

Do Giant Oilfield Discoveries Fuel Internal Armed Conflicts?

Yu-Hsiang Lei and Guy Michaels¹

This version: 13 October 2011

Abstract:

We use new data to examine the effects of giant oilfield discoveries around the world since 1946. On average, these discoveries increase per capita oil production and oil exports by up to 50 percent. But these giant oilfield discoveries also have a dark side: they increase the incidence of internal armed conflict by about 5-8 percentage points. This increased incidence of conflict due to giant oilfield discoveries is especially high for countries that had already experienced armed conflicts or coups in the decade prior to discovery.

Keywords: Natural Resources, Resource Curse, Petroleum, Armed Conflict, Civil War

JEL Classification: Q34, Q33, O13

¹ Department of Economics and Centre for Economic Performance, London School of Economics.
Corresponding author: Michaels (g.michaels at lse.ac.uk). We thank Christian Siegel for excellent research assistance and Emek Basker, Tim Besley, Francesco Caselli, Ethan Ilzetzki, Marco Manacorda, Steve Pischke, Steve Redding, Daniel Sturm, and John Van Reenen for helpful comments. We also thank the Centre for Economic Performance for its generous financial support.

1. Introduction

Do natural resource windfalls, such as those arising from the discovery of giant oilfields, increase the risk of internal armed conflict? Anecdotal evidence from Nigeria, Angola, and Iraq lead us to suspect that they may, and recent research (Dal Bó and Dal Bó 2011, Besley and Persson 2009, 2011, and Acemoglu et al. 2010) even sheds light on the mechanisms that underlie some of these conflicts over resources. But as Norway, Canada, and Brazil show, not all oil rich countries experience conflict. Careful surveys of the literature on conflicts and natural resources (e.g. Ross 2004, 2006 and Blattman and Miguel 2010) show how difficult it has been to estimate the causal effect of oil on armed conflict in all but a handful of countries.² The goal of this paper is to examine whether giant oilfield discoveries really do fuel internal armed conflicts around the world, and if so – in which settings.³

We begin with a simple model, following Besley and Persson (2009), which guides our empirical analysis. In this model, giant oilfield discoveries increase oil revenues, generating windfall income for the incumbent. When the incumbent cannot credibly commit to share this windfall, the opposition may mobilize to challenge him, and this may lead to an internal armed conflict. Such conflicts over resources are especially likely in countries where political violence tends to translate into political and economic gains.

To investigate this model's predictions, we ideally require exogenous variation in resource windfalls. Finding such variation in multiple countries is challenging, since cross-country (or cross-conflict) comparisons may be contaminated by omitted variables bias. Using panel data to absorb country fixed effects is not straightforward either, because the quantity of natural resources extracted is a choice and oil prices may be affected by violent conflict. To overcome this difficulty, we focus on the discovery of giant oilfields, each of which contained ultimate recoverable reserves (URR) of 500 million barrels (bbl) equivalent or more before extraction began (data on these giant oilfields are reported in Horn 2004).⁴ Of

² Studies of the causal effect of natural resources on conflict tend to focus on specific countries. For example, Angrist and Kugler (2008) and Dube and Vargas (2010) study the effect of resource windfalls on conflict in Colombia, and Bellows and Miguel (2009) study this effect in Sierra Leona. Perhaps closer to our paper is contemporaneous work by Cotet and Tsui (2010) on oil and conflict, which we discuss below.

³ Unless otherwise specified, we use “oil” as a shorthand for both oil and gas.

⁴ For brevity we use a broad definition of oil, which also includes condensate and natural gas. To determine whether an oilfield has estimated ultimate recoverable reserves of 500 Million bbl of oil equivalent or more, the estimated reserves of oil and condensate are summed up. These are then added to the amount of natural gas, which is converted to oil at a ratio of 6,000 cu ft/bbl (Horn 2004). Note that ultimate recoverable reserves include the amount already extracted and the amount that has not yet been extracted.

the 910 giant oilfields that were known as of 2003, we focus on the 782 giants that were discovered since 1946 in 65 different countries.

We argue that in a panel of countries with country and year fixed effects, giant oilfield discoveries are plausibly exogenous. To see why, consider how important giant oilfields are as a global source of hydrocarbons. Horn (2007) concludes that giant oilfields account for over 40 percent of the world's URR of oil and gas. Discoveries of these fields are therefore economically important events, which are rare in all but a handful of countries: in only about 4.6 percent of the country-year observations in our data was one or more giant oilfield discovered. It is true that countries can influence the prospecting efforts within their territory, and thus affect the discovery rate. But prospecting for oil is highly uncertain, and the odds of finding a giant oilfield are usually low, which is not surprising given the size of the prize. We therefore assume that the timing of giant oilfield discoveries is plausibly exogenous and discuss below some robustness checks that support this interpretation.⁵ But before we further discuss our causal interpretation of the findings, we first describe them.

We find using a panel of 193 countries from 1946-2008 that on average oil production increases by about 35-50 percent within 4-10 years of a giant discovery. Giant oilfield discoveries similarly increase oil exports by about 20-50 percent within 6-10 years.

Having found evidence suggesting a large impact of giant oilfield discoveries on oil output, we examine the model's predictions regarding conflict. We find that on average giant oilfield discoveries increase the incidence of internal armed conflicts (measured as a year with 25 or more conflict casualties) by about 5-8 percentage points within 4-8 years of discovery, compared to a baseline probability of about 10 percentage points.

We also find that the discovery of giant oilfields is especially likely to fuel internal conflicts in countries with recent histories of political violence. For example, giant oilfield discoveries increase the incidence of internal armed conflict by about 11-18 percentage points (compared to a baseline probability of about 37-39 percent) when a country experienced at least one such conflict in the decade prior to discovery. The effect of discovery on the incidence of internal armed conflict is similarly 11-14 percentage points (compared to a baseline probability of about 19-20 percent) in countries that experienced at least one coup in the decade prior to discovery. By contrast, in countries that experienced no internal conflicts or

⁵ Interestingly, we find no significant correlation between the number of giant oilfields discovered in a year and the inflation-adjusted price of oil in that year from 1946-2003. Still, if forward-looking agents search for oil and find it when they expect to have good opportunities to take advantage of giant discoveries, it is more likely that peace, rather than conflict, should follow discoveries. But our findings suggest that giant discoveries are actually followed by conflict.

coups in the decade before a discovery, there is no significant effect of giant oilfield discoveries on the incidence of internal armed conflicts.

Turning to the effect of giant oilfield discoveries on economic outcomes, we find suggestive evidence that per capita GDP and government spending may increase by about 6 percent within the decade following a giant discovery. But unlike our results on conflict, these results are not robust to all the specifications we consider, which we discuss below. Moreover, we find no evidence that giant oilfield discoveries significantly affect private consumption or spending.

To support our interpretation that the findings described above are the causal consequences of giant oilfield discoveries, we report results from a number of robustness checks. First, we address the concern that giant oilfield discoveries may have resulted from economic or political changes that preceded them. Reassuringly, we find no evidence of significant economic or political changes in the five years leading up to giant oilfield discoveries. We also test whether giant oilfield discoveries follow lulls in previous conflicts, and find no evidence to support this hypothesis. Second, we tackle the concern that giant oilfield discoveries are serially correlated over time, because some oilfields are close together, so one finding one may lead to another. While it is true that giant oilfield discoveries in a country's recent past increase the odds that it finds a giant oilfield in a given year, controlling for these past discoveries leaves the findings discussed above essentially unchanged. Our results are also robust to excluding country-year observations within a decade or less of previous giant discoveries. Observations with giant oilfield discoveries account for only about 1 percent of the remaining sample, making them especially difficult to anticipate. Third, we address concerns that economic or political conditions shortly before discovery may affect our estimates, by showing that our results are robust to controlling for (instrumented) lagged dependent variables, lagged institutional quality (Polity 2), and lagged aggregate private investment. Fourth, we tackle the concern that observations with oil discoveries are different from others in ways that are difficult to measure directly. To do so, we use the Oil and Gas Journal Data Book (2008) to restrict our sample to observations where at least one oil discovery – not necessarily of a giant oilfield – was made. Regressions using this sample compare the effect of giant oilfield discoveries to the effect of smaller oilfield discoveries. Remarkably, even in this restricted sample we find that our results hold.

Our finding that giant oilfield discoveries fuel internal conflicts in countries prone to violence has some implications for policy. For example, those who strive to reduce armed conflict should be concerned about oil rents that incumbents obtain in conflict-prone areas, especially

if those rents encourage challenges to the incumbents' power. At the same time, the firms that prospect for oil in conflict-prone areas and those who regulate them ought to be concerned about negative externalities for many locals, who have little to gain from giant oilfield discoveries but may suffer from conflicts over the oil.

The remainder of the paper is organized as follows. Section 2 discusses the related literature, Section 3 presents a model of conflict over oil revenues, Section 4 discusses the data, Section 5 presents our results, and Section 6 concludes.

2. Related Literature

Concerns that some natural resources - including oil - may fuel internal armed conflicts arise from looking at oil-rich countries, such as Angola, Colombia, Iraq, Sudan, and Indonesia. A number of influential papers, including Collier and Hoeffler (1998, 2004) and Reno (1999) have investigated the relationship between natural resources and conflict, sparking considerable interest among social scientists and policy makers. Surveys of the developing literature on this topic, including Ross (2004, 2006), Humphreys (2005), and Blattman and Miguel (2010), tend to conclude that there is evidence linking oil to some instances of internal armed conflict. At the same time, not all oil-rich countries experience internal armed conflict, so these conflicts are clearly not inevitable.⁶

Theoretical studies of the links between natural resource rents and conflict have focused on the possibility that these conflicts are the result of competition over resources. Summarizing this literature, Blattman and Miguel (2010) point out that models of armed conflict typically consider the cases where property rights are not well-protected, contracts are imperfectly enforced, and rulers are not always replaced by fair elections. Recent contributions to the literature on conflicts over resources include Garfinkel and Skaperdas (2007), Dal Bó and Dal Bó (2011), Besley and Persson (2009, 2011), Caselli and Cunningham (2009), and Acemoglu et al. (2010).

But despite all this research on the relation between natural resources and armed conflict, establishing the causal effect of resource windfalls on conflict around the world has proved difficult. Some of the best-identified studies examine causality using regional variation within

⁶ For example, Michaels (2011) and Caselli and Michaels (2009) find no evidence of armed conflict in the U.S. South and in Brazil.

countries. For example, Bellows and Miguel (2009) find that chiefdoms with more diamond wealth in Sierra Leona experienced more armed clashes, and studies of Colombia find that high coca prices increase conflict in coca producing regions (Angrist and Kugler 2008) and high oil prices increase conflict in areas where oil is extracted from or shipped through in pipelines (Dube and Vargas 2010).

Collectively these and other within-country study suggest that natural resource windfalls can fuel armed conflicts, at least in some countries and settings. But in order to generalize these findings to the rest of the world and to better understand in what settings natural resource windfalls are more likely to cause armed conflict, it seems useful to look beyond the boundaries of specific countries. It turns out, however, that using variation from multiple countries to identify the effect of natural resource on conflict is not straightforward. To see why this is a challenge, consider first comparisons of resource rich countries with resource scarce ones, or of conflicts that occur in resource rich and resource scarce parts of the world. The main concern about this approach is that resource-rich areas may differ from others in ways that are difficult to measure and control for. For example, the Middle East is rich in oil but it also differed from other parts of the world in important ways before oil was discovered. These differences, which are notoriously hard to measure, along with oil abundance, may have caused subsequent conflicts, and telling apart the causes is difficult. To overcome the problem of fixed differences between countries, we could consider a second approach, which interacts country-specific measures of oil abundance with variation over time in oil prices. But this approach suffers from a problem of reverse causality, since conflicts may raise oil prices, as they probably did during the Arab-Israeli War in 1973, the Iranian Revolution in 1979, the Kuwait War in 1990, and the Libyan Civil War in 2011, making the direction of causality between conflict and resource revenues difficult to ascertain.⁷ A third approach we could try instead uses time-varying measures of oil production or exports in each country. But this approach also has problems in shedding light on causality, since countries choose the amount of oil they extract, and potential buyers may also choose how much to buy from whom. These choices may respond, directly or indirectly, to armed conflicts or their underlying causes.⁸

⁷ The possibility that internal conflict in Libya increased oil prices was discussed by the media. See for example: <http://www.bbc.co.uk/news/business-12522291>. When we regress an indicator for internal conflict on an interaction of an indicator for countries with at least one giant oilfield and the log of inflation-adjusted oil price, controlling for country and year fixed effects, we get a coefficient of 0.044 (s.e. 0.024), suggesting a positive and marginally significant relation between the two.

⁸ For example, the recent internal armed conflict in Syria appears to have reduced its oil production: <http://www.ft.com/cms/s/0/c9d67952-e823-11e0-9fc7-00144feab49a.html#axzz1aOqrle6u>

Since identifying the causal effect of natural resources on conflict is difficult using the approaches described above, our paper focuses on the discovery of giant oilfields as a more plausibly exogenous source of variation. The closest paper to ours in this respect is contemporaneous work by Cotet and Tsui (2010), which examines the relationship between oil reserves and discoveries and conflict in using a panel of countries. Our paper differs from theirs in a number of ways. First, our analysis is guided by a theoretical model, and this leads us to examine predictions that differ from theirs. Second, we examine oil output and exports, in addition to conflict incidence as our main outcomes of interest, while they focus primarily on conflict onset. Third, we use different data from theirs - our analysis is focused on giant discoveries and we use smaller discoveries only as controls. Fourth, in addition to examining the main effects of giant discoveries we investigate interactions of giant oil discoveries with measures of past violence, while they examine interactions of (not necessarily giant) oil discoveries with democratic and non-democratic countries. Finally, we use different empirical specifications, which examine lagged outcomes within a decade of discovery, while they focus on contemporaneous outcomes. Despite these differences, our findings, like theirs, suggest no contemporaneous relation between oil discoveries and the onset of conflict. But we also look a few years after giant discoveries and do find an increase in the onset of conflict. When we focus on conflict incidence, rather than onset, the effects we find are even bigger. And as we note above, the higher incidence of conflict following discovery appears to be due to discoveries in countries with recent histories of violence. But before we turn to our empirical evidence, we first outline a simple model that helps guide our empirical investigation.

3. A Model of Conflict for Resources

To guide our empirical analysis, we begin with a simple model of conflict over resources, following Besley and Persson (2009).⁹ The model focuses on two potentially conflicting groups denoted by J : an incumbent I and an opposition O . Each group makes up half of the population and can mobilize a fraction A^J of its citizens to serve in its army. The decision of each group whether or not to mobilize an army is discrete, and is denoted by $\delta^J \in \{0, A^J\}$. The probability that power transitions from the incumbent to the opposition is determined by

⁹ As we explain below, we depart from their model only in relatively minor details.

a conflict function: $\text{Prob}(\text{change of power}) = \frac{1}{2} + (1/\mu)[\delta^O - \delta^I]$. The parameter μ captures the degree to which the country can resist political violence, and low values of μ mean that it is more prone to political violence as a means of transferring power. We assume that $A^I/\mu \leq \frac{1}{2} \leq 1 - A^O/\mu$, which holds as long as μ is sufficiently large.

The winning group has access to government revenue denoted by R , which comes from natural resources. These resources must be shared according to an institutional rule, which stipulates that the incumbent gets $(1-\theta)2R$ and the opposition gets $2\theta R$, where $\theta \in [0, \frac{1}{2}]$. In other words, we consider sharing rules that range from institutions that lead to complete equality ($\theta = \frac{1}{2}$) to institutions where the winner takes all ($\theta = 0$).

In addition to any revenues they may receive from the government, each citizen supplies one unit of labor to the market, earning a real wage w . A group that wants to finance its army does so by taxing its population. Since we are interested primarily in bilateral internal conflicts (as opposed to one-sided conflicts), we depart from Besley and Persson (2009) by assuming that the opportunity cost of fighting is equal for the opposition and the incumbent.¹⁰

The timing of events within each period is as follows. First the amount of resources at the government's disposal, R , is determined randomly. We assume that if a giant oilfield is discovered then $R = R^H$, and otherwise $R = 0$.¹¹ Second, the opposition decides whether to mobilize its army to fight the incumbent. Third, the government decides whether to mobilize its own army to fight the opposition. We assume that both the opposition and the incumbent only mobilize if the net expected returns to mobilization are strictly positive, and an internal conflict takes place if at least one party mobilizes an army. Fourth, these choices and the probabilistic conflict technology then determine who wins power. Finally, the winner allocates the resources R .

Given our assumptions, the expected per capita payoff to incumbent members is:

$w(1 - \delta^I) + [\frac{1}{2} - (1/\mu)(\delta^O - \delta^I)(1 - 2\theta)]2R$, where the first term wages net of taxes, and the second is the expected size of the transfer. Similarly, the expected payoff to opposition members is: $w(1 - \delta^O) + [\frac{1}{2} + (1/\mu)(\delta^O - \delta^I)(1 - 2\theta)]2R$.

¹⁰ Besley and Persson (2009) study repression as one-sided violence by an incumbent, which has lower opportunity cost of fighting than the opposition since he can finance part of his army by taxing that opposition. In our empirical analysis (Subsection 5.3) we therefore examine the possibility of repression in the aftermath of giant oil discoveries.

¹¹ Besley and Persson (2009) do not focus on oil discoveries but on rents in general. Our assumption of two states of the world – with and without giant oil discoveries – makes the model more closely related to our empirical analysis.

To solve for the equilibrium we identify the sub-game perfect Nash equilibrium in the sequential game where the opposition moves first. It turns out that this game has two equilibria:

Peace (when neither side mobilizes): $\delta^O = \delta^I = 0$, which occurs when $2R(1-2\theta)/w \leq \mu$.

Internal conflict (when both sides mobilize): $\delta^I = A^I$ and $\delta^O = A^O$, which occurs when $2R(1-2\theta)/w > \mu$.

This model guides our empirical analysis of the effect of giant oilfields in a number of ways. First, the model assumes that giant oilfield discoveries increase oil revenues. While this assumption seems very plausible, it may take time to start generating revenues from newly discovered oilfields, especially if it is difficult to extract the oil or if the discovering country lacks the appropriate technology, capital, or infrastructure. Our first empirical challenge is therefore to determine whether oil production and oil exports increase significantly within a few years of discovery, and if so – by how much.

Second, we investigate the effect of oil discovery on internal armed conflict. The model predicts that in countries where $2R^H(1-2\theta)/w > \mu$, the discovery of a giant oilfield ends peace and sets off an internal conflict. This can happen when the incumbent receives most of the oil revenues, and cannot commit to sharing them with the opposition. If conditions are otherwise ripe for conflict, a giant oilfield discovery can fuel conflict over the oil revenues.

Third, giant oilfield discoveries are likely to set off conflict only in countries where political violence is seen as effective, namely where μ is sufficiently low. Empirically, we identify countries with low μ as those with a history of internal conflicts or coups in the decade prior to the discovery of a giant oilfield. It is in those countries that we expect giant oilfield discoveries to trigger armed conflicts over the control of the oil revenues. The model also allows for the possibility that giant oilfield discoveries fuel conflicts in countries with low wages (which imply a low opportunity cost of fighting), poor institutions that increase inequality (θ close to zero represents “winner takes all” societies, where it pays to fight for control), or ethnic fractionalization that creates conflicting groups to begin with. In practice, however, underlying characteristics such as income, institutions, and ethnic fractionalization may be interrelated with each other and with the degree to which political violence pays off (μ). In the empirical analysis below (Section 5) we focus on the interaction of giant oilfield discoveries with empirical measures of μ , but we also examine other possible interactions related to the model.

Finally, the discovery of a giant oilfield increases government revenues, R , and total per capita GDP, $R + w$. The increase of log per capita GDP in this model is $\partial \ln(R + w) / \partial \ln(R) =$

$R/(R + w)$, or in other words the proportional increase in GDP as a result of an oilfield discovery is less than the proportional increase in oil revenues as a result of this discovery. Moreover, as we discussed, in some cases oil discoveries cause mobilization, and this may reduce civilian per capita GDP. Any additional factors which are not modeled, such as the cost of conflict or any distortionary effect of oil on the rest of the economy, may further reduce the net benefits of giant oilfield discovery. Given these caveats, Subsection 5.3 investigates the effect of giant oil discoveries on GDP and its components.

4. Data on Oil, Conflicts, and Economic Outcomes

To analyze the effects of giant oilfield discoveries we require panel data on the timing of these discoveries in addition to outcome measures and control variables. Since country definitions differ across datasets, we use the country definitions from the Penn World Table, (Heston et al. 2009), a commonly used dataset, as the basis for our analysis.¹² The Penn World Table reports data on countries from 1950-2007, but since we examine all the conflicts that took place after the end of the Second World War (see below), some of the variables we match in from other datasets span the period from 1946-2008.¹³

Data on oil discovery and production. Our main regressor of interest is an indicator for the discovery of (at least one) giant oil field in a given country in a given year. We use data from Horn (2003, 2004), which reports the date of discovery, the name of the discovering country, and a number of other variables, for 910 giant oilfields discovered both onshore and offshore from 1868-2003. This dataset builds on previous datasets (e.g. Halbouty et al. 1970), and attempts to include every giant oilfield discovered around the world. To qualify as a giant (and thus be included in the dataset), an oilfield must have contained ultimate recoverable reserves (URR) of at least 500 million barrels of oil equivalent (MMOBE) before any oil was extracted.¹⁴ One limitation of these data is that the oilfields it describes differ considerably in the identity of those who estimated the URR and in the way the URR was estimated. Moreover, the estimated URR of various oilfields was gradually updated, depending on the

¹² We add three Communist countries which existed until the early 1990s: the USSR (until 1991), Yugoslavia (until 1991), and Czechoslovakia (until 1992); the countries that emerged from these three are covered in our dataset from the year following the corresponding collapse. We also add North Korea, Myanmar, and Netherlands Antilles. Our results are robust to excluding these countries.

¹³ We construct crosswalks for matching PWT countries to country names used in other datasets.

¹⁴ This estimate sums up oil-equivalent measures of oil, natural gas, and condensate.

estimators and their methods.¹⁵ Since this process may induce measurement error issues across oilfields, we simply construct an indicator for whether a country is mentioned in the dataset as having discovered at least one giant oilfield in each given year. This does not avoid all forms of measurement error, as some oilfields may have been incorrectly included in the dataset or excluded from it, but we consider this a reasonable compromise given the limitations of the data.

Of the 910 giant oilfields covered in Horn (2004), 782 were discovered from 1946 onwards, and these discoveries took place in 65 different countries. The 461 country-year observations with giant discoveries account for about 4.6 percent of all the observations in our data, so in all but a few countries giant oilfield discoveries are rare events (Table A1 lists the number of observations with discoveries in each discovering country). The rate of giant oilfield discoveries peaked during the 1960s and 1970s, and country-year pairs with discoveries were most common in Asia (41%), followed by Europe (18%), Africa (16%), North America (12%), South America (9%), and Oceania (4%).¹⁶ Table 1 reports summary statistics for our measure of giant oilfield discoveries and for other variables that we describe below.

We complement our data on giant oilfield discoveries with data on the timing of other oilfield discoveries from the Oil and Gas Journal Data Book (2008). This source reports more discoveries than our main dataset, since it is not limited to giant oilfield discoveries, but its main drawback is that the quantity of oil discovered is not reported for most oilfields. In addition, these data seem to focus on oil-producing fields, so they may exclude some gas fields.¹⁷ But these data are still useful, since they allow us to restrict parts of our analysis to observations with oil discoveries, and compare the effect of giant oilfield discoveries to discoveries of smaller fields.

Our final source for data on oil is Ross (2010), which reports the value of production of oil and gas by country and year from 1932 onwards.¹⁸ These data allow us to examine whether giant oilfield discoveries affect the value of oil and gas which a country produces. We

¹⁵ For example, some oilfields' URR was updated from an earlier version of the dataset we use (compare Horn 2004 and Horn 2003).

¹⁶ The continent classification follows that of the United Nations Statistical Division. The country-year distribution of discovery by decades is 3% for 1946-1949, 15% for 1950s, 22% for 1960s, 22% for 1970s, 14% for 1980s, 17% for 1990s, and 7% for 2000-2003.

¹⁷ Some fields covered in Horn (2004) do not appear in the Oil and Gas Journal Data Book (2008), even though this latter source covers smaller fields, so it reports more fields overall. This may be because the coverage of the Oil and Gas Journal Data Book is uneven across countries, whereas Horn (2004) attempts to cover all giant oilfields in all countries.

¹⁸ Details of data construction can be found in Ross (2010).

convert this variable into US\$2005, in line with our other variables below, using the CPI index from US Bureau of Labor Statistics.

Data on economic outcomes. The Penn World Table (PWT 6.3, 2009) is our source for GDP-related measures and population from 1950-2007. We use this dataset to construct Purchasing Power Parity (PPP) adjusted per-capita GDP in constant US\$2005, and to decompose it into private consumption, private investment, and government expenditure.¹⁹ We also construct a measure of the real exchange by taking the ratio of the nominal exchange rate (XRAT, which measures dollars per local currency unit) to PPP. Using this definition, a decrease in the real exchange rate corresponds to a real exchange rate appreciation. In addition, we supplement the PWT data with International Monetary Fund (IMF) data (Abbas et al. 2010) on public debt as a percentage of GDP.

To measure countries' international trade, we use the NBER-UN trade data (Feenstra et al. 2005), which reports trade outcomes from 1962-2000. We construct per capita measures of oil exports and non-oil exports. This last measure is constructed by summing up the exports in SITC Revision 2 categories 33 (Petroleum, petroleum products and related materials) and 34 (Gas, natural and manufactured). We convert all these measures into US\$2005 as described above.

Data on political violence. We use UCDP/PRIO dataset (Gleditsch et al. 2002) to measure the incidence of internal armed conflicts from 1946-2008.²⁰ One of our main outcomes of interest is an indicator for whether a given country experiences an internal conflict, which claims the lives of 25 people or more, in each given year. About 10 percent of our country-year observations involve such conflicts, and these conflicts take place in 97 different countries. Almost half of the internal conflicts in our data took place during the 1980s and 1990s, and the continent with the most conflicts was Asia (47% of conflict observations), followed by Africa (33%), South America (8%), North America (7%), Europe (6%), and Oceania (1%).²¹

¹⁹ PPP-adjusted GDP per capita is constructed using *rgdpl* (real GDP per capita, Laspeyres) and the components of GDP are constructed by multiplying each share, *kc* (private consumption), *ki* (private investment), *kg* (government spending), to *rgdpl*. All variables are from PWT 6.3.

²⁰ Conflicts are classified into four types in the UCDP/PRIO dataset: interstate, internal, internationalized internal, and extra-systemic (conflicts between a state and a non-state group outside its territory). Our main incidence measure is constructed using internal conflicts, but we consider others below.

²¹ The country-year distribution of conflict incidence by decades is 2% for 1946-1949, 5% for 1950s, 10% for 1960s, 15% for 1970s, 22% for 1980s, 27% for 1990s, 17% for 2000-2008.

For our robustness checks, we construct five other measures of internal armed conflict. The first is an intensity-scaled measure of internal armed conflicts, which takes on the value of one if the internal conflict's death toll in a given year was 25-999, two if it was 1000 or more, and zero otherwise. The second is an indicator for having either an internal or an internationalized internal conflict, since conflicts may switch from one type to the other. The third is an indicator for having any type of armed conflict (internal or not). The final two measures, following Collier and Hoeffler (2004) and Cotet and Tsui (2010), are an indicator for onset of internal armed conflict (having an internal conflict and no internal armed conflict in the preceding year) and a measure of internal armed conflict transitions (an indicator for an internal armed conflict in the current year minus the indicator for the previous year).

Another measure of political violence that we use is an indicator for having at least one coup in a given year, using data from the Polity IV project (Marshall and Marshall 2011). A coup is defined as a forceful seizure of executive authority and office by a dissident or opposition faction within the country's ruling or political elites that results in a substantial change in the executive leadership and the policies of the prior regime, or an attempt to do so (we do not distinguish between successful or unsuccessful coups). About 5.5 percent of the observations are classified as having at least one coup, and coups thus defined took place in 116 different countries from 1946-2008. Coups were fairly evenly distributed from the 1960s onwards (and rarer before), and the continent with the most country-year observations with coups is Africa (51%), followed by Asia (25%), South America (9%), North America (9%), Europe (5%), and Oceania (1%).²²

As an indicator for repression, we use the measure for purges from Banks (2010). This indicator takes on a value of one if a country experiences at least one purge in a given year, and zero otherwise. A purge is defined as systematic murders and eliminations of political opponents by incumbent regimes. About 8.6 percent of our observations are classified as having had at least one repression, and repression thus defined took place in 112 countries from 1946-2008. Repressions peaked during the beginning of the sample period - the 1940s and 1950s - and gradually declined over time.²³

We also follow Besley and Persson (2011) in constructing an indicator for countries with strong institutions. They use the fraction of time spent having the highest score for executive constraints variable (XCONST) from Polity IV project (Marshall et al. 2010) as the criterion

²²The country-year distribution of coup incidence by decades is 4% for 1946-1949, 7% for 1950s, 19% for 1960s, 20% for 1970s, 20% for 1980s, 18% for 1990s, 12% for 2000-2008.

²³ The country-year distribution of repression incidence by decades is 32% for 1946-1949, 30% for 1950s, 17% for 1960s, 11% for 1970s, 3% for 1980s, 2% for 1990s, and 1% for 2000-2008.

for having strong institutions.²⁴ In our analysis, we also use the Polity 2 score from the Polity IV project as a measure for institutional quality. This is a common measure of a country's political institutions, taking on values from -10 (strongly autocratic) to 10 (strongly democratic). Finally, we use the ethnic fractionalization measure from Alesina et al. (2003).

5. Results

This section begins by discussing our baseline empirical specifications and estimates of the effect of giant oilfield discoveries on oil production and exports and on internal armed conflicts (Subsection 5.1). We then discuss the robustness of our estimates using a number of alternative specifications (Subsection 5.2). We conclude this section by discussing the estimates of giant oilfield discoveries on other economic and political outcomes (Subsection 5.3).

5.1. Baseline Specifications and Results

In order to examine the effect of giant oilfield discoveries, we use our panel data on countries over time to estimate the following specification:

$$Y_{it+j} = \beta_{1j} \text{Disc}_{it} + \text{Country}_i + \text{Year}_t + \varepsilon_{it}, \quad (1)$$

where Y_{it+j} is the outcome in country i in year $t+j$, Disc_{it} is an indicator for the discovery of a giant oilfield in country i in year t , Country_i and Year_t are country and year fixed effects, and ε_{it} is a stochastic error. We estimate this specification for different lags j , where in most cases for $j \in \{2, 4, 6, 8, 10\}$. This allows us to non-parametrically trace the effect of discovery on outcomes over a decade. In addition, some of our specifications add controls, as explained in below and in the various tables.

As we discuss later on, we find that after controlling for country and year fixed effects, the timing of giant oilfield discoveries is largely uncorrelated with countries' economic and political performance in the five preceding years. One notable exception to this, however, is

²⁴ Details can be found in Besley and Persson (2011, pp. 1430-1431). There are 26 countries they define as having strong institutions: Australia, Austria, Belgium, Canada, Costa Rica, Denmark, Estonia, Finland, Germany, Greece, Iceland, Ireland, Israel, Italy, Japan, Latvia, Lithuania, Luxembourg, Netherlands, New Zealand, Norway, Sweden, Switzerland, South Africa, United Kingdom, and the United States.

that giant oilfield discoveries in a country's recent past raise the odds of additional discoveries in its near future. Specifically, we find that the unconditional probability of a giant discovery in year t rise from about 1 percent when there were no giant oilfield discoveries from $t-10$ to $t-1$ to 87 percent if there was a giant discovery in every year from $t-10$ to $t-1$. Controlling for country and (year fixed effects) significantly reduces the predictive power of past discoveries, though it remains statistically significant. In a regression where the dependent variable is $Disc_{it}$ and the regressor of interest is the number of years with giant oilfield discoveries from $t-10$ to $t-1$, controlling for country and year fixed effects, the estimated coefficient is 0.032 (s.e. 0.004).²⁵

These results suggest that giant oilfield discoveries in a country's recent past have some predictive power for whether it discovers yet another giant. We account for this serial correlation in the timing of giant oilfield discoveries by repeating our estimates of specification (1) with another specification, which we call (1a), and which includes the number of years with giant oilfield discoveries from $t-10$ to $t-1$ (labeled $PDisc_{it}$) as a control. In addition to reporting estimates from specifications (1) and (1a) in tables, we also plot the regression estimates and 95 percent confidence intervals for $j \in \{-5, -4, \dots, 9, 10\}$ using figures. These figures allow us to economize on space when we examine whether pre-discovery years differed from post-discovery years in terms of the outcomes of interest and to display outcomes in years $t+j$ where j is zero or positive and odd.

Having specified our baseline estimation strategy, we now examine the model's first prediction (or rather its assumption) that giant oilfield discoveries increase oil revenues. As Panel A of Table 2 shows, oil production increases by about 25-30 percent within two years of a giant discovery. This effect of discovery on oil production rises to about 35-45 percent within four years, and remains stable (at least) until 10 years after discovery. Past discoveries also matter, and their effect declines from about 26 percent two years after discovery to about 16 percent ten years after discovery. These large and precise estimates confirm that giant oilfield discoveries have an important economic impact on the discovering countries, as we can expect from the sheer size of these fields.

Panel B of Table 2 reports the effect of giant oilfield discoveries on oil exports. These discoveries raise oil exports by about 20-30 percent after six years, and this increases to about 40-50 percent after ten years. Past discoveries again matter, increasing oil exports by about

²⁵ One implication of this is that part of the effect of giant discoveries on subsequent outcomes may operate through a (slightly) increased change of making further discoveries.

10-20 percent. These estimates are similar to the effects we find on oil production, although oil exports appear to take a bit longer to respond to giant discoveries.

Our finding that giant oilfield discoveries increase per capita production and exports of oil lead us to investigate the second prediction of the model, that internal armed conflict increase after giant oilfields are discovered. Panel A of Table 3 documents the effect of giant oilfield discoveries on the subsequent incidence of internal armed conflict. In line with the second prediction of the model, we find that giant oilfield discoveries increase the incidence of internal armed conflict by about 5-8 percentage points during the years $t+4$ to $t+8$. This is a large effect, since the mean of this variable is just 10 percent, as we report in Table 1. Table 3 also shows that our estimates of the effect of giant oilfield discoveries on internal armed conflict do not change much when we control for the number of discoveries in $t-10$ to $t-1$.

Our finding that giant oilfield discoveries increase the incidence of internal conflicts is also not sensitive to the way these conflicts are measured. For example, Panel B of Table 3 shows that giant oilfield discoveries have a similar impact on an intensity-scaled measure of armed conflict, which gives more weight to conflict years with 1,000 casualties or more, as described in Section 4. Giant discoveries increase this scaled measure of internal conflict by about 6-9 percentage points within 4-6 years (or 4-8 years for specification 1a). Panel C of Table 3 shows an increase of 5-8 percentage points within the same lags from discovery year in a measure of conflict, which includes both internal and internationalized internal armed conflicts.²⁶

Given this evidence that giant oilfield discoveries increase the incidence of internal armed conflict, we now ask: which countries are particularly likely to experience internal conflicts after giant oilfield discoveries? According to Prediction 3 of our model, armed conflict over the oil revenues is prevalent in countries where political violence pays off. In order to identify these conflict-prone countries, we use past violence as an indicator. As Panel A of Table 4 shows, countries that experienced one or more coups from $t-10$ to $t-1$ were more likely to plunge into internal conflict following discoveries. In fact, in these countries giant oilfield discoveries raised the incidence of an internal conflict by as much as 11-14 percentage points from $t+4$ to $t+8$. This figure is high, but we should bear in mind that the mean incidence of an internal conflict following a coup is about 19-20 percent. By contrast, for countries that

²⁶ We also estimate similar regressions for the onset of internal armed conflict and for internal armed conflict transitions, find no contemporaneous relation between these and giant oilfield discoveries, a finding that is similar to Cotet and Tsui (2010). These measures increase significantly 4 years after discovery.

experienced no coups from $t-10$ to $t-1$, oil discoveries have no significant effect on the incidence of internal conflict, again consistent with the model's predictions.

Another indicator that political violence pays off is that a country already experienced internal conflict at some point from $t-10$ to $t-1$. In those countries, giant oilfield discoveries raise the probability of conflict by as much as 11-18 percentage points. The baseline level of violence in these cases is also very high, with a mean about 37-39 percent. Panel D of the table shows that we again find no significant effect of giant oilfield discoveries in countries with little previous history of violence.

5.2. Robustness of Our Main Results

The results discussed so far suggest that giant oilfield discoveries increase oil output and the incidence of internal armed conflict, and that the latter increase particularly worrisome for countries with recent histories of violence. We now examine the robustness of each of these results, and we start by looking at what happens in the years leading up to giant oilfield discoveries.

Sub-Figure A of Figure 1 shows estimates of specification (1) for our measure of oil production before and after discovery. The figure suggests that oil production did not change much in the five years leading up to giant oilfield discoveries. Similarly, Sub-Figure B of Figure 1 shows that oil exports also did not increase during the lead-up to the discovery of giant oilfields, and we can again see that oil exports took longer to respond to giant oilfield discoveries than oil production. Sub-figure C of Figure 1 shows that internal armed conflict also did not change significantly in the years leading up to giant discovery.²⁷ Finally, Sub-Figure D of Figure 1 shows that in countries that experienced at least one internal armed conflict from $t-10$ to $t-1$, conflicts did not systematically flare up in the years prior to giant oilfield discoveries.

The four sub-figures of Figure 2 similarly show estimates to the corresponding sub-figures of Figure 1, except that this time we control for the number of years with giant oilfield discoveries from $t-10$ to $t-1$ (these estimates are generated using specification (1a) instead of

²⁷ Another way to examine whether giant oilfield discoveries occur during periods of lull following conflict is to estimate a regression where the dependent variable is an indicator for giant discovery and the regressor of interest is an indicator for having no internal armed conflict in periods $t-j$ to $t-1$ and conflict in period $t-j-1$, for $j \in \{1, 2, 3, 4, 5\}$, controlling for country and year fixed effects. The coefficient of interest in all these regressions is small and imprecise, providing further evidence that lulls in fighting do not predict the timing of giant oilfield discoveries.

specification 1). These results once again show no evidence of significant trends before discovery. Moreover, the estimates are quantitatively very similar to those in Figure 1. From this point on, to save on space, we discuss only estimates that control for discoveries before t , as in specification (1a).

In Figure 3 we examine the changes before and after discovery in some of the alternative measures of conflict discussed above. Sub-Figures A and B of this figure correspond to Panels B and C of Table 3, using as outcomes internal armed conflicts scaled by intensity and internal armed conflict including ones that were internationalized. These outcomes, like our main measure of armed conflict, show little change in the years leading up to discovery, and become positive and significant within 4-8 years after discovery. Sub-Figure C of the figure corresponds to panel A of Table 4, showing that in countries that had at least one coup from $t-10$ to $t-1$, internal armed conflicts increase more with giant oilfield discoveries, and there were no significant changes in the years leading up to discovery. Sub-Figure D shows similar results for countries that experienced any type of armed conflict from $t-10$ to $t-1$.

The finding that our key variables of interest do not change systematically in the years leading up to giant oilfield discoveries is consistent with our interpretation that the estimates what takes place following giant oilfield discoveries reflects these discoveries' effect on the discovering countries. In the following paragraphs we address further potential concerns about this interpretation.

One concern that may linger, for example, is that there might be serial correlation not only in the timing of giant oilfield discoveries but also in the outcomes we examine. To address this concern, panel B of Table 5 re-estimates specification (1a), but this time controlling for the dependent variable in $t-1$, which is instrumented by the dependent variable in $t-2$. The outcome here is our measure of oil production, and the estimates are smaller than the baseline, but still positive and significant.

Another related concern is that political conditions in the discovering country may have changed shortly before discovery. But Panel C of the table adds to specification (1a) a control for polity 2 (a common measure of institutional quality) in $t-1$ and this does not change the estimates much. Since we do not have a measure of investment in the oil sector, Panel D reports estimates of specification (1a) with a control for log PPP-adjusted per capita private investment in 2005 US dollar in $t-1$, and again the estimates remain statistically significant. Panel E adds together all the controls from Panels B-D, and again the estimates remain significant for $t+2$ through to $t+10$, this time with the exception of the estimate for $t+8$, which is marginally significant.

While the results discussed so far include all discoveries of giant oilfields since 1946 and control for discoveries in countries' recent past, a concern remains that the odds of discovery are not the same in all countries and in all years. More specifically, the regressions discussed so far include country-year observations where the probability of discovery is relatively high given the history of past discovery, along with many (most) observations where the odds of discovery were low. Panel F of Table 5 focuses on giant oilfield discoveries that were especially surprising, since no giant oilfield was discovered in the country from $t-10$ to $t-1$. When we exclude those observations, giant oilfield discoveries are made in just over 1 percent of the remaining observations, so these discoveries were in all likelihood highly unexpected. The results show that the effect of these unexpected discoveries on oil production are actually about twice as large as in the baseline, and precisely estimated. This is probably because in the countries that make these discoveries, oil production prior to the giant discovery was usually very low.

Another potential concern regarding our identification is that the countries that discover giant oilfields differ from others in ways that change over time and are therefore not fully controlled for by country fixed effects. To address this concern, Panel G re-estimates the baseline specification using only countries that make at least one giant discovery in the period we analyze (from 1946 onwards). The estimates in this specification are similar to those in the baseline, but slightly larger.

Finally, we address the concern that country-year observations with oil discoveries are somehow different from others not only across countries, but also within countries, and in ways that we cannot observe and control for directly. To mitigate this concern, we use data from the Oil and Gas Journal Data Book (2008), which records country-year pairs where some oil discoveries, not necessarily giant, were made. Estimating specification (1a) using only these country-year observations, we essentially compare observations with (and following) giant oilfield discoveries to years with (and following) smaller discoveries. As Panel H of the table shows, even using this variation alone, the effect of giant oilfield discoveries on oil production is positive and significant, albeit smaller, for $t+6$ to $t+10$.

Table 6 repeats the robustness checks described above for our main result, that giant oilfield discoveries increase the probability of internal armed conflict from $t+4$ to $t+8$. Controlling for the (instrumented) lagged dependent variable, lagged polity 2 score and lagged investment, or all of these together, tends to increase the coefficients very slightly, and they remain statistically significant. Excluding observations that follow one or more discoveries in $t-10$ to $t-1$ makes the estimate for $t+4$ imprecise, but the coefficients for $t+6$ and $t+8$ are still precise

– the latter is even larger than in the baseline specification. Restricting the sample to countries with giant oilfield discoveries leaves the baseline coefficients almost unchanged. And using only observations with some oil discoveries tends to increase both the point estimates and the standard errors, leaving the estimates for t+4 and t+6 positive and statistically significant.

Table 7 reports estimates for the same robustness checks as before, but this time for the effect of giant oilfield discoveries on internal armed conflicts in countries that experienced at least one year of conflict from t-10 to t-1. As before the controls we include make little difference to our estimates when they are included separately or simultaneously: the estimates for t+4 to t+8 remain significant and change little in magnitude. Excluding observations with recent past discoveries makes the estimate for t+4 imprecise, but the estimates for t+6 and t+8 are still precise. Restricting our analysis to the set of countries with giant oilfield discoveries again makes almost no difference relative to the baseline. Finally, using only observations with some oil discoveries, while restricting our sample to about 400 observations, still gives positive and significant estimates for t+4 and t+6.

Our regressions so far have emphasized that the effect of oil discovery on conflict are larger in countries with a history of conflict. We now compare the interaction of discovery and recent conflicts with interactions of discovery with other features of the discovering country. To do so, we begin by estimating the following equation:

$$Y_{it+j} = \beta_{2j}Disc_{it} + \gamma_{2j}PConf_{it} + \delta_{2j}PDisc_{it} + \theta_{2j}Disc_{it} \times PConf_{it} + Country_i + Year_t + \varepsilon_{it}, \quad (2)$$

where $PConf_{it}$ measures the number of years from t-10 to t-1 in which country i experienced internal armed conflict. Panel A of Table 8 reports estimates of β_{2j} and θ_{2j} for $j \in \{2, 4, 6, 8, 10\}$. As the table shows, θ_{2j} is positive and significant for 4, 6, 8, and even 10 years after discovery, confirming again that oil discoveries spell trouble in countries with recent histories of violence.

We now add to this specification interactions of discovery with other country characteristics, following our discussion in Section 3. First, we consider the possibility that in countries with strong institutions, giant oilfield discoveries lead to less conflict. To test this, we add to specification (2) an interaction of discovery with strong institutions, which may proxy for an

institutionalized commitment to share revenues with the opposition (θ close to $\frac{1}{2}$).²⁸ Second, much of the literature (see survey in Blattman and Miguel 2010) finds that conflicts are more prevalent in poor countries. In the model this corresponds to low-wage countries, and given our data limitations we proxy this using lagged per capita GDP. Specifically, we examine whether giant oilfield discoveries are more likely to tip poor countries into internal conflict by further adding to specification (2) controls for log per capita GDP in t-1 (as discussed in the data section) and its interaction with our indicator for discovery, $Disc_{it}$. Finally, we consider the possibility that in countries with higher ethnic fractionalization, oil discoveries are more likely to cause conflict, possibly because those countries are more prone to be divided into opposing factions that willing to fight each other. We test this hypothesis by further adding to specification (2) an interaction of our measure of ethnic fractionalization (again see data section) with $Disc_{it}$. Panel B of Table 8 shows that none of the three interactions we added is statistically significant in any of the regressions, while the interaction of discoveries and past conflict is still positive and significant from 4 years after discovery onwards. This suggests that the countries that should be most concerned about tipping into violent conflict over resources are those with recent histories of conflict, but not necessarily with other institutional features.

5.3. Additional Results

The results discussed so far suggest that giant oilfield discoveries may have two opposing effects on the discovering countries economy: they increase its revenues (from oil) and at the same time increase the incidence of a costly internal conflict. We now turn our attention to the fourth and last we discuss in the model section – whether these discoveries have a positive or a negative effect on per capita GDP and its components. For these outcomes, Table 9 reports estimates of specification (1a), and an augmented specification, which includes controls as in Panel E of Tables 5-7, namely the dependent variable in t-1 (instrumented by that same variable in t-2) and polity 2 and log PPP-adjusted per capita private investment, both also measured in t-1.

Panel A of Table 9 suggests that giant oilfield discoveries may increase per capita GDP by about 4-6 percent. But as Panel B shows, this estimate is imprecise when more controls are

²⁸ As discussed above, the definition of strong institutions indicator we use comes from Besley and Persson (2011). Since the indicator itself is time invariant, its main effect is absorbed by the country fixed effects.

added. Having also experimented with similar specifications with various controls, we conclude that the positive effect we find in Panel A is not very robust, so we are unable to say conclusively whether giant oilfield discoveries have a small positive effect on per capita GDP, or whether this effect is zero.

The next two panels of Table 9 show similar results for the effect of giant oilfield discoveries on per capita government spending. Once again the effect is either significant and around 4-6 percent (Panel C), or insignificantly different from zero (Panel D) when more controls are added.

The rest of Table 9 shows that giant oilfield discoveries have no significant effect on per capita private consumption and (with the exception of one negative estimate for t+6 in panel H), also no effect on per capita private investment.²⁹

We conclude this section with an investigation of other possible economic and political consequences of giant oilfield discoveries in Table 10. In Panels A and B of this table we test one of the mechanisms often discussed in the “Dutch Disease” literature (e.g. Corden and Neary 1982), that natural resource booms may cause a real exchange rate appreciation. This may happen, for example, if an oil-producing country exports more than it imports, spending some of its revenues on local non-tradable goods. As a result of such spending, the nominal exchange rate may appreciate (if the exchange rate is flexible) or local prices may rise. Either (or both) of these can cause a real exchange rate appreciation, which can hurt the non-oil exporting industries. Panels A and B of Table 10, however, show that giant oilfield discoveries decrease the real exchange rate only for some years after discovery, and even then the effect is quite small and imprecisely estimated. Panels C and D similarly show that non-oil exports are not significantly reduced by giant oilfield discoveries. A more thorough investigation of various related “Dutch Disease” mechanisms is, however, outside the scope of this paper and we leave it for future work.

We next examine an alternative hypothesis on a potential cost of oil production, namely that it may lead, in some cases, to over-spending and indebtedness by the government (for related discussions see Tornell and Lane 1999 and Manzano and Rigobón 2008). As Panels E and F of Table 10 show, we find no evidence to support this hypothesis for our full sample of countries. Again, we leave further investigations of this issue for particular countries or regions for future work.

²⁹ We do not report figures for the effect of giant oilfield discoveries on per capita GDP and its components, but these are available on request from the authors, and they also suggest that the changes before and around discoveries are small and imprecisely estimated.

Turning to other political economy hypotheses on the effect of natural resources, Panels G and H of Table 10 examine whether competition over oil revenue takes the form of coups to replace the incumbent. As the table shows, we find no evidence that giant oilfield discoveries increase the odds of coups in the subsequent decade. Finally, the last two panels of Table 10 test the prediction of Besley and Persson (2009, 2011), that resource windfalls increase repression. Our estimates, however, show no significant increase in repression in the aftermath of giant oil discoveries.

In sum, the results discussed in this subsection suggest that while the economic gains from giant oilfield discoveries to the local population may be limited, we do not identify other costs to from discovery, except for our main result of an increased likelihood of internal armed conflict.

6. Conclusions

We began this paper by asking whether natural resource windfalls fuel internal armed conflicts, and if so – in which settings. To answer this question, we use new data on giant oilfield discoveries to identify the effect of oil on economic and political outcomes around the world. We find that within a few years of giant oilfield discoveries, per capita oil production and oil exports in discovering countries increase by up to 50 percent. But we also find that discovering giant oilfields increases the incidence of internal armed conflict by about 5-8 percentage points. This increase is driven predominantly by countries with recent histories of political violence – those that experienced coups or armed conflicts during the decade prior to discovery. We show that these findings are robust to a number of different specification checks and the addition of various controls.

Our findings shed light on the questions we began with. Giant oil and gas field discoveries in Norway, Canada, and Australia, are unlikely to fuel internal armed conflicts, since these countries' recent histories include little political violence. But in countries where political disputes are often resolved by violence (or remain unresolved despite violence), giant oilfield discoveries can fuel the flames of internal conflicts.

Our finding that giant oilfield discoveries fuel internal conflicts in countries that are prone to violence also has implications for policy. Those who strive to reduce armed conflict should be concerned about oil rents that incumbents obtain in conflict-prone areas, especially if those rents encourage challenges to the incumbents' power. At the same time, the firms that

prospect for oil in conflict-prone areas and those who regulate them ought to be concerned about negative externalities for many locals, who have little to gain from giant oilfield discoveries but may suffer from conflicts over the oil.

References

- Abbas, S.M. Ali, Nazim Belhocine, Asmaa El-Ganainy, and Mark Horton, (2010) "A Historical Public Debt Database", IMF Working Paper WP/10/245, Washington, DC.
- Acemoglu, Daron, Ticchi, Davide, and Vindigni, Andrea (2010) "A Theory of Military Dictatorships," *American Economic Journal: Macroeconomics* 2010, 2:1, 1–42
- Alesina, Alberto, Arnaud Devleeschauwer, William Easterly, Sergio Kurlat, and Romain Wacziarg. (2003) "Fractionalization". *Journal of Economic Growth* 8 (June): 155-194.
- Angrist, Joshua D. and Kugler, Adriana D. (2008) "Rural Windfall or a New Resource Curse? Coca, Income, and Civil Conflict in Colombia," *Review of Economics and Statistics*, vol. 90(2), pages 191-215, 03.
- Banks, Arthur S. 2010. Cross-National Time-Series Data Archive. Databanks International. Jerusalem, Israel; see <http://www.databanksinternational.com>
- Bellows, John, and Miguel, Edward (2009). "War and Local Collective Action in Sierra Leone." *Journal of Public Economics*, 93(11–12): 144–57.
- Besley, Timothy and Persson, Torsten (2009). "Repression or Civil War?," *American Economic Review*, vol. 99(2), pages 292-97, May.
- Besley, Timothy and Persson, Torsten (2011). "The Logic of Political Violence," *Quarterly Journal of Economics*, forthcoming
- Blattman, Christopher and Miguel, Edward (2010). "Civil War," *Journal of Economic Literature*, vol. 48(1), pages 3-57, March.
- Caselli, Francesco, and Michaels, Guy (2009) "Do Oil Windfalls Improve Living Standards? Evidence from Brazil," NBER Working Papers 15550.
- Collier, Paul and Hoeffler, Anke, (1998) "On Economic Causes of Civil War," *Oxford Economic Papers*, Oxford University Press, vol. 50(4), pages 563-73, October.
- Collier, Paul and Hoeffler, Anke (2004) "Greed and grievance in civil war," *Oxford Economic Papers*, Oxford University Press, vol. 56(4), pages 563-595, October.

- Corden, W Max and Neary, J Peter, (1982) "Booming Sector and De-Industrialisation in a Small Open Economy," *Economic Journal*, Royal Economic Society, vol. 92(368), pages 825-48, December.
- Cotet, Anca and Tsui, Kevin K. (2010) "Oil and Conflict: What Does the Cross-Country Evidence Really Show?," Working Papers 201002, Ball State University, Department of Economics, revised Mar 2010.
- Dal Bó, Ernesto, and Dal Bó, Pedro (2011). "Workers, Warriors and Criminals: Social Conflict in General Equilibrium," *Journal of the European Economic Association*, forthcoming
- Dube, Oeindrila and Vargas, Juan. (2010). "Commodity Price Shocks and Civil Conflict: Evidence from Colombia", unpublished working paper, NYU.
- Feenstra et al. (2005), World Trade Flows: 1962-2000. NBER Working Paper No. W11040 <http://www.internationaldata.org/data/undata/undata.html>.
- Garfinkel, Michelle R. and Skaperdas, Stergios "Economics of Conflict: An Overview," in T. Sandler and K. Hartley (eds.), *Handbook of Defense Economics*, Vol. II, 2007, 649-709.
- Gleditsch, Nils Petter, Peter Wallensteen, Mikael Eriksson, Margareta Sollenberg, and Håvard Strand. 2002. "Armed Conflict 1946-2001: A New Dataset." *Journal of Peace Research* 39(5).
- Guidolin, Massimo and La Ferrara, Eliana (2007). "Diamonds Are Forever, Wars Are Not: Is Conflict Bad for Private Firms?," *American Economic Review*, vol. 97(5), pages 1978-1993, December.
- Halbouty, Michel T., A. Meyerhoff, Robert E. King, Robert H. Dott, Douglas Klemme and Theodore Shabad, 1970, World's Giant Oil and Gas Fields, Geologic Factors Affecting Their Formation, and Basin Classification: Part I: Giant Oil and Gas Fields. in AAPG Memoir 14: Geology of Giant Petroleum Fields, p. 502-528
- Heston, Alan, Robert Summers, and Bettina Aten Penn World Table, Version 6.3, Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania, August 2009.
- Horn, Myron K. (2007). "GIANT FIELD TRENDS-2: Giant fields likely to supply 40%+ of world's oil and gas." *Oil and Gas Journal* 105.13
- Horn, Myron K. (2003) "Giant Fields 1868-2003 (CD-ROM)." In Halbouty, M.K., ed., *Giant Oil and Gas Fields of the Decade, 1990-1999*. Houston: AAPG Memoir 78, 2003.

- Horn, Myron K. (2004) "Giant Fields 1868-2004 (CD-ROM)." Revision to 2003 version. Houston: AAPG/Datapages Miscellaneous Data Series. Version 1.2, 2004.
- Humphreys, Macartan. 2005. "Natural Resources, Conflict, and Conflict Resolution: Uncovering the Mechanisms." *Journal of Conflict Resolution*, 49(4): 508–37.
- Manzano, Ozmel, and Rigobón, Roberto (2008). "Resource Curse Or Debt Overhang?" in Lederman, Daniel and Maloney, William F. (eds.) *Natural Resources Neither Curse Nor Destiny*. Stanford University Press and the World Bank.
- Marshall, Monty G., Gurr, Ted R., and Jaggers, Keith (2010), "Polity IV: Political Regime Characteristics and Transitions, 1800-2009", The Center for Systemic Peace.
- Marshall, Monty G. and Marshall, Donna R. (2011), "Coups d'Etat, 1946-2010" from Polity IV: Regime Authority Characteristics and Transitions Datasets, The Center for Systemic Peace.
- Michaels, Guy (2011) "The Long Term Consequences of Resource-Based Specialisation," *Economic Journal*, vol. 121(551), pages 31-57, March.
- Oil and Gas Journal Data Book (2008). Tulsa, Oklahoma: Petroleum Publishing Company.
- Reno, William. (1999) *Warlord Politics and African States*. Boulder: Lynne Rienner Publishers.
- Ross, Michael L. (2004) "What Do We Know about Natural Resources and Civil War?" *Journal of Peace Research*, 41(3): 337–56.
- Ross, Michael L. (2006) "A Closer Look at Oil, Diamonds, and Civil War." *Annual Review of Political Science*, 9: 265–300.
- Ross, Michael L. (2011) "Replication data for: Oil and Gas Production and Value, 1932-2009", <http://hdl.handle.net/1902.1/15828> UNF:5:Hwe3jAjxG7fgOMzpGQXOxw== V4 [Version]
- Tornell, Aaron and Lane, Philip R. Lane, (1999) "The Voracity Effect," *American Economic Review*, American Economic Association, vol. 89(1), pages 22-46, March.
- U.N. Statistics Division, "Composition of macro geographical (continental) regions, geographical sub-regions, and selected economic and other groupings" <http://unstats.un.org/unsd/methods/m49/m49regin.htm>
- U.S. Bureau of Labor Statistics, Consumer Price Index History Table, (<ftp://ftp.bls.gov/pub/special.requests/cpi/cpiiai.txt>).