break test. Another way to validate our event-date is to conduct a data-driven test of the presence and location discontinuous breaks in the time series of militant attacks without imposing our ex-ante known event-date. To do this, we follow the literature on structural breaks, as summarized in Perron (2006). The for the presence of a structural break, we estimate the following model

$$m_t = \alpha + \rho m_{t-1} + e_t$$

And test whether the estimated coefficient $\hat{\alpha} = E[m_t|m_{t-1}] - \hat{\rho} m_{t-1}$, varies over time. In other words, we test whether the mean of the time series, after subtracting the AR(1) persistence term, varies over time. To test both the existence and identify the location of a structural break, we use a supremum Wald test where the test statistic over a sample of size T is $\sup_t S_T(t)$, where $S_T(t)$ is the Wald statistic evaluated at a potential break date t. In other words, the estimated structural break corresponds to the largest estimated Wald statistic over all possible dates, with the p-value derived from the limiting behavior of this supremum (see Perron (2006) for details).

The results are summarized for event-windows varying from 20 to 45 in Figure A8 and Table B4. Figure A8 plots histograms of date-specific Wald statistics for each bandwidth, with a vertical line indicating the magnitude of the Wald statistic on our known amnesty date. The results indicate that the Wald statistic corresponding to the amnesty date is in the far-right tail of this distribution for every bandwidth. In Table B4, we test the significance of these estimates, displaying the corresponding maximal break date in event-time. The break date corresponds to the amnesty in 4 out of 6, all of which are significant at the 1% level. In the 20-month window test, the break date corresponds to one month after the amnesty, significant at 5%. Only in the 25-month window does the estimated date not correspond to the amnesty period. However, as shown in Figure A8, the Wald statistic at the true amnesty date is still 2nd highest in this distribution.

Optimal lag length. Table 1 allows only persistence of the AR(1) form. In this section, we estimate optimal lag-length across a variety of criteria and then re-estimate. Following the literature on optimal order selection in VAR models (see Lütkepohl 2005 and Nielsen 2006), of which our univariate setting is a special case, we estimate the optimal lag length p for a generic autoregressive model of the form

$$m_t = \alpha + \sum_{i=1}^{p} \beta_i m_{t-i} + \epsilon_t$$

by evaluating several model fit statistics for varying p, taking 10 as our maximum lags. In particular, we consider the likelihood ratio (LR) test, which jointly tests the null hypothesis that the coefficient on the pth lag is zero. The optimal lag is then selected by the largest lag with a 5% significant likelihood-ratio test statistic. Alternatively, we also use several additional measures of model fit, including the final prediction error (FPE), as well as likelihood-based metrics such as the Akaike information criterion (AIC), the Hannan and Quinn information criterion (HQIC), and Schwarz's Bayesian information criterion (SBIC). All model fit statistics are calculated without including the constant.

Table B5 provides the results of re-estimating the main linear event-study specification without fixed-effects with optimally-selected lags. The columns give the results of each different model fit statistics, while the each sub-table re-estimates for a different event-window, as in Table 1. The optimal lag-length, indicated in each sub-table footer, varies substantially across event-windows and model fit criteria, varying from 0 to 9 lags. However, for any optimally-selected p, the results remain negative and significant at the 1% or 5% levels.

E.2 Robustness of difference-in-differences

In this section, we address several additional alternative explanations that might account for the post-amnesty rise in bunkering concentrated in areas where active militant groups received Federal amnesty: rising oil prices, generic violations of parallel trends, endogenous militant location choice and rebounding oil output. We also test robustness to the definition of the treatment and measurement error in sabotage spills.

Oil prices: The post-amnesty period was one of steadily rising oil prices after a crash in early 2009. This increased the profitability of bunkering, and may have lead to a differential increase in oil theft among amnestied communities if we believe ex-militant leaders have a comparative advantage in the illegal market, due perhaps to sunk fixed costs, knowledge of oil company operations, or networks of ex-combatants. However, even in absence of amnesty oil prices could generate these incentives for existing militant groups, suggesting a violation of parallel trends. We test for this explanation in columns (4) and (6) of Table 2 by including the interaction between oil prices and T_i^d, the amnesty treatment indicator. Comparing columns (4) to (1) and (6) to (5), we find that this inclusion actually somewhat increases the main coefficients' size. It is clearly not responsible for the main effect, but the $Price_t T_i^d$ term is significant, indicating that oil theft in militant controlled-areas responds positively to increasing profitability, as expected.

Choice of control group: Militants choose their location endogenously. They are more likely to locate near oil assets, in coastal riverine communities, and are disproportionately located in the core oil-producing states of Delta, Bayelsa, and Rivers. It may therefore be the case that there oil theft outcomes are affected by fundamental differences between communities that are militant- and non-militant, or oil-producing and non-oil producing. To explore whether comparability on unobservables can be enhanced we consider the following battery of alternative control groups: The synthetic difference-in-difference approach of Arkhangelsky et al. (2019), controlling by using non-amnestied militant regions as a control group, and restricting the sample to respectively the core Niger Delta, oil producing, and pipeline communities only.

Results are given in Figure A19, which displays parallel trends plots, and in Table B7, which provides estimates of the main TWFE model, for each control group or sample restriction. The parametric estimates in Table B7 suggest that changing the control group or the sample does not materially affect the estimates. Column (6), labeled, "placebo," takes as the control group only those village within 30 kilometers of a *non*-amnestied militant camp, that is camps controlled by commanders who were either defeated in battle or not amnestied. While this does not address with variable militant capabilities that might be correlated with amnesty take-up, at the very least it holds militant the extensive margin of activity constant. Across all six alternative specifications the estimated coefficient of interest ψ_τ is of comparable magnitude to the baseline model.

In Figure A19, the parallel trends assumption appears to be met for all of the different choices of control group. Whereas in the full sample baseline approach there are small but significant differences between treatment and control in some of the pre-periods, these differences disappear completely when using different control groups or samples.[39] This improvement in balance on pre-treatment outcomes suggests that there is some benefit in refining the control group to increase comparability.

Oil output: A plausible but uninteresting explanation for the result, which does not necessarily suggest a violation of any key assumptions but is nonetheless outside the model, is that rebounding oil production, as see in Figure A3 Panel C, increased the value of bunkering and thus the level of theft. In this interpretation, the growth in the black market is entirely – or at least partly – an adverse consequence of bringing oil output back to pre-conflict levels.

We test this hypothesis by estimating a pipeline-level TWFE model that controls directly for oil production, in addition to one that normalizes sabotage by annual oil flow.[40] Oil flow

[39] In the case of the synth-DD, this is by construction.

[40] We use a pipeline-level specification rather than an oilfield-level one because, as shown in the next section, sabotage is concentrated on the pipelines that connect fields to larger flowstations and export terminals. Sabotage

is calculated as the average daily flow volume from all fields connected to a pipeline segment through that pipeline segment and toward the nearest export terminal. All segments are normalized to be of equal length, at xx kilometers. If the impact on oil production is driving the bunkering rise, then controlling for or normalizing the outcome by oil flow should kill the effect. In Figure A20, we re-create the parallel trends plot at the pipeline-level, comparing raw sabotage (Panel A) with sabotage normalized by 000s of barrels per day (Panel B). The trends look broadly similar. Table B9 shows the regression results, which indicate that moving to the pipeline level preserves the main result (column (1)), that this is not sensitive to the inclusion of oil flow (column (2)), and that the effect remains significant at 1% for the normalized outcome variable (column (5)). They are 63 and 93% of the sample means, respectively, which is of the same order of magnitude as the main village-level estimate.

Measurement error: Measurement error might contaminate the causal estimate of ψ if data collection is a function of amnesty and militancy over time. In a plausible story, the capacity of the environmental regulator may be growing over time, so they conduct more JIVs and leave fewer spills uninvestigated. Further, this secular trend may not be captured by year fixed effects if it is precisely in militant-controlled regions where collecting data is most difficult pre-amnesty, and therefore improves the most after the security situation improves. If this is the case, we should observe that malfunctions also rise in amnestied territories post-amnesty, since all types of spills would be subject to the same measurement error. Figure 6 Panel B, Table B8 and Table B9 columns (3) and (4) all show precisely estimated null effects with no pre-trends.

Treatment definition: The estimated effect may be sensitive to the use of different thresholds for d. We should expect to see that as d rises, the estimated ψ falls in magnitude, as the radius begins to include villages that are only marginally under militant control at the time of the amnesty. This is borne out in Figure A21 Panel A. With the exception of $d = 10km$, which has an estimated coefficient near zero, the effect falls as the radius from an amnestied camp expands as expected.[41] However, eve at $d = 100km$ we still observe a positive and significant amnesty effect. Panel B provides the equivalent estimates for malfunction spills, which are consistently near zero for any d.

Maximum likelihood estimation: We also consider the robustness of the results to non-linear specifications. In particular, the outcome data is measured as a count and contains zeroes for many of the village-years. To account for these facts, in Table B10, we re-estimate the main specification using fixed-effects (Panel A) and zero-inflated (Panel B) Poisson regressions. In the first-stage of the zero-inflated model, we estimate the probability of having nonzero theft as a function of distance to the nearest oilfield and pipeline using a Logit

on wellheads within fields thus captures only a small share of the relevant bunkering incidents.

[41] There are several explanations for a null effect at 10 kilometers. Most likely, militant camps are typically deep in the Niger Delta creeks and thus at a slight geographic remove from bunkering targets.

model. We estimate the models for the full sample in columns (1)-(3) using various combinations of controls interacted with year dummies as well as oil prices interacted with the treatment. In columns (4)-(6), we test robustness of the the results to using sub-samples of core Niger Delta and oil-producing communities, as well as the placebo control group. For columns (1)-(5) the results are positive and significant in all specifications across both models and broadly similar in magnitude. The coefficient more than doubles in size when we consider the restricted placebo control group in column (6).

E.3 Infrastructure type

A final piece of evidence suggests that the increase in pipeline vandalism is driven by the economic incentives for organized crime generated by amnesty. We can disaggregate third-party sabotage across four general categories of oil infrastructure – trunklines, flow lines, delivery lines, and wellheads. Because of their width, trunklines – which funnel large volumes of crude from intermediate inland flow stations to coastal export terminals – are by far the most economically valuable oil assets to bunkerers. In contrast, smaller flowlines require similar tapping techniques but yield less crude, while wellheads are of little value for oil thieves. Delivery lines are larger than flowlines but smaller than trunklines, yielding intermediate profits. Organized criminal bunkering motivated by profit maximization should be concentrated primarily on trunklines.

FIGURE A23 HERE

We re-estimate the main village-level TWFE equation, including oil intensity interacted with post-amnesty, decomposing the sabotage outcome variable into its four additive infrastructure type components. As a further falsification test, we also run the same regressions for the components of equipment failure spills. The results are in Figure A23. Each portion of the plot shows the corresponding sabotage and equipment failure coefficients for the subset of spills occurring on trunklines, delivery lines, flowlines, or wellheads. The results show a large and significant effect of amnesty on trunkline sabotage. They also indicate a smaller – though still significant – positive effect on delivery lines. As hypothesized, flowlines and wellheads see null effects. Across all infrastructure types, malfunction spills are not affected. These results suggest that organized crime – rather than unorganized acts of sabotage stemming from community discontent or attempts to gain compensation – is driving the aggregate observed effect of amnesty on bunkering.

E.4 Rebel fighting threats

A central assumption in our model is that rebel groups are able to immediately resume fighting in the event that government reneges on the bargain, for instance due to an exogenous shock to beliefs. The credible threat of renewed violence generates incentives for government to maintain payments, see the derivation of (9) in Section 5.3. The assumption of retained military capability post-amnesty may appear unreasonable since a stated goal of the amnesty was to disarm and demobilize militant groups, pushing post-amnesty military capability κ to zero. However, several pieces of evidence suggest that this aspect of the amnesty was largely toothless, and rebels maintained the ability to remobilize quickly. Qualitatively, there are ample reports of powerful militant commanders surrendering very few weapons (Ebiede, Langer, and Tosun 2020).[42]

In May of 2016, newly elected President Muhammadu Buhari announced that amnesty funding would be slashed by 70%.[43] Plausibly, President Buhari was more optimistic on the prospect of military victory than his predecessor, so his election constituted a change to the rebels strategic environment. We use this announcement to test whether rebel fighting capabilities are sufficient to adapt and respond. Figure A11 presents an RDiT plot for the Buhari announcement. The results indicate an immediate and large jump in oil-related militant attacks directly following the announcement, without any noticeable lag. In Table B6 we estimate the size of this short-run effect to be between 2.1-10.5 additional monthly militant attacks. This estimate is always positive and nearly always significant irrespective of the specification choice or event-window, though these choices affect the magnitude.

Figure A12 again conducts the RDiT placebo test across 211 potential event dates, and shows that the announcement date is in the far right tail of the distribution of coefficient estimates (Panel A) and t-statistics (Panel B). Taken together, these results suggest an immediate, short-run causal effect of Buhari's announcement on conflict. The fact that this increase is similar in magnitude to the reduction in militancy observed after the amnesty provides strong evidence in favor of symmetric fighting capabilities before and after amnesty.

[42] Ebiede, Langer, and Tosun (2020) note that the militant leader known as Young Shall Grow surrendered just one AK-47, despite controlling multiple camps in our data.

[43] https://www.ft.com/content/6d7f2766-15bf-11e6-b197-a4af20d5575e