Does Oil Cause Civil War Because It Causes State Weakness?*

Adam N. Glynn†

Original Draft: October 22, 2008
This Draft: October 21, 2009

Abstract

Conflict scholars have devoted considerable attention to the natural resource curse, and specifically to connections between natural resources, state weakness, and civil war. Many have posited a state weakness mechanism– that significant oil production causes state weakness, and state weakness consequently increases the likelihood of civil war onset. Using standard measures, this paper demonstrates that the state weakness mechanism *does not exist* in the short or medium term. The methods developed in this paper show that in only two cases is there the possibility of a medium term effect, and the state weakness mechanism is unlikely to be operative even in these two cases. Furthermore, these methods do not rely on assumptions about unmeasured confounders, so this result is robust to the consideration of other risk factors for civil war onset. The state weakness mechanism may still exist in the form of long term effects or an effect that reinforces pre-existing war and/or state weakness. However, the null hypothesis of no long-term and/or reinforcing effect cannot be rejected without the use of additional assumptions.

*An earlier version of the results in this paper can be found in a conference paper presented at the 2008 Summer Meeting of the Society for Political Methodology. The author thanks Kevin Quinn, Nahomi Ichino, Oliver Bevan, Patrick Egan, Gary King, Jeff Frieden, Jorge Dominguez, Torben Iversen, Michael Hiscox, Kenneth Shepsle, Daniel Ziblatt, John Patty, Maggie Penn, Peter Hall, and the participants of the Summer Meeting for their helpful comments and discussion. The usual caveats apply.

†Department of Government and The Institute for Quantitative Social Sciences Harvard University, 1737 Cambridge Street, Cambridge, MA 02138. aglynn@fas.harvard.edu

1 Introduction

Does oil wealth lead to the onset of civil conflict? Ross (2004) notes a robust correlation between these variables across a number of observational studies, but many plausible explanations exist. One explanation consistent with this correlation is that countries with oil wealth have less incentive to build state capacity– "a political 'Dutch Disease'," and the resulting lower state capacity increases the likelihood of civil war (Fearon and Laitin, 2003; Fearon, 2005). Furthermore, the two causal effects comprising this state weakness mechanism have appeared in various guises throughout the literature on the natural resource curse (Chaudhry, 1997; Karl, 1997; Moore, 2001; Wantchekon, 2002; Sørli, 2002; Humphreys, 2005).

Although the state weakness mechanism is theoretically compelling, the empirical evidence for this mechanism is limited. Fearon (2005) regresses a proxy for state capacity on a measure of fuel exports and finds a negative relationship. Humphreys (2005) regresses civil war onset on interacted measures of oil production and lagged measures of state weakness, and finds a positive relationship between state weakness and civil war onset at all positive levels of oil production.[1] However, as these studies acknowledge, state capacity is difficult to measure, and these results only suggest the possibility of a state weakness mechanism (each study separately examines one of the two constituent causal effects).

Moreover, the likelihood of heterogeneous causal effects makes the combination of such evidence potentially misleading. Even if we could use randomized experiments to establish the two constituent effects (an effect of oil on state capacity and an effect of state capacity on the likelihood of civil war), this would not be sufficient to establish an indirect effect (of oil on civil war through state capacity) when causal effects are heterogeneous (Robins and Greenland, 1992).

The following stylized example demonstrates the insufficiency of regression results for the establishment of mechanisms. Suppose there are two types of countries. For the first type, oil production

[1] This result was implied by Table 3 of Humphreys (2005) but not discussed explicitly in the text of that paper. Also, due to the fact that the state weakness measures were lagged in this study, this result would be most accurately described as a controlled direct effect.

causes state weakness, but state weakness has no effect of the likelihood of civil war. Therefore, for the first type of country, the first constituent cause is present but not the second, and there is no state weakness mechanism. For the second type of country, oil production has no effect on state weakness, but state weakness increases the likelihood of civil war. Therefore, for the second type of country, the second constituent cause is present but not the first, and there is also no state weakness mechanism. Furthermore, note that in this example, the average effect of oil production on state weakness is positive and the average effect of state weakness on the likelihood of civil war is positive. Therefore, this stylized example is consistent with regression evidence for the two constituent effects, even though there is no state weakness mechanism for any country. This insufficiency of regression evidence is important for our understanding of the natural resource curse, because the established existence of the weak states mechanism would provide evidence of oil as a causal risk factor distinct from other primary commodities,[2] and might justify the adoption of certain policies.[3]

However, while we will generally be unable to establish the exact magnitude of the indirect effect, we may still be able to learn *something* about this effect. Counterfactual causal models provide a framework for causal learning, and a number of studies have shown that implausible causal parameter estimates can be eliminated with bounding techniques (Manski, 1990; Balke and Pearl, 1997; Chickering and Pearl, 1997; Manski, 2003; Quinn, 2008; Cai et al., 2007). Furthermore, counterfactual analysis has a long history in comparative politics (Fearon, 1991; Lebow, 2000; Przeworski, 2004; Sekhon, 2004), and while it is well known that counterfactuals can be poorly defined (Quine, 1953; Dawid, 2000), they have been shown to be useful for the study of causal effects in both the medical and the social sciences (Lewis, 1973; Robins and Greenland, 2000; Morgan and Winship, 2007). Counterfactual causal models have been particularly effective in

[2] This would partially settle the debate between Fearon (2005) and Collier and Hoeffler (2004).

[3] For example, Humphreys (2005) suggests that oil revenue management laws might work by, "tackling the ways in which natural resource revenues weaken state structures." It is important to note that the absence of a weak states mechanism would not necessarily obviate all benefits of such policies, because the policies might have effects through other channels.

highlighting ambiguity in causal research questions.

In the following sections, I extend the bounding approach to a counterfactual model for causal mechanisms in order to assess the empirical evidence for the weak states mechanism. Section 2 discusses the use of counterfactual models to provide bounds for the short-term and medium-term effects of oil production on the probability of civil war onset. These total effects are depicted graphically by the reduced form model in Figure 1 (a). Section 3 presents a similar analysis for the total effects of oil on state weakness, and these effects are depicted graphically by the arrow between OIL and SW in Figure 1 (b). Section 4 presents the key results of this paper, providing a method for bounding mechanism-specific effects (depicted by the pathway from OIL through SW to WAR in Figure 1 (b)). Using standard measures for oil production, state weakness, and civil war, I find in Sections 2 and 3 that the bounds on the effects of OIL on WAR and OIL on SW are consistent with previously published results, but in Section 4, I find that *there is no evidence for a short-term state-weakness mechanism, and the evidence for a positive upper bound on the medium-term effect is based on only two cases. Furthermore, all assumptions made throughout these sections are chosen to make the upper bounds large.* In Section 5, a cursory examination of these two cases reveals that the state weakness mechanism is unlikely to be operative, and given this lack of evidence for a medium-term effect, I discuss the remaining evidence for the weak states pathway as a long term effect or an effect that reinforces previous war or weakness. Section 6 concludes.

2 Oil and Civil War in the Second Half of the 20th Century

In studying the effects of oil production on the onset of civil war, we might ask the following two basic questions, "Do countries experience an onset of civil war after they start producing significant quantities of oil?" If so, "How long does it take?" Of course, even these questions depend on the definition of "civil war" and "starting significant oil production," and in this paper I utilize the "significant oil" and "civil war" variables as coded in the seminal paper, Fearon and Laitin (2003).[4]

However, in order to establish causal effects of oil production on the onset of civil war, we must also ask the counterfactual question of whether countries that produced oil would have experienced a civil war if they had not produced oil.[5] In regression or matching methods, this counterfactual question is addressed by utilizing non-producing countries as proxies for the counterfactual outcomes that we would have observed for the oil producing countries (had they not produced oil). For example, in a one-to-one matching difference estimator, the oil producing countries are matched with non-producing countries that are similar on all other pre-oil production variables. A difference is taken between the civil war indicators for each matched pair, and these differences are averaged. Roughly speaking, the quality of the inference produced by these approaches depends primarily on the quality of the matches.

In the bounding approach, we utilize the dichotomous nature of the outcome variable to match each oil producing country with *all* possible values for the counterfactual outcomes that we *could* have observed for the oil producing countries (had they not produced oil). Hence we do not need to consider the non-producing countries in our analysis. This will limit the precision of our conclusions, but it is often a useful first step because it does not rely on the assumptions implicit in the regression and matching approaches (e.g. no unmeasured confounders). Furthermore, we will sometimes learn enough from the bounding approach to at least partially answer the question of interest. This will be demonstrated in Section 4.

If we index the 35 countries that were significant oil producers (at some point in the data set) with $i = 1, ..., 35$ (by convention we define the variable $OIL_i = 1$ for these 35 countries and $OIL_i = 0$ for the other countries in the data set), we can define the binary variables WAR_i^1 (which

[4]To summarize the definitions of the key variables, significant oil production is defined as fuel exports exceeding one-third of export revenues, and civil war is defined as a conflict between state and non-state agents that killed at least 1,000 individuals. While these measures are certainly not definitive, they are reasonable, well documented, and the deficiencies of these measures are well understood by subject matter experts. Furthermore, this choice allows the analysis in this paper to be connected more directly to previously reported results.

[5]We may also be interested in the question of whether countries that did not produce oil would have experienced a civil war if they had produced oil. However, this query has fewer policy implications, and I will omit this analysis in order to simplify the presentation.

takes the value 1 if the country experienced a civil war in its first year after first significant oil production)[6] and WAR_i^9 (which takes the value 1 if the country experienced a civil war sometime in its the first nine years after the first significant oil production). The restriction to the nine year window allows the inclusion of the former Soviet States, which had only nine years in the data set. All other countries with fewer than nine post-oil years experienced a civil war onset within their time in the data set.[7] Note that these WAR_i^1 and WAR_i^9 variables are well defined and observable for the oil producing countries in the data set.

However, in order to address the causal effect of oil on civil war, we must consider the more difficult question of whether the 35 countries would have experienced a civil war had they not been significant oil producers. We will denote the answer to this question with the counterfactual binary variables $WAR_i^1(0) \equiv WAR_i^1(OIL_i = 0)$ and $WAR_i^9(0) \equiv WAR_i^9(OIL_i = 0)$. It is important to note that we have assumed there are right answers to these counterfactual questions. This could be problematic because non-significant oil production can mean different things (e.g. a little oil vs. no oil at all), because a particular level of oil production can mean different things (e.g. less oil in the ground vs. less drilling), and because oil production in one country might affect civil war outcomes in another country (so that we would actually need to specify the oil production levels for all countries in our study before the counterfactual variables would be well defined). However, it is also important to note that all empirical studies of the effect of oil on civil war implicitly make some version of this assumption, and the paucity of conflict data requires that some modeling assumptions be made.[8] Furthermore, even when these counterfactual variables are well defined, we

[6] I have followed the convention of (Fearon and Laitin, 2003) by not lagging the oil production variable. Therefore, significant oil production and civil war onset in the same year will be interpreted as oil production preceding civil war onset. This choice can be seen as conservative, because this paper attempts to establish an upper bound for the weak states mechanism-specific effect.

[7] The only slight exception to this rule was the Yemen Arab Republic, which had only five post-oil years and no post-oil civil wars prior to its consolidation into Yemen. I have appended the first four years of Yemen's history to the end of the Yemen Arab Republic's history in order to achieve the nine post-oil years. The removal of either of these countries from the analysis does not change the substantive findings.

[8] For example, if oil production is measured in dollars per capita, then a parametric model will typically be necessary, and if we adjust the model to allow oil production in one country to affect civil war onset in other countries, we require some knowledge of the causal structure between countries. In either of these cases, the counterfactual statements become complex.

do not observe their values.

With the relevant counterfactuals specified, we can define country specific effects to be a contrast between the observed variables and the counterfactual variables: $WAR_i^1 - WAR_i^1(0)$ and $WAR_i^9 - WAR_i^9(0)$.[9] Notice that due to the binary coding of the war variables, these contrasts can only take the values $\{-1, 0, 1\}$. Table 1 shows the possibilities for the 35 oil producing countries, where the counterfactual variables and the country specific effect are not observed (the set notation $\{\cdot\}$ indicates the possible values in this table). However, notice that for the oil producing countries that did not experience war, the country specific effect can only take the values $\{-1, 0\}$. That is, oil production either prevented war, or had no effect, hence we can say that for these countries, oil "didn't cause" war. For the oil producing countries that experienced war, the country specific effect can only take the values $\{0, 1\}$, hence we can say that for these countries, oil "didn't prevent" war. This fundamental lack of identification is a well known problem for unit-specific causal inference, and often analysts resign themselves to the identification of average causal effects (note that even randomized experiments only identify average causal effects, not unit specific effects). Because we are primarily interested in the effects of oil for the 35 oil producing countries, this average can be restricted to an average over those countries "treated" with oil production, and in the causal inference and matching literature, this type of average effect is called the Sample Average Treatment effect on the Treated (SATT). SATT can be interpreted as the probability that oil production caused a civil war for a randomly sampled oil producing country among this sample.

Using the possible values for the country specific effects, we can derive bounds for SATT by averaging over the lower possible values for the lower bound, and averaging over the upper possible values for the upper bound. For the first year analysis, there were 31 countries with -1 as the lower possible value and 4 countries with 0 as the lower possible value, so the lower bound for SATT must be -31/35. Furthermore, those 31 countries had an upper possible value of 0 while the 4 countries have an upper possible value of 1, so the upper bound for first year SATT must be 4/35. Using

[9] Sometimes these contrasts are written as $WAR_i^t(1) - WAR_i^t(0)$ for $t = 1, 9$, but for the oil producing countries $WAR_i^t = WAR_i^t(1)$ for $t = 1, 9$.

similar reasoning, we find that the average "within nine years" effect must be between -27/35 and 8/35. Note that these results are consistent with previous analysis (Ross (2004) summarizes) in that they allow for the *possibility* of positive oil effects.

It is important to note three things about this bounding methodology. First, this analysis requires no assumptions about omitted variables, unmeasured confounders, or ignorability (see Manski (2003) or Quinn (2008) for a full discussion). Second, using the counterfactual model, we have ruled out some possibilities for the SATT. Without observing any data, SATT would have been bounded between -1 and 1 (for a width of 2), while the observational bounds on SATT have a width of 1. Third, if SATT is the parameter of interest, then there is no need for additional assessments of uncertainty (e.g. confidence intervals or p-values). If one wants to posit a population of countries (for which this sample is representative), then traditional techniques can be utilized because the bounds are defined by observable proportions in the data set. This is discussed in Appendix A.

However, despite the invariance of bounding analysis to confounders, when past values of the outcome variable could affect the current value of the treatment variable, we may want to stratify on these past values in order to separately consider the initializing effects of the treatment versus the reinforcing/feedback effects of the treatment. Of course, an attempt to stratify on past values poses a problem, because we must choose the relevant definition of "past value" for stratification. In the civil war example, we could stratify on war status in the *first* year prior to first-oil, or we could stratify on whether a country had *any* prior history of civil war, or we could attempt to stratify on *all* possible combinations of civil war history (see Beck et al. (1998) for a discussion on the difficulties of this task).

In this paper, I follow the convention of Fearon and Laitin (2003) by stratifying only on war status in the previous year,[10] and we find that only Colombia had pre-existing war in the year prior

[10] Even if we stratify the analysis based on *any* prior history of civil war in a nine-year pre-oil window, we find no additional countries, and if we allow for *any* prior history of civil war in the dataset, we find that only four of the countries experienced civil war prior to first significant oil. None of these experienced a civil war onset during the post-oil nine-year window. This type of analysis becomes more complicated if we attempt to count each country as more than one observation (e.g. an analysis based on country-years).

to first-oil, and this country did not experience a civil war onset after first oil (Colombia will be removed from the analysis for the remainder of the paper).[11] By considering only countries that were not experiencing war in the year prior to first oil, we restrict our attention to the initializing effects of oil production, instead of the potentially reinforcing effects of oil production.

3 Oil and State Weakness

The previous section established at least the possibility of a large, positive, initializing effect of oil on civil war. However, even if we could establish the existence of such an effect, we would like to ascertain the portion of this effect that is due to the state weakness mechanism. The state weakness mechanism corresponds to the indirect path from OIL through SW to WAR in Figure 1 (b), so to begin studying this effect, we might focus on the first stage of the path (the effect of oil production on state weakness).

To analyze the effect of oil on state weakness, we ask questions analogous to those from the previous section (with state weakness as the outcome variable): "Do countries experience state weakness after they start producing significant quantities of oil?" If so, "How long does it take?" Of course, the answers to these questions depend on the definition of "state weakness" and as noted in the introduction, this concept has been particularly difficult to measure. For the purposes of this paper, I will employ the most inclusive measure of state weakness (utilized in Humphreys (2005)) that can be calculated for each year in the data set: whether a country was either unstable or anocratic as coded in Fearon and Laitin (2003).[12] If we index the 34 remaining countries (with Colombia removed due to preexisting war) that were significant oil producers with $i = 1, ..., 34$, we can denote the answers to the questions above with the binary variables SW_i^1 (which takes the value

[11] This analysis also assumes that countries that were significant oil producers in their first year in the data set (new states and 1945 oil producers) had *no history* prior to their first year. Stratification of the analysis into first year oil producers and later oil producers does not affect the substantive findings of the paper.

[12] The Fearon and Laitin (2003) analysis includes the instability and anocracy variables entered separately in their regression, but the Humphreys (2005) combination of these measures will increase the prevelance of state weakness in the data set. Additionally, I assume that oil production is causally prior to state weakness when observed in the same year. As before, these coding choices can be seen as conservative, given the substantive goals of the paper, because they will tend to maximize the upper bounds of effects.

one if the country experienced state weakness in its first year after first significant oil production) and SW_i^9 (which takes the value one if the country experienced state weakness sometime in its the first nine years after the first significant oil production). Note that these SW_i^1 and SW_i^9 variables are well defined and observable in the data set.

Again, in order to begin addressing the effect of oil on state weakness, we must again consider the more difficult question of whether the remaining 34 countries would have experienced state weakness if they had not been significant oil producers. We will denote the answer to this question with the counterfactual binary variables $SW_i^1(0)$ and $SW_i^9(0)$. As before, we have assumed that there are right answers to these counterfactual questions, and even when these questions have well defined answers, these answers are unknown. With the relevant counterfactuals specified, we again specify country specific effects to be the contrasts between the observed variables and the counterfactual variables: $SW_i^1 - SW_i^1(0)$ and $SW_i^9 - SW_i^9(0)$.

Table 2 shows the possibilities for the 34 oil producing countries, where the counterfactual variables are unobserved. For the oil producing countries that did not experience state weakness, the country specific effect can only take the values $\{-1, 0\}$, hence we can say that for these countries, oil "didn't cause" weakness. For the oil producing countries that experienced state weakness, the country specific effect can only take the values $\{0, 1\}$, hence we can say that for these countries, oil "didn't prevent" weakness. Using the procedure described in Section 2, the first year SATT among these countries must be between -19/34 and 15/34 and the "within nine years" SATT must be between -16/34 and 18/34. These results are consistent with other results in the literature (Fearon, 2005), showing that oil *may* have highly positive short-term and medium-term effects on state weakness.

Again, if we want to (at least partially) separate the "initializing" versus "reinforcing" effects of oil on state weakness, we should stratify the analysis on whether each country was coded as weak in the year prior to first significant oil production. Table 3 shows the possibilities for the 28 countries that did not have war or state weakness in the year prior to first significant oil production.[13] The first year SATT among these countries must be between -19/28 and 9/28 and the "within nine years" SATT must be between -16/28 and 12/28, so that even when we consider only "initializing/non-feedback" effects, we retain the *possibility* of large oil on state weakness effects. In the next section we will consider only these 28 countries.

4 The State Weakness Mechanism

The previous sections established at least the possibility of large, positive, initializing oil effects on both state weakness and civil war onset. However, in order to establish even the possibility of a state weakness mechanism we must consider the following more complicated counterfactual question, "Would the 28 remaining oil producing countries have experienced a civil war onset if their state capacity after first oil production had been as it would have been had they not been significant oil producers?" If this counterfactual question is well defined, then the state weakness mechanism can be defined by the contrast between the observed war outcome and the answer to this counterfactual question. Furthermore, if state weakness is causally prior to civil war onset within each year of the data set[14] and if the binary measurement of state weakness is sufficient for the definition of the counterfactual statements, then the country level mechanism-specific effects, can be written as the product of the country level effect of oil on state weakness and the country level effect of state weakness on civil war (See Appendix B for details). It is important to note that this product methodology only works because the analysis is at the country level, the OIL and SW variables are binary, and the analysis only considers the oil producing countries. Pearl (2001) and Robins (2003) demonstrate how to extend the counterfactual definition of mechanisms to non-binary settings. Bounding in these settings is more complicated, but analogous.

Table 4 reports the possible country-level state-weakness mechanism-specific effects as the pos-

[13]Again, this analysis also assumes that countries that were significant oil producers in their first year in the data set (new states and 1945 oil producers) had *no history* prior to their first year.

[14]This assumption is made to maximize the upper bound on the average mechanism-specific effect.

sible products between the country level effects of oil on state weakness and the country level effects of state weakness on civil war. For example, the 16 oil producing countries in the first row of Table 4 did not experience state weakness in the first year after oil production, therefore oil did not cause weakness (the effect of OIL on SW^1 must be in the set of $\{-1, 0\}$). Furthermore, these countries also did not experience civil war after not experiencing state weakness, and therefore state weakness would not have prevented civil war (the effect of SW^1 on WAR^1 must be in the set of $\{0, 1\}$). Finally, because the country-level state-weakness mechanism-specific effect must be the product of these two country level effects, we know that for these countries, one of the following four things must be true:

1. Oil prevents state weakness, and state weakness has no effect on war, so the mechanism specific effect is zero ($-1 \cdot 0 = 0$).
2. Oil prevents state weakness, and state weakness causes war, so the mechanism specific effect is negative one ($-1 \cdot 1 = -1$).
3. Oil has no effect on state weakness and state weakness has no effect on war, so the mechanism specific effect is zero ($0 \cdot 0 = 0$).
4. Oil has no effect on state weakness and state weakness causes war, so the mechanism specific effect is zero ($0 \cdot -1 = 0$).

Therefore, for the 16 countries in this category, the mechanism specific effect must be -1 or 0. The other rows of the table can be similarly explained.

If we average over the possible products in Table 4, then we derive the bounds for the average mechanism-specific effect among these oil producing countries, which we will call a Sample Average Mechanism specific effect on the Treated (SAMT). Note that we cannot rule out the possibility of positive, initializing SAMT effects (the short-term upper bound is $> 10\%$ and the medium-term upper bound is $> 17\%$). However, notice as well that these upper bounds depend a great deal on the possibility of negative oil on state weakness effects and negative state weakness on civil war effects.

While it seems at least plausible that oil production *might* prevent state weakness, it is less plausible that state weakness would prevent civil war onset. In fact, if we make the seemingly reasonable assumption that state weakness cannot prevent civil war onset among the oil producing countries,[15] we will greatly reduce the upper bounds for the average state weakness mechanism. Table 5 shows the restrictions based on this weak monotonicity assumption, and the effect that this assumption has on the bounds. In particular, under this assumption the first year mechanism-specific effect *cannot be positive*. This is because among the countries that did not exhibit war or state weakness at the time of first significant oil production, no country developed state weakness and civil war in the same year as first oil production. Therefore, there is no "initializing" short-term weak states mechanism. Even when we allow for the possibility of delayed effects by using the nine-year window, we find that only two countries fit the profile of significant oil production followed by state weakness and then civil war onset. Therefore, the medium-term weak-states mechanism represented by the nine year SAMT can be *at most* 2/28. It is important to note that this upper bound, while logically possible, is not necessarily a reasonable estimate, and therefore this is very weak evidence for even the possibility of a positive medium-term effect. The next section discusses the strength of this evidence and the possibility for long-term and "reinforcing" effects.

5 Azerbaijan, Indonesia, and the Possibility for Long Term and Reinforcing Effects

The analysis in the previous section showed that there is no evidence for a short-term "initializing" state weakness mechanism, and that the evidence for the medium-term non-feedback mechanism is extremely limited. The upper bound of 2/28 demonstrates that the evidence for even the possibility of a small, positive medium-term effect rests on only two cases: Azerbaijan and Indonesia. Specifically, proponents of the medium-term initializing weak states mechanism must argue that, in these countries, the effect of oil on state weakness was positive and the effect of state weakness on civil war was also positive.

This argument will be hard to maintain in the case of Azerbaijan. Even if one were willing to

[15] Strictly speaking, we need only assume that there is no country for which oil production prevents state weakness and state weakness prevents civil war.

assume that the long running oil production produced state weakness in this case, it seems unlikely that extra state capacity (generated by a lack of oil wealth) would have prevented the Nagorno-Karabakh conflict in the early 1990s. Three quarters of the residents of the Karabakh region claimed Armenian ethnicity, and these former Soviets had been agitating for the inclusion of this region into the Armenian SSR since well before independence from Russia (De Waal, 2003, pp.11-12). Hence the motivation for war was already present. Furthermore, the mountainous terrain and the removal of Soviet forces (that had been enforcing Azerbaijani authority in the region) provided the opportunity for conflict (De Waal, 2003, pp.160-166). Hence unless we believe that not producing oil would have prevented the removal of the Soviet forces, or that not producing oil would have increased the state capacity of Azerbaijan to the point that the new Azerbaijani state would have been able to suppress the conflict, we must conclude that the state weakness mechanism was not operative for this case.

In the case of Indonesia (Vickers, 2005, pp.99-112), the state weakness observed during the early years of independence was likely the result of the long and bloody war of independence with the Dutch (an entirely different sort of "Dutch Disease"). One could argue that this war of independence was the result of oil (i.e. the Dutch wouldn't have tried to regain control of their colony after World War II if Indonesia had not been a significant oil producer). However, even if this were true, this story does not fit the popular version of the state weakness mechanism. If oil caused state weakness in this case, it did not seem to do so by reducing the incentives for bureaucratic development.

If the historical record from the only two possible cases of a medium-term initializing mechanism-specific effect casts doubt on the possibility of a state weakness mechanism, the evidence is stronger for the possibility of long-term effects and reinforcing effects. Five oil producing countries exhibited state weakness followed by conflict onset sometime after the first decade of significant oil production. However, three of these cases also showed some history of civil conflict prior to the exhibited state weakness, so it is unclear how much of this should be interpreted as evidence for an initializing effect. There is considerably more evidence for reinforcing effects (even medium-term reinforcing effects). Of the six countries that experienced state weakness in the year prior to first significant oil production, three experienced state weakness followed by civil war sometime in the decade after first-oil. Therefore, if significant oil production has an effect on the probability of civil war onset (in its relation to state weakness), it most plausibly functions to reinforce state weakness, and/or to increase the probability of civil conflict in states that are already weak (this is consistent with the findings in Humphreys (2005) and Snyder (2006)).

6 Conclusion

The analysis in this paper has shown that the state weakness mechanism, if it exists, most plausibly functions *only* to reinforce pre-existing war or weakness (or must take longer than a decade to form). Using standard measures for oil production, state weakness, and civil war, the analysis in Sections 2 and 3 shows that previously reported results for the effects of oil production on state weakness and oil production on civil war are plausible. However, using these same measures and the bounding methodology developed in this paper, Sections 4 and 5 show that the state weakness mechanism-specific effect does not cause civil war in the first decade after significant oil production. Furthermore, the bounding methodology ensures that this result is robust to unmeasured confounding. However, while the analysis used in this paper does not rely on assumptions about alternative risk factors, three important assumptions were made.

First, causal independence was assumed between countries. Specifically, we have assumed that oil production in one country does not affect state weakness in other countries, and that the level of oil production and state weakness in one country does not affect the civil war outcomes for other countries. This assumption is clearly questionable, however, it is important to note that nearly all other empirical studies of oil, civil war and the weak states mechanism rely on these assumptions as well.

Second, as in previous studies, we have assumed that binary measurement of oil production (Fearon and Laitin, 2003) and state weakness (Humphreys, 2005) was sufficient. This may be problematic, and one could conclude that the results from Section 4 merely demonstrate the inadequate measurement of the variables in this application. However, the results from Sections 2 and 3 were consistent with previous results, so if the results from Section 4 are due to measurement inadequacy, this inadequacy is special to the question asked in this section.

Third, we have assumed a causal order to the variables (within a given year, oil production is assumed to precede state weakness which is assumed to precede civil war). While this assumption may be problematic, it was chosen to maximize the upper bounds for the mechanism-specific causal effect, and can therefore be seen as conservative.

Even taking these assumptions into account, the results of this paper demonstrate the necessity for a more refined notion of the state weakness mechanism. If better measurement of the key variables is needed in order to establish an "initializing" effect, then this will surely have implications for the construction of policy. If instead, oil production causes civil war onset primarily through the reinforcement of preexisting state weakness, then the policies recommended by conflict scholars should be tailored specifically for weak states. Future studies on the state weakness mechanism should focus on distinguishing between these scenarios and producing the corresponding refinements.

Appendix A: Inference for the Bounds

All of the SATT and SAMT lower and upper bounds in this paper correspond to observable proportions in the data set. Therefore, if we are willing to posit a population from which our sample is representative (and a corresponding Population Average Treatment effect on the Treated (PATT) or a Population Average Mechanism specific effect on the Treated (PAMT)), then we can make inference to the bounds for PATT or PAMT that would be represented by the population proportions. Confidence intervals and hypothesis tests will follow the usual rules of inference for proportions.

However, for the application considered in this paper, the typical asymptotic approximations will not be appropriate due to the limited sample sizes and the consideration of population proportions that are near the boundary of the parameter space. If an infinite population is considered, then the binomial distribution should be utilized for the creation of intervals. For example, while the upper bound for the mechanism-specific nine-year SAMT considered in Section 4 is 2/28, we could consider forming a 95% one sided confidence interval to assess our uncertainty about the upper bound for the PAMT. This analysis finds the lower confidence bound for the upper PAMT bound to be 0.013 (i.e. we cannot reject the hypothesis that the upper bound on the PAMT is 0.013). If the population is considered finite, then the hypergeometric distribution should be used instead of the binomial distribution. Considering the same result from Section 4, we find that the 95% lower confidence bound for the upper PAMT bound cannot be larger than 0.016, and this corresponds to a population size of 123. Therefore, if we assume that the relevant population of countries is no larger than 123, we cannot reject the hypothesis that the two in-sample countries that exhibit the *possibility* of a positive effect are the only such countries in the population.

In general, it will be more philosophically defensible (and more useful for the prediction of policy effects) to use a Bayesian approach to make uncertainty statements about the causal effect bounds. Quinn (2008) provides such an approach for the effects described in Sections 2 and 3. In principle, it is possible to extend this approach to the effects described in Section 4.

Appendix B: Decomposition of Mechanisms Specific Effects with Binary Variables

We define SW_i^{9*} to take the value one only when state weakness occurs *before* a civil war onset, and we assume that the following counterfactual binary variables are well defined for the oil producing countries,

$$WAR_i^1(SW_i^1 = SW_i^1(OIL_i = 0)) = \\ WAR_i^1(SW_i^1 = 1) \cdot SW_i^1(1) + WAR_i^1(SW_i^1 = 0) \cdot (1 - SW_i^1(1)) \quad (1)$$

$$WAR_i^9(SW_i^{9*} = SW_i^{9*}(OIL_i = 0)) = \\ WAR_i^9(SW_i^{9*} = 1) \cdot SW_i^{9*}(1) + WAR_i^9(SW_i^{9*} = 0) \cdot (1 - SW_i^{9*}(1)) \quad (2)$$

so that the war outcomes, if state weakness had been as if not an oil producer, are equal to either the war outcome if weak, or the war outcome if not weak, depending on the state weakness outcome if not an oil producer.

In order for the first-year counterfactual to be well defined, we must assume that within each year of the data set, oil production is causally prior to state weakness which is causally prior to the war outcome. Furthermore, we must assume that the binary measurement of oil production is sufficient for the year-one state weakness outcome to be well defined had the countries not been oil producers, and we must assume that the binary measurement of state weakness is sufficient for the counterfactual war outcome to be well defined if the weak states had not been weak and if the non-weak states had been weak. In order for the nine-year counterfactual to be well defined, we must make similar assumptions, however, note that the state weakness variable now covers multiple years, so that a country will be counted as weak if it exhibits *any* weakness after oil production but prior to civil war onset.

For the oil producing countries, $WAR_i^1 = WAR_i^1(SW_i^1 = SW_i^1(1))$, and given the decomposition defined in (1), the year-one mechanism-specific effect can be written as the following:

$$
\begin{aligned}
WAR_i^1 - WAR_i^1(SW_i^1 = SW_i^1(0)) &= WAR_i^1(SW_i^1 = SW_i^1(1)) - WAR_i^1(SW_i^1 = SW_i^1(0)) \\
&= WAR_i^1(SW_i^1 = 1) \cdot SW_i^1(1) + WAR_i^1(SW_i^1 = 0) \cdot (1 - SW_i^1(1)) \\
&\quad - WAR_i^1(SW_i^1 = 1) \cdot SW_i^1(0) - WAR_i^1(SW_i^1 = 0) \cdot (1 - SW_i^1(0)) \\
&= \left\{SW_i^1(1) - SW_i^1(0)\right\} \\
&\quad \cdot \left\{WAR_i^1(SW_i^1 = 1) - WAR_i^1(SW_i^1 = 0)\right\}.
\end{aligned}
$$

so that the year-one state weakness mechanism can be represented as the product of the oil on state weakness effect and the state weakness on civil war effect.

Additionally, $WAR_i^9 = WAR_i^9(SW_i^{9*} = SW_i^{9*}(1))$ for the oil producing countries, and the decomposition defined in (2) allows the nine-year mechanism-specific effect to be written as the following:

$$
\begin{aligned}
WAR_i^9 - WAR_i^9(SW_i^{9*} = SW_i^{9*}(0)) &= WAR_i^1(SW_i^{9*} = SW_i^{9*}(1)) - WAR_i^1(SW_i^{9*} = SW_i^{9*}(0)) \\
&= WAR_i^1(SW_i^{9*} = 1) \cdot SW_i^{9*}(1) + WAR_i^1(SW_i^{9*} = 0) \cdot (1 - SW_i^{9*}(1)) \\
&\quad - WAR_i^1(SW_i^{9*} = 1) \cdot SW_i^{9*}(0) - WAR_i^1(SW_i^{9*} = 0) \cdot (1 - SW_i^{9*}(0)) \\
&= \left\{SW_i^{9*}(1) - SW_i^{9*}(0)\right\} \\
&\quad \cdot \left\{WAR_i^1(SW_i^{9*} = 1) - WAR_i^1(SW_i^{9*} = 0)\right\}.
\end{aligned}
$$

so that the nine-year state weakness mechanism can also be represented as the product of an oil on state weakness effect and a state weakness on civil war effect. See Glynn (2009) for details and for decompositions in the more general case of individual-level linear heterogeneous-effects models with interactions. These decompositions can also be used to extend the bounding approach if bounds are placed on the effects in that model.

References

Balke, A., and J. Pearl. 1997. "Bounds on Treatment Effects from Studies with Imperfect Compliance." *Journal of the American Statistical Association* 92(439).

Beck, N., J.N. Katz, and R. Tucker. 1998. "Taking Time Seriously: Time-Series-Cross-Section Analysis with a Binary Dependent Variable." *American Journal of Political Science* 42:1260–1288.

Cai, Z., M. Kuroki, J. Pearl, and J. Tian. 2007. "Bounds on Direct Effects in the Presence of Confounded Intermediate Variables." *Biometrics* .

Chaudhry, K.A. 1997. *The Price of Wealth: Economies and Institutions in the Middle East*. Cornell University Press.

Chickering, D.M., and J. Pearl. 1997. "A clinicians tool for analyzing non-compliance." *Computing Science and Statistics* 29(2):424–431.

Collier, P., and A. Hoeffler. 2004. "Greed and grievance in civil war." *Oxford Economic Papers* 56(4):563–595.

Dawid, A.P. 2000. "Causal inference without counterfactuals (with discussion)." *Journal of the American Statistical Association* 95(450):407–424.

De Waal, T. 2003. *Black garden: Armenia and Azerbaijan through peace and war*. New York University Press.

Fearon, James D., and David D. Laitin. 2003. "Ethnicity, Insurgency, and Civil War." *American Political Science Review* 97:75–90.

Fearon, J.D. 1991. "Counterfactuals and Hypothesis Testing in Political Science." *World Politics* 43(2):169–195.

Fearon, J.D. 2005. "Primary Commodity Exports and Civil War." *Journal of Conflict Resolution* 49(4):483.

Glynn, A.N. 2009. "The Product and Difference Fallacies for Indirect Effects." *Submitted for Publication* .

Humphreys, M. 2005. "Natural Resources, Conflict, and Conflict Resolution: Uncovering the Mechanisms." *Journal of Conflict Resolution* 49(4):508.

Karl, T.L. 1997. *The Paradox of Plenty: Oil Booms and Petro-States*. University of California Press.

Lebow, R.N. 2000. "What's So Different about a Counterfactual?" *World Politics* 52(4):550–585.

Lewis, D.K. 1973. *Counterfactuals*. Oxford, U.K.: Blackwell.

Manski, C. 1990. "Nonparametric Bounds on Treatment Effects." *American Economic Review* 80(2):319–323.

Manski, C.F. 2003. *Partial Identification of Probability Distributions*. Springer.

Moore, M. 2001. "Political Underdevelopment: What causes' bad governance'." *Public Management Review* 3(3):385–418.

Morgan, S.L., and C. Winship. 2007. *Counterfactuals and Causal Inference*. Cambridge University Press.

Pearl, J. 2001. "Direct and indirect effects." *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence* pp. 411–420.

Przeworski, Adam. 2004. "Institutions Matter?" *Government and Opposition* 39(4):527–540.

Quine, W.V. 1953. *From the logical point of view*. Harvard University Press Cambridge, Mass.

Quinn, K.M. 2008. "What Can Be Learned from a Simple Table? Bayesian Inference and Sensitivity Analysis for Causal Effects from 2 x 2 and 2 x 2 x K Tables in the Presence of Unmeasured Confounding." *Working Paper* .

Robins, J.M. 2003. "Semantics of causal DAG models and the identification of direct and indirect effects." *Highly Structured Stochastic Systems* pp. 70–81.

Robins, J.M., and S. Greenland. 1992. "Identifiability and Exchangeability for Direct and Indirect Effects." *Epidemiology* 3(2):143–155.

Robins, J.M., and S. Greenland. 2000. "Comment on" Causal Inference Without Counterfactuals" by AP Dawid." *Journal of the American Statistical Association* 95:431–435.

Ross, M.L. 2004. "What Do We Know about Natural Resources and Civil War?" *Journal of Peace Research* 41(3):337.

Sekhon, J.S. 2004. "Quality Meets Quantity: Case Studies, Conditional Probability, and Counterfactuals." *Perspectives on Politics* 2(02):281–293.

Snyder, R. 2006. "Does Lootable Wealth Breed Disorder?: A Political Economy of Extraction Framework." *Comparative Political Studies* 39(8):943.

Sørli, M.E. 2002. "Resources, Regimes and Rebellion: A Critical Assessment of the Greed & Grievance Model of Civil Conflict."

Vickers, A. 2005. *A history of modern Indonesia*. Cambridge University Press.

Wantchekon, L. 2002. "Why Do Resource Abundant Countries Have Authoritarian Governments?" *Journal of African Finance and Economic Development* 2:57–77.

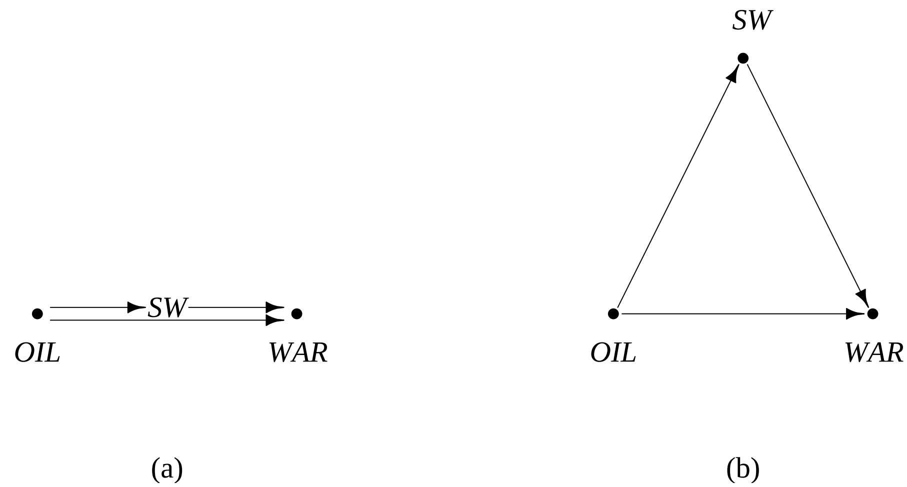

Figure 1: *Graphical model depicting the state weakness mechanism.* Panel (a) represents the reduced form model where the direct and indirect effects have been combined into a total effect. Panel (b) represents the full model where the potential effect of oil production (OIL) on civil war onset (WAR) can be decomposed into an indirect effect though state weakness (SW) and a direct effect (representing the aggregation of all other possible pathways).

Table 1: Effect of first oil production on civil war onset for the oil producing countries. The set notation $\{\cdot\}$ indicates the possible values for the counterfactual variables and effects in this table.

	# of Obs	OIL	WAR^1	$WAR^1(0)$	$WAR^1 - WAR^1(0)$ Effect	
Short	31	1	0	$\{0, 1\}$	$\{-1, 0\}$	"didn't cause"
Term	4	1	1	$\{0, 1\}$	$\{0, 1\}$	"didn't prevent"
	35				Avg Effect Bounds = $[-31/35, 4/35]$	
	# of Obs	OIL	WAR^9	$WAR^9(0)$	$WAR^9 - WAR^9(0)$ Effect	
Medium	27	1	0	$\{0, 1\}$	$\{-1, 0\}$	"didn't cause"
Term	8	1	1	$\{0, 1\}$	$\{0, 1\}$	"didn't prevent"
	35				Avg Effect Bounds = $[-27/35, 8/35]$	

Table 2: Effects of first oil production on state weakness for the oil producing countries that were not in a state of war prior to first oil production.

	# of Obs	OIL	SW^1	$SW^1(0)$	$SW^1 - SW^1(0)$ Effect	
Short	19	1	0	$\{0,1\}$	$\{-1,0\}$	"doesn't cause"
Term	15	1	1	$\{0,1\}$	$\{0,1\}$	"doesn't prevent"
	34				Avg Effect Bounds = $[-19/34, 15/34]$	
	# of Obs	OIL	SW^9	$SW^9(0)$	$SW^9 - SW^9(0)$ Effect	
Medium	16	1	0	$\{0,1\}$	$\{-1,0\}$	"doesn't cause"
Term	18	1	1	$\{0,1\}$	$\{0,1\}$	"doesn't prevent"
	34				Avg Effect Bounds = $[-16/34, 18/34]$	

Table 3: Effects of first oil production on state weakness for the oil producing countries that were not in a state of war or weakness prior to first oil production.

	# of Obs	OIL	SW^1	$SW^1(0)$	$SW^1 - SW^1(0)$ Effect
Short	19	1	0	$\{0, 1\}$	$\{-1, 0\}$ "doesn't cause"
Term	9	1	1	$\{0, 1\}$	$\{0, 1\}$ "doesn't prevent"
	28				Avg Effect Bounds = $[-19/28, 9/28]$
	# of Obs	OIL	SW^9	$SW^9(0)$	$SW^9 - SW^9(0)$ Effect
Medium	16	1	0	$\{0, 1\}$	$\{-1, 0\}$ "doesn't cause"
Term	12	1	1	$\{0, 1\}$	$\{0, 1\}$ "doesn't prevent"
	28				Avg Effect Bounds = $[-16/28, 12/28]$

Table 4: Effect bounds for the Sample Average Mechanism specific effect on the Treated (SAMT)– the average effect of first oil production on civil war onset that goes through state weakness for the oil producing countries that did not experience a civil war or state weakness in the year prior to first oil production.

# of Obs	OIL	SW^1	WAR^1	OIL on SW^1	SW^1 on WAR^1	Possible Products
16	1	0	0	$\{-1,0\}$	$\{0,1\}$	$\{-1,0\}$
3	1	0	1	$\{-1,0\}$	$\{-1,0\}$	$\{0,1\}$
9	1	1	0	$\{0,1\}$	$\{-1,0\}$	$\{-1,0\}$
0	1	1	1	$\{0,1\}$	$\{0,1\}$	$\{0,1\}$
28						AVG= $[-25/28, 3/28]$
# of Obs	OIL	SW^9	WAR^9	OIL on SW^9	SW^9 on WAR^9	Possible Products
13	1	0	0	$\{-1,0\}$	$\{0,1\}$	$\{-1,0\}$
3	1	0	1	$\{-1,0\}$	$\{-1,0\}$	$\{0,1\}$
10	1	1	0	$\{0,1\}$	$\{-1,0\}$	$\{-1,0\}$
2	1	1	1	$\{0,1\}$	$\{0,1\}$	$\{0,1\}$
28						AVG= $[-23/28, 5/28]$

Table 5: Effect bounds for the Sample Average Mechanism specific effect on the Treated (SAMT)– the average effect of first oil production on civil war onset that goes through state weakness for the oil producing countries that did not experience a civil war or state weakness in the year prior to first oil production. Second stage monotonicity was assumed (i.e. state weakness cannot prevent war).

# of Obs	OIL	SW^1	WAR^1	OIL on SW^1	SW^1 on WAR^1	Possible Products
16	1	0	0	$\{-1, 0\}$	$\{0, 1\}$	$\{-1, 0\}$
3	1	0	1	$\{-1, 0\}$	$\{\cancel{-1}, 0\}$	$\{0, \cancel{1}\}$
9	1	1	0	$\{0, 1\}$	$\{\cancel{-1}, 0\}$	$\{\cancel{-1}, 0\}$
0	1	1	1	$\{0, 1\}$	$\{0, 1\}$	$\{0, 1\}$
28						$AVG = [-16/28, 0/28]$

# of Obs	OIL	SW^9	WAR^9	OIL on SW^9	SW^9 on WAR^9	Possible Products
13	1	0	0	$\{-1, 0\}$	$\{0, 1\}$	$\{-1, 0\}$
3	1	0	1	$\{-1, 0\}$	$\{\cancel{-1}, 0\}$	$\{0, \cancel{1}\}$
10	1	1	0	$\{0, 1\}$	$\{\cancel{-1}, 0\}$	$\{\cancel{-1}, 0\}$
2	1	1	1	$\{0, 1\}$	$\{0, 1\}$	$\{0, 1\}$
28						$AVG = [-13/28, 2/28]$