Meta-analysis of Commodity Prices and Conflict[†]

Pre-Analysis Plan

Graeme Blair[‡] Darin Christensen[§] Aaron Rudkin[¶]

1 Introduction

Political scientists and economists have assembled a body of research assessing the causal relationship between commodity prices and conflict (Dube and Vargas 2013; Bazzi and Blattman 2014; Christensen 2018). Results have been mixed. On the one hand Bazzi and Blattman (2014) state definitively "Price shocks have no effect on new conflict, even large shocks in high-risk nations." Yet, Dube and Vargas (2013), to take a prominent example, find that conflict increases in Colombia's oil-producing municipalities as the international price for oil goes up. These findings could be due to theoretical choices about what types of conflicts or commodities to study, or design decisions about what cases to sample or the most credible empirical strategy.

Both the outpouring of recent research and its conflicting findings motivate this meta-analysis. We believe there is now a critical mass of studies that address the question, how do commodity price changes affect civil conflict? Moreover, these studies share a concern with causal identification and employ methods that exploit exogenous, inter-temporal variation in food, fuel, or mineral prices. Despite these commonalities, these studies also differ in ways that might explain their divergent findings. We seek to define the theoretically relevant dimensions and specify when and how we'll test for heterogeneity across studies.

This document serves as a pre-analysis plan for a meta-analysis to adjudicate the state of research on commodity prices and conflict. We intend for this meta-analysis to exhibit best methodological practices by adhering to the principles and requirements of the PRISMA statement (Moher et al. 2009). The PRISMA checklist is included as an appendix.[1]

This pre-analysis plan is designed to pre-register our hypotheses, data collection process, and analyses. It is registered at the EGAP design registry.

2 Theory and Hypotheses

We study two conflict outcomes: armed intrastate conflict and coups d'état. We distinguish between intrastate conflict over an incompatibility about who controls a territory and conflict over who controls the government of the state (Themnér 2014). We consider two forms of heterogeneity in the effects of commodity prices on

[†]The latest version of this preanalysis plan and associated replication materials is archived at https://dx.doi.org/10.17605/OSF.IO/DY9UF. The PAP is registered with EGAP. We thank Maryam Aljafen, Macartan Humphreys, Michael Ross, Valerie Wirtschafter, and the UCLA Improving Designs in the Social Sciences workshop for helpful comments.

[‡]Assistant Professor of Political Science, UCLA (graeme.blair@ucla.edu). https://graemeblair.com

[§]Assistant Professor of Public Policy and Political Science, UCLA (darinc@luskin.ucla.edu). http://darinchristensen.com

[¶]Ph.D. candidate in political science, UCLA (rudkin@ucla.edu). http://rudkin.ca

[1]The checklist includes items that are reported on after data collection; we only report on pre-data collection items in this document.

conflict and coups: whether extraction of the commodity is labor-intensive or capital-intensive; and whether the state can easily profit from sale of the commodity, whether through taxation or direct ownership or management of the commodity (we label these commodities "taxable" as a shorthand).

Why do changes to commodity prices affect violence? In what follows, we review prominent theoretical accounts and summarize their empirical implications in a table of hypotheses (Table 1).

	Outcome				Incompatibility[b]		Commodity Type			
	Armed Conflict									
Effect	Onset	Intensity	Duration	Coups	Territory	Government	Labor Intensive	Capital Intensive	Taxable[a]	Non-Taxable
1. Rapacity	+	+	?	0	+	0	+	+	+	+
2. State prize	+	+	?	+	0	+	+	+	+	0
3. Opportunity cost	0	+/-	0	0	+/-	+/-	-	+	+/-	+/-
4. State capacity	-	-	-	-	-	-	-	-	-	0
5. Dal Bó—Dal Bó	+	+/-	?	0	+/-	0	-	+	+/-	+/-

Legend: + positive effect; - negative effect; +/- effect direction depends on mix of commodity types (labor- and capital-intensive; taxable and non-taxable); ? no clear implication of theory.

[a] Taxable is shorthand for commodities the state can profit from, whether through taxation or direct ownership or management of the commodity.

Table 1: Hypotheses. The pattern of expected effects of an increase in commodity prices is presented for each theory described in Section 2, according to outcome, type of commodity, and the incompatibility of the conflict (either to control the territory surrounding the resource, *territory*, or to change the *government* of the state).

First, the **rapacity** effect: as the price of a commodity increases, so too do the returns to seizing control of its production. Price increases should increase violence at the site of production (or point of extraction) or along the route to market. The rapacity effect implies an increase in the onset and intensity of conflict in the territory where the commodity is produced or transported (Keen 1998; Collier and Hoeffler 2004; Dal Bó and Dal Bó 2011; Dube and Vargas 2013). This effect should hold for both labor- and capital-intensive commodities. The rapacity effect does not affect coups, and the effect on the duration of conflict is ambiguous.[2]

A variant is the **state prize** effect (Soysa 2002; Fearon and Laitin 2003; Fearon 2005). In this account, the state is the primary target; control of government implies control over the revenues derived from royalties, taxes on the resource sector, or state ownership. If commodity price increases translate into a larger state prize (e.g., more royalties), then this motivates conflict. By this logic, only taxable commodities are likely to affect conflict. The implications of the state prize effect are similar to those of the rapacity effect: onset and intensity are expected to increase, with ambiguous effects on duration. Yet, conflict should focus on control of the state, not control of the sites of extraction, production, and transportation (Ross 2012; Paine 2016). The state prize hypothesis also implies an increase in coup risk.

Second, the **opportunity cost** effect. Individuals choose to work in either productive sectors or in the appropriation sector (e.g., as rebels or bandits), aiming to control commodity profits. When employment and wages are high, individuals face a greater opportunity cost to leave productive sectors; conversely, when economic fortunes are poor, individuals have little to lose (Becker 1968; Grossman 1991). If commodity prices affect employment and wages, they thereby affect rates of participation in the appropriation sector. Under the opportunity cost effect, increases in prices of capital-intensive (labor-intensive) commodities increase (decrease) the intensity of conflict (Dal Bó and Dal Bó 2011). Price changes do not affect conflict onset or

[2]For a discussion, see Ross (2004b), which suggests that for a rapacity mechanism to result in longer conflicts, the weaker side must be able to finance itself from the resource. Several accounts discount this possibility empirically (Ross 2004a, Fearon (2005)).

duration or the risk of coups. These effects occur whether or not the commodity is taxable, and predict changes both in conflict over territory and government control.

Increases in commodity prices increase returns to predation, but may also drive up wages and thus the opportunity cost to fighting. Thus, a key theoretical question is when rapacity effects will dominate opportunity cost effects for potential fighters. Dal Bó and Dal Bó (2011) argue that the rapacity effect dominates for commodities with capital-intensive production, which raise returns to fighting faster than wages. For labor-intensive commodities, price increases will lead to wage increases that dominate rapacity effects. The Dal Bó and Dal Bó (2011) hypothesis is, thus, a combination of the rapacity and opportunity cost effects (see also Dube and Vargas 2013). The risk of conflict onset is increased, while the effect on conflict intensity is increased by rises only in the prices of capital-intensive commodities. These effects obtain both for armed conflict over territory and who governs, but do not affect the risk of coups (which typically do not involve labor outside the government and the military). Whether or not the commodity is taxable does not moderate these effects.

Third, the **state capacity** effect. When states earn revenues from commodities, price changes affect the state's capacity to buy off or suppress (or deter) opponents through spending on the military or police (Snyder and Bhavnani 2005; Cotet and Tsui 2013; Paine 2016).[3] Increases in commodity prices uniformly affect conflict under the state capacity effect, reducing the onset, intensity, and duration of conflict. However, only the prices of taxable commodities that the state derives revenues from yield the state capacity effect.

We exclude several notable mechanisms. Many believe some commodities like alluvial diamonds open opportunities for armed groups to finance themselves; evidence suggests this rarely if ever has happened (Ross 2004a; Fearon 2005). We also exclude mechanisms that operate only in the long-run. Most notably, early studies of the relationship between natural resources (a subset of the commodities we study) and conflict focused on the role of grievances. Theoretical work has cast doubt on this claim, given that in general grievances are widespread yet conflict is not (Fearon and Laitin 2003). Recent accounts of how grievances mediate the effects of natural resources focus on increased income inequality in resource-rich states, a long-run process (Cederman, Gleditsch, and Buhaug 2013). Export commodities may harm non-export sectors of the economy, and the resulting depressed levels of internal trade may lead to violence (Humphreys 2005).

In addition to studying armed conflict and coups, we will study social conflict (e.g., protests and riots). The theoretical literature is less developed in this area, so we will not present hypotheses about expected effects.

3 Search Criteria

Studies will be identified using the following strategy: First, we will search an online academic database (Google Scholar) with the query described below. Next, we track any paper marked on Google Scholar as citing focal papers in this literature: specifically, Dube and Vargas (2013), Bazzi and Blattman (2014), and Collier and Hoeffler (2004). Finally, we will ask authors who we have identified as studying this topic for additional working papers not yet listed online.

[3] The relationship between state capacity and commodities is often cast in the opposite direction: oil-rich states should have *less* state capacity (el-Beblawi and Luciani 1987; Ross 2001). This refers to a long-term process of oil revenues substituting for tax revenues and reducing fiscal capacity. This alternative, often dubbed the rentier state mechanism, is not relevant for the short-term price fluctuations we study here.

We will search paper titles and descriptions for the following query, subject to syntax changes depending on the database of interest:[4] ("resource" OR "commodity" OR "oil" OR "mineral" OR "diamond" OR "gold" or "drug" OR "mine") AND ("price" OR "value" OR "shock") AND ("conflict" OR "protest" OR "war" OR "violence" OR "attack").

We will record each study's title and metadata information at this stage, then apply the methodological filter described below. At each stage we will record information about the number of qualifying papers.

4 Inclusion Criteria, First Stage

At this stage, we will collect an exhaustive super-set of all studies we will later include in the meta-analysis. Studies are eligible for inclusion at this stage if they satisfy the following criteria:

- English language;
- Published in a peer reviewed journal or a circulating working paper of professional quality.
- Include empirical results that are not just simulations and use real-world data.
- Describe the relationship between the value of natural resources (e.g., commodity prices, including illicit and non-mineral commodities) and conflict (e.g., civil war, protest), including data on value of natural resources and conflict.
- Paper is available to us after a good faith effort to obtain access, in a standardized computer format such as DOCX, PDF, HTML, or TXT and has metadata information including a title and author(s) attached to the document.

Although it is impossible to anticipate every design that may or may not satisfy these requirements, our intent is to include only those studies which use real-world measures of price and conflict; this would, for instance, exclude studies which use individual-level surveys or survey experiments which measure theoretical willingness to participate in conflict or offer theoretical economic conditions to respondents.

5 Study Selection, Second Stage

At this stage, we cull from all eligible studies those which pass a methodological and substantive filter. The study's outcome must relate to social conflict, armed conflict, or coups. We permit different measures of conflict, both in terms of the scale of violence (e.g. civil war vs. protest), and how conflict is measured (e.g. onset, incidence, duration, or extent).

Our methodological filter builds on this substantive filter to select only those studies which use a credible strategy to identify a causal relationship between resource commodity prices and conflict outcomes. In short, studies needed to be estimating versions of the following, general model:

$$\text{Conflict}_{it} = \alpha_i + \theta_t + \beta \text{Value of Natural Resource}_{it} + \epsilon_{it}$$

[4] Journal search engines may prevent us from gathering an exhaustive sample with such a broad query; thus, we expect to disaggregate the query piece by piece. We will report the final search queries in a technical appendix.

In this model, i indexes geographic units (e.g., countries, provinces, or districts); t, time; and β estimates the effect of changes in the value of resources. This model implies the inclusion of unit and time fixed effects to absorb time-invariant characteristics of geographic units and temporal shocks that affect all units in a given time period. Papers need not report the magnitude of fixed effects; they need only include such terms.

This design would exclude papers that use purely cross-sectional data or a pooled estimator with panel data, unless authors credibly argue that $\mathbb{E}[\epsilon_{it}|\text{Price}_{it}] = 0$. We will note papers that are included on the basis of credible identification strategies that nevertheless depart from the above general model in a technical appendix to the resulting meta-analysis.

We admit papers that incorporate the Value of Natural Resource_{it} variable in two ways: direct measures of commodity prices included in the regression or as instrumented first-stage outcomes in instrumental variables designs.

In order to be included, studies must report i) effect estimates and ii) measures of uncertainty (i.e. $\widehat{\beta}$ and its standard error in the above model). If measures of uncertainty are not provided, we will admit the study if it provides sufficient information to compute standard errors. We do not insist on any specific level of clustering of ϵ_{it}.

We are interested in papers affecting any level of geographic aggregation and any temporal period for which data exists. Studies may be cross-national or rely only on within-country variation.

6 Data Collection Process

Our general model guides our data collection: we focus on tables that estimate the relationship between changes in the value of a natural resource and conflict.

Our experience from prior knowledge of papers written on this topic guides considerations about which models to select within a given paper. We have attempted to exhaustively describe which specifications we will gather:

- Papers generally include *multiple (conflict) outcomes*. By "conflict outcomes", we mean measurements of: onset, intensity, and duration of armed internal conflict; count of coups; and count of social conflict events. We will extract estimates of the treatment effect for all conflict outcomes supplied. Conflict outcomes are estimated separately in our meta-analysis.
- Papers may employ *multiple specifications*, including and excluding covariate adjustment ("controls"). When this occurs, we will select specifications that hew closest to our general model to improve direct comparability between studies. Concretely, we favor simpler models over those that include additional covariate adjustment. We will choose the simplest model, unless the authors of the paper clearly indicate a preferred specification on a design basis, in which case we will choose that model.
- Papers sometimes employ *multiple standard error calculations*. We prefer standard error calculations that are justified based on the research design or, if none are justified by the design, are heteroskedasticity-robust and clustered at the appropriate geographic level of variation. We will extract and report only one standard error and justify choices in supplementary notes.

- Papers may include *multiple levels of geographic aggregation.* In general, we prefer to extract the estimates based on the lowest level of geographic aggregation for which full data is available, unless authors raise a clear preference for a different level of geographic aggregation on the basis of theory, data integrity, or causal identification. We will extract and report only one level of geographic aggregation and justify choices in supplementary notes.
- Some papers may aggregate or disaggregate *multiple commodities.* We will admit four types of estimates: overall measures that aggregate multiple primary commodities; measures that aggregate only labor-intensive commodities (as labeled by authors); measures that aggregate only capital-intensive commodities (as labeled by authors); and measures for individual primary commodities (which we will classify as either primarily labor- or capital-intensive). We discuss rules for coding these commodities in an appendix below.
- Some papers may present multiple *tax statuses* – i.e. whether the commodity in question is taxable or non-taxable. We will admit three types of estimates: overall measures which combine taxable and non-taxable commodities, measures which contain only taxable commodities, and measures which contain only non-taxable commodities. We describe the process of labeling commodities in an appendix below.
- Some papers may present multiple *incompatibility types*: whether a conflict is territorial or governmental. We admit both, as well as measures that combine both and discuss all conflicts. We describe the process of labeling conflict types in an appendix below.
- Some papers use *multiple temporal lags.* There is not a strong a priori reason to believe that all effects of commodity prices are realized instantaneously. Where authors report a single temporal lag in their main paper and offer others as "extra specifications" or "robustness checks", we will collect the main temporal lag. In cases where authors report multiple temporal lags in the same model, we will prefer an aggregation of all studied temporal lags into a single effect estimate; if one is not available, we commit to calculating one. If it is not possible to calculate one because replication data or code are not readily available, we will choose the shortest reported temporal lag in the model.
- Some papers report *different model link functions.* If a paper reports models which otherwise are identical on the above criteria, but differ in the link function they use, we select linear models, including linear probability models, in order to, when possible, compare estimated effect sizes on a common scale

In the event that we have failed to specify other possible sources for multiple specifications, we offer the following principles for choosing between them: we will collect the effect estimate that most closely hews to our general model; that is no more complex than it needs to be; that is the author's "primary" or "preferred" model if noted in text; and that if all else fails is the model that we judge to be the best identified model with respect to the relationship of interest, regardless of effect size or direction.

In the event that unexpected study characteristics appear to create a "researcher degree of freedom" with respect to choosing which estimates to extract, we will follow the above rules as closely as possible, document deviations, and select coefficients without respect to magnitude or direction of the effect. If there are instances where it is impossible to adjudicate the best selection, we will retain each possible choice in the replication materials for readers to examine (but will not present results except for the single number that was selected, given the high number of possible combinations).

7 Data Items

We extract the estimate and standard error for the estimates of β in concert with the specification choice rules specified in the section above. In addition, we extract reported p-values for use in our p-curve diagnostic.

If no standard error is present, but replication data and code are available, we will re-estimate the model to calculate a standard error.

If this is not possible, but the design suggests a normal approximation can be used to deduce the approximate standard error from other measures of uncertainty (i.e. t-statistic or p-value), we will do so. If it is not possible to accurately estimate uncertainty in this manner, the paper will not be included in keeping with our guidelines above. We will record any intervention required to calculate a standard error in an appendix.

If no p-value is available, we will calculate one.

In addition, in cases where reported effect sizes are not standardized, we will record information sufficient to standardize the effect size (the standard deviations of the treatment variables and outcome variables). Where the model selected uses a non-linear link function, we seek replication data sufficient to re-estimate the model on the same data with a linear link function and contact authors if no data is publicly available. If sufficient information to standardize effects is not available and we cannot source replication data or code, we will not include the study.

In addition to effect size and precision, we will also extract effect level metadata including: country, commodity labor/capital-intensiveness, commodity taxability, DV type, and conflict incompatibility type.

8 Risk of Bias in Individual Studies

We use a methodological filter, described above, chosen such that the risk of bias in individual studies is minimized by the common causal identification strategy used. We do not expect specific risks involved in this meta-analysis that are not common to all research.

9 Summary Measures

Our general model defines our principal summary measure: the coefficient on the value of natural resources variable in a generalized difference-in-differences (i.e., an extension of difference-in-differences to allow more than 2 time periods). We are comparing changes in conflict among units that experience differential price changes between areas that have natural resources and those that do not.

We expect the authors will typically report unstandardized regression coefficients; the formula used to implement standardized coefficients is:

$$\hat{\beta}_{std} = \hat{\beta} \times \frac{sd_x}{sd_y}$$

In designs with fixed effect structures, such as those we plan to analyze, we would prefer to use the mean of within-group standard deviations for these variables in order to better represent plausible changes in treatment intensity when interpreting overall effect sizes (Mummolo and Peterson 2018). However, we are not aware of any published studies where authors report sufficient data to calculate standardized coefficients in this manner. As a result, we are required to use pooled standard deviations, which are typically reported in study descriptive statistics tables.

In R, the following function implements the standardization:

```
standardize_coefficient = function(data) {
  data$effect_size * data$sd_x / data$sd_y
}
```

We expect it may be difficult or impossible to obtain the necessary two standard deviations to standardize effects for some studies. We will make all reasonable efforts, including analyzing replication data and contacting authors, to obtain these statistics or to estimate them on our own. We discontinue efforts to obtain replication or standardization data when authors request that their paper not be included, or when authors do not respond to good faith attempts to contact them. If we cannot obtain such data, we will exclude the study from analyses given that outcomes are likely to be on vastly different scales and thus uncomparable without standardization.

10 Synthesis of Results: Estimator

Given the differing numerical scales and data inputs in this literature, it is appropriate to compare standardized effect sizes (Cohen 1988; Morris and DeShon 2002). Authors typically report unstandardized regression coefficients, which in a regression context represent the change in the dependent variable associated with a one unit change in the predictor. Standardized effect sizes, by contrast, represent the change in standard deviations of the dependent variable associated with a change of one standard deviation in the treatment. Standardized effect sizes are thus unitless.

Having standardized regression coefficients, our primary estimator for our synthesized result is a Bayesian multilevel model. We have two goals with our estimator: the first is to produce a weighted average that takes into account the precision of the constituent studies; and the second is to formally account for clustering of results by country. In an appendix below, we will present simulation results in order to demonstrate that this result recovers accurate parameters assuming the data generating process.

Our model is:

$$\mu_{country} = \mu + \eta_{country}$$
$$\theta_{study} = \mu_{country} + \mathbf{X}\beta + \eta_{study}$$
$$Y \sim \mathbb{N}(\theta, \sigma^2)$$

$$\eta_{study} \sim \mathbb{N}(0, \tau^2)$$
$$\eta_{country} \sim \mathbb{N}(0, \gamma^2)$$

In this model, Y, σ^2, and $\mathbf{X}$ are observed data. Other parameters are estimated. Our four covariates $\mathbf{X}$ are: taxable, non-taxable, labor-intensive, and capital-intensive commodity types (each binary dummy variables).

We have a noninformative prior on each of our estimated parameters, μ, τ, and γ. Other parameters estimations are derived from these. For μ, we use a generally noninformative prior. For the other parameters, we use the default Stan priors, which are (improper) uniform priors over the space of admissible values. We constrain τ and γ to be non-negative, giving us the priors:

$$\tau \sim Unif(0, \infty)$$
$$\gamma \sim Unif(0, \infty)$$
$$\mu \sim \mathbb{N}(0, 10)$$

We fit this model using Stan via the **rstan** `R` package. The code is as follows:

```
data {
  int<lower=0> N_studies; // number of studies
  int<lower=0> N_countries; // number of countries (several studies per country)
  int<lower=0> N_covariates; // number of covariates
  int country[N_studies]; // country identifier (upper level of hierarchy)
  real y[N_studies]; // estimated effects
  matrix[N_studies, N_covariates] X; // matrix of data without intercept
  real<lower=0> sigma[N_studies]; // s.e. of effect estimates
  matrix[N_countries, N_studies] in_country;
}

parameters {
  real mu; // grand mean effect
  real<lower=0> gamma; // between-country variability
  real<lower=0> tau; // between-study variability (fixed across countries)
  real zeta[N_countries]; // random effect for each country
  real eta[N_studies]; // random effect for each study
  vector[N_covariates] mu_X; // coefficents on covariates
}
transformed parameters {
  real theta[N_studies]; // study effect prediction
  real tau_sq = tau^2;
  real mu_0[N_countries]; // country effect prediction
  real mu_country[N_countries];

  real muoverall;
  real mutaxablelabor;
  real mutaxablecapital;
  real mutaxablelaborcapital;
```

```
  real munontaxablelabor;
  real munontaxablecapital;
  real munontaxablelaborcapital;
  real mutaxablenontaxablelabor;
  real mutaxablenontaxablecapital;
  real mutaxablenontaxablelaborcapital;

  for (i in 1:N_countries)
    mu_0[i] = mu + gamma * zeta[i];
  for (j in 1:N_studies)
    theta[j] = mu_0[country[j]]+ row(X, j) * mu_X + tau * eta[j];
  for (k in 1:N_countries)
    mu_country[k] = mu_0[k] + 1/sum(row(in_country, k)) * row(in_country, k) * X * mu_X;

  muoverall = mu + mean(X * mu_X);

  mutaxablelabor                  = mu;
  mutaxablecapital                = mu + mu_X[1];
  mutaxablelaborcapital           = mu + mu_X[2];
  munontaxablelabor               = mu + mu_X[3];
  munontaxablecapital             = mu + mu_X[4];
  munontaxablelaborcapital        = mu + mu_X[5];
  mutaxablenontaxablelabor        = mu + mu_X[6];
  mutaxablenontaxablecapital      = mu + mu_X[7];
  mutaxablenontaxablelaborcapital = mu + mu_X[8];

}
model {
  mu ~ normal(0, 10);
  target += normal_lpdf(zeta | 0, 1);
  target += normal_lpdf(eta | 0, 1);
  target += normal_lpdf(y | theta, sigma);
}
```

We will run 4 chains of length 2,000 for each model. We will include a footnote reporting model diagnostics. We will report potential scale reduction factors ($\hat{R}$) and ensure each parameter of interest is below the standard 1.1 threshold (Gelman and Rubin 1992; Gelman and Rubin 2014). Failure to meet this threshold indicates non-convergence. We will also ensure the model does not produce "divergent transitions": an indication of nonconvergence due to trouble approximating the posterior function. If these diagnostics are not satisfactory after 2,000 draws, we will increase chain length to ensure convergence and document changes to the chain length specified above. In addition, we will ensure that effective sample size ratios for our parameters of interest pass the standard of $\hat{n}_{eff} >= 5m = 20$, where m is the number of chains (Gelman and Rubin 2014). Stan reports each of these statistics by default. We will not report traceplots or other visual diagnostics for our Bayesian model.

Our results rely on four summaries of the posterior for each model: the mean ("estimate"); the standard deviation ("standard deviation"); the 2.5th and 97.5th percentiles ("95% credible interval"); and an equivalence test value of the null hypothesis that the effect size is larger than .2 standard deviations in absolute value terms at the $\alpha = .05$ confidence level. Following Wellek (2010), we construct the hypothesis test by calculating the proportion of the posterior estimate that falls within the equivalence range between $-.2$ and .2. If the proportion of the posterior between the equivalence level is $\geq 1 - \alpha$ we reject nonequivalence. We construct analogous statistics for each of the subgroup analysis described, for taxable and nontaxable and labor- and capital-intensive commodities.

The code to prepare the data, fit the model in Stan, and calculate these statistics are included in the replication materials for this PAP.

This model is fit for ten outcome variables defined in Table 1:

- Conflict onset
- Conflict intensity
- Conflict duration
- Coups
- Social conflicts
- Conflict onset, territorial incompatibilities
- Conflict intensity, territorial incompatibilities
- Conflict duration, territorial incompatibilities
- Conflict onset, government incompatibilities
- Conflict intensity, government incompatibilities
- Conflict duration, government incompatibilities

We have two primary quantities of interest. First, the average effect of commodity prices on the outcome, i.e. $\mu + \overline{\text{X}\beta}$. Second, we report heterogeneous effects by commodity type, one effect for each pair of taxability { taxable, non-taxable, and combinations of both } and mode of production { labor-intensive, capital-intensive and combinations of both } (nine total combinations). If there is insufficient data to fit the model with all nine combinations (determined by an examination of the model fit), we will collapse them into a simpler model without interactions between the two variables taxabiliy and mode of production. Each effect is calculated as $\mu + \beta_{\text{pair}}$ (with the left-out category of taxable and labor-intensive calculated simply as μ).

In addition, of secondary interest, we present country mean effects in the appendix, calculated as follows:

$$\mu_{country-cov} = \mu_{country} + \frac{1}{n} \sum_{i \in country} \mathbf{X}_i \beta$$

11 Assessing the Meta-Analysis Research Design

In what follows, we describe the results of a simulation study to assess the properties of the Bayesian multilevel model we use as our estimator as well as the conventional approach of using a random effects weighted mean estimated via REML. The purpose of this study is to show unbiased recovery of coefficients under a range of controlled data-generating processes, as well as the pathologies of a single level random effects weighted mean when there is country-level clustering of effect sizes and variation in effects by type of commodity.

We use `DeclareDesign` (Blair et al. 2017) to conduct simulations of our design: each simulation consists of 250 iterations of drawing data from a data generating process, calculating estimator estimates using our Bayesian multilevel model and random effects weighted means, comparing these estimates to the desired estimand quantities, and recovering four diagnosand summaries of the model properties: estimation bias, root mean-squared error, and statistical coverage.

The first simulation we run compares the Bayesian multilevel model estimator to the random-effects weighted mean estimated via maximum likelihood when the true data generating process does not have country-level clustered effects or covariate intercept shifts. The second simulation compares the two under the presence of country-level clustered effects, but no covariates. The third compares the two under the presence of both country-level clustered effects and covariates. In each case, we use the Bayesian multilevel model to extract estimates for the grand mean; country-level heterogeneity (γ); within-country heterogeneity (τ); covariate effects (β); and country-level means. We use the random effects weighted mean estimator to extract estimates for the grand mean and between-study heterogeneity (τ). We use a simplified model with a single covariate for these simulations.

Scenario	Estimator	Coefficient	Bias	RMSE	Coverage
None	multilevel	mu	-0.001	0.016	0.956
None	multilevel	tau	-0.002	0.010	0.956
None	reml	mu	-0.001	0.011	0.956
None	reml	tau	-0.001	0.010	0.952
Country-level	multilevel	mu	-0.005	0.080	0.948
Country-level	multilevel	tau	0.000	0.010	0.944
Country-level	reml	mu	-0.005	0.079	0.360
Country-level	reml	tau	0.354	0.358	0.000
Country and Commodity Type	multilevel	mu	-0.002	0.081	0.940
Country and Commodity Type	multilevel	tau	0.000	0.010	0.948
Country and Commodity Type	reml	mu	0.163	0.182	0.032
Country and Commodity Type	reml	tau	0.369	0.372	0.000

From these results, we see that the model we adopt for this study, the Bayesian multilevel model, recovers both mu (the mean effect size) as well as tau (representing cross-study variation in effect sizes) without bias under all three data-generating processes described above. The frequentist coverage of the estimator is nominal in all three scenarios. In addition, we note that these characteristics are only shared with the standard random-effects estimator when there is no clustering of effect size either by country or commodity type. Importantly, given that we do not know whether in fact either type of true heterogeneity exists, our Bayesian multilevel model is unbiased, exhibits nominal coverage, and has low root mean squared error under the no-heterogeneity scenario in which the random effects model performs well. The RMSE is slightly lower for the REML model than the multilevel model for the μ mean effect size parameter; this should be expected given the large number of parameters estimated in the multilevel model. However, this difference is relatively small. For these reasons, we adopt the Bayesian multilevel model.

12 Risk of Bias Across Studies

We expect that null results have been published in this particular literature. There is a strong belief that prices do affect conflict, which has enabled the publication of prominent null results, including Bazzi and Blattman (2014). This does not preclude publication bias where particularly well-executed or considered null results are published, but far weaker positive results are published. As a result, we will address publication bias if it is present.

In addition to data gathering steps which might mitigate this issue (i.e., by including working papers) and our qualitative claim that prominent null results have indeed been published, we will conduct standard steps for identifying publication bias. We will present funnel plots and p-curves.

A funnel plot (Light and Pillemer 1984; Egger and Minder 1997) is used as a diagnostic tool to investigate the likelihood of publication bias (and the "file drawer problem"). We report a standard funnel plot from `metafor`. The expectation is that the observed points should support the inverted funnel shape plotted on the grid. In plots where there is an asymmetry in the plotted points, we have suspicion that publication bias may interfere with result inference. We present a funnel plot for our collected data.

A p-curve (Simonsohn and Simmons 2014) is a diagnostic tool made to evaluate the presence of either publication bias or p-hacking (wherein only those studies, or those analyses within studies, which pass a bright-line statistical significance test or published or submitted for publication). It relies on a basic assumption about the distribution of reported p-values. If no true effect exists, then the distribution of reported p-values below the traditional significance threshold of 0.05 will be uniform. When a true effect exists, the distribution will be right-skewed (with the lowest p-values having the most associated studies). The extent of the right skew depends on the power of the resulting studies, with well-powered studies of large effects having the largest skew. We present p-curves for each outcome and sample below.

13 Mock Analysis

In addition to specifying the general form of our analyses, we register mock analyses in order to make many small decisions about data analyses and visualization before data is collected. We expect minor aesthetic changes to the code and plots in the final analyses, and will catalog all changes in a pre-registration reconciliation document in an online appendix associated with the paper.

Our mock analysis makes use of the `R` package `DeclareDesign` (Blair et al. 2017), which enables researchers to simulate executing research design before data is collected and to diagnose its properties. The below analyses are all based on simulations of mock data. The accompanying Rmd file registered with this pre-analysis plan contains code describing the simulation of this data.

To clarify the analyses that follow, we present a snippet of the simulated data used to construct the mock data:[5]

Country	DV	Taxable?	Nontaxable?	Taxable and nontaxable?	Labor-intensive?	Capital-intensive?	Labor- and capital-intensive?	Incompatibility	Effect Size	Effect SE	p-value
39	intensity	0	0	1	1	0	0	government	4.30	0.45	0.00
25	onset	0	1	0	1	0	0	government	-0.22	0.72	0.76
27	social	1	0	0	0	1	0	government	0.77	0.56	0.17
30	duration	0	0	1	0	0	1	both	0.80	0.38	0.04
16	onset	1	0	0	0	0	1	government	2.81	0.92	0.00
46	duration	0	0	1	0	1	0	both	5.91	0.17	0.00
16	onset	0	0	1	1	0	0	both	6.47	1.73	0.00
06	coups	0	0	1	0	1	0	territory	2.07	0.80	0.01
42	social	1	0	0	0	0	1	territory	1.85	0.49	0.00
07	duration	0	1	0	0	0	1	government	0.05	1.01	0.96

[5] The data will not be presented in this way in the paper.

13.1 Mock summary table

We will report primary results in a table similar to the format of Table 1. We summarize some relationships with + in Table 1, but present each possible estimated effect size in this table. To do so, we construct effect summaries, with: + representing a positive effect for which the null hypothesis of nonequivalence is rejected; - representing a negative effect for which the null hypothesis of nonequivalence is rejected; and 0 representing an effect for which we fail to reject the null hypothesis of nonequivalence. In addition to the hypotheses we test about armed conflict, we also present average effects of commodity price changes on social conflict.

	Armed Conflict			
	Onset	Intensity	Duration	Coups
Topline results	+	+	-	0

	Incompatibility	
	Territory	Government
Onset	+	-
Intensity	+	-
Duration	+	-

	Heterogeneous Effects on Onset			
	Labor-intensive	Capital-intensive	Combination	Difference (Labor - capital)
Taxable	-	-	-	0
Non-taxable	+	-	0	+
Combination	0	0	0	-

	Heterogeneous Effects on Intensity			
	Labor-intensive	Capital-intensive	Combination	Difference (Labor - capital)
Taxable	-	-	-	0
Non-taxable	+	0	0	+
Combination	0	+	-	-

	Heterogeneous Effects on Duration			
	Labor-intensive	Capital-intensive	Combination	Difference (Labor - capital)
Taxable	+	0	-	0
Non-taxable	-	+	0	-
Combination	0	0	-	+

Table 2: Results summary. In the top panel, we present results for all studies for the four topline outcomes. In the bottom panel, we examine at left the conflict outcomes for each incompatibility. At right in the bottom panel we examine the effects by commodity type for each topline conflict outcome. Legend: +, positive effect; $-$, negative effect; 0, null effect based on test of equivalence.

13.2 Mock Figures

Below we include mock examples of the figures that will result from our analysis.

First, we present plots of overall effects (μ in our model) for each outcome, where conflict can be either over territories or governments.

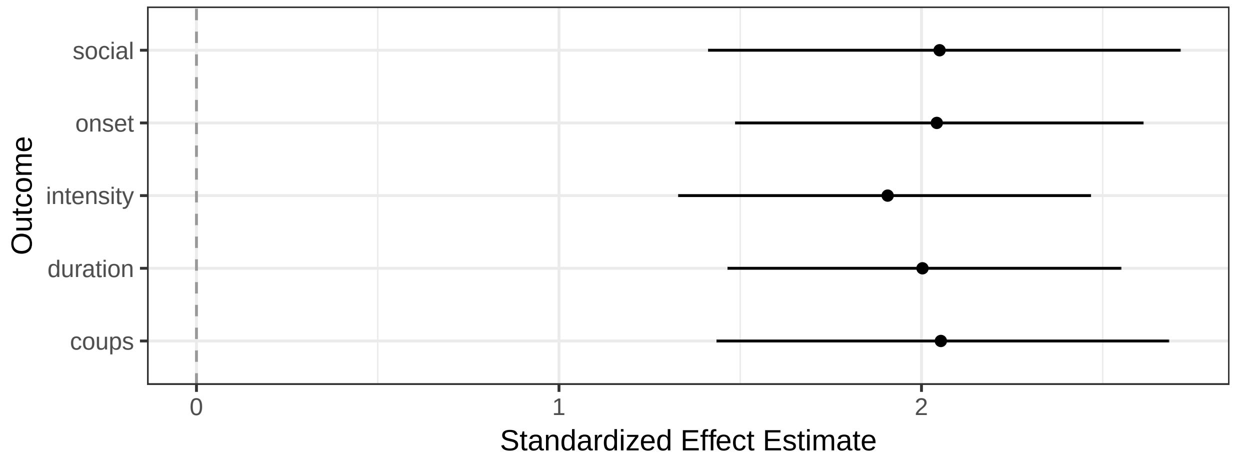

Figure 1: Average Effects of Commodity Prices on Armed Conflict and Coups.

Next, we present a plot of each conflict outcome disambiguated by source of incompatibility.

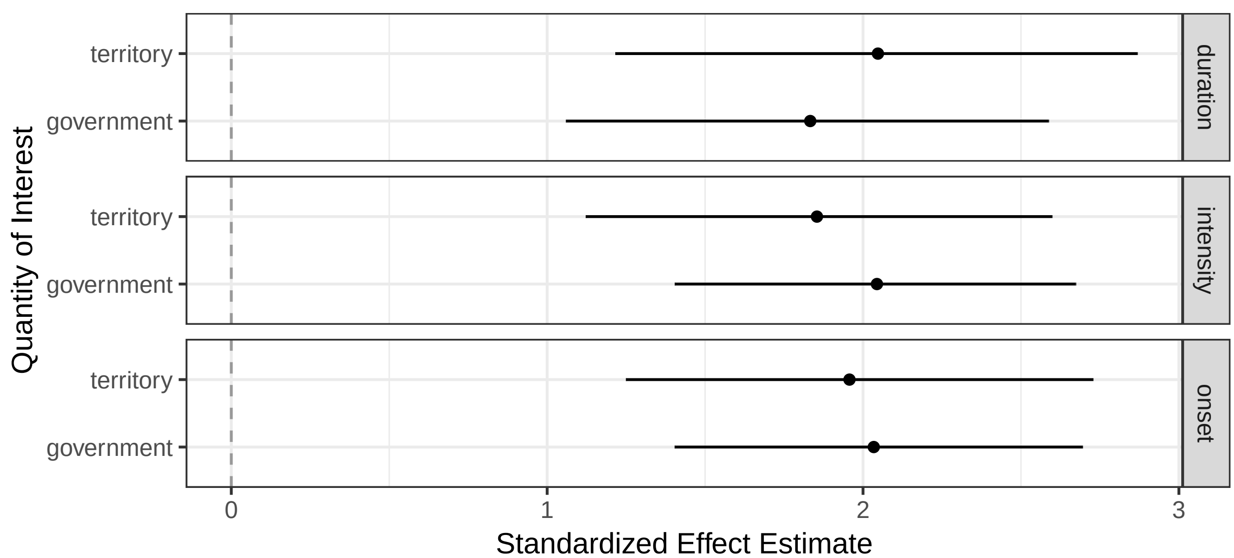

Figure 2: Heterogeneous effects of commodity prices on armed conflict by conflict incompatibility, distinguishing between conflicts over territory and government control.

Next, we present a plot of the effect of on conflict for each combination of taxability status (taxable, nontaxable, or a combination) and mode of production (labor-, capital-intensive, or a combination).

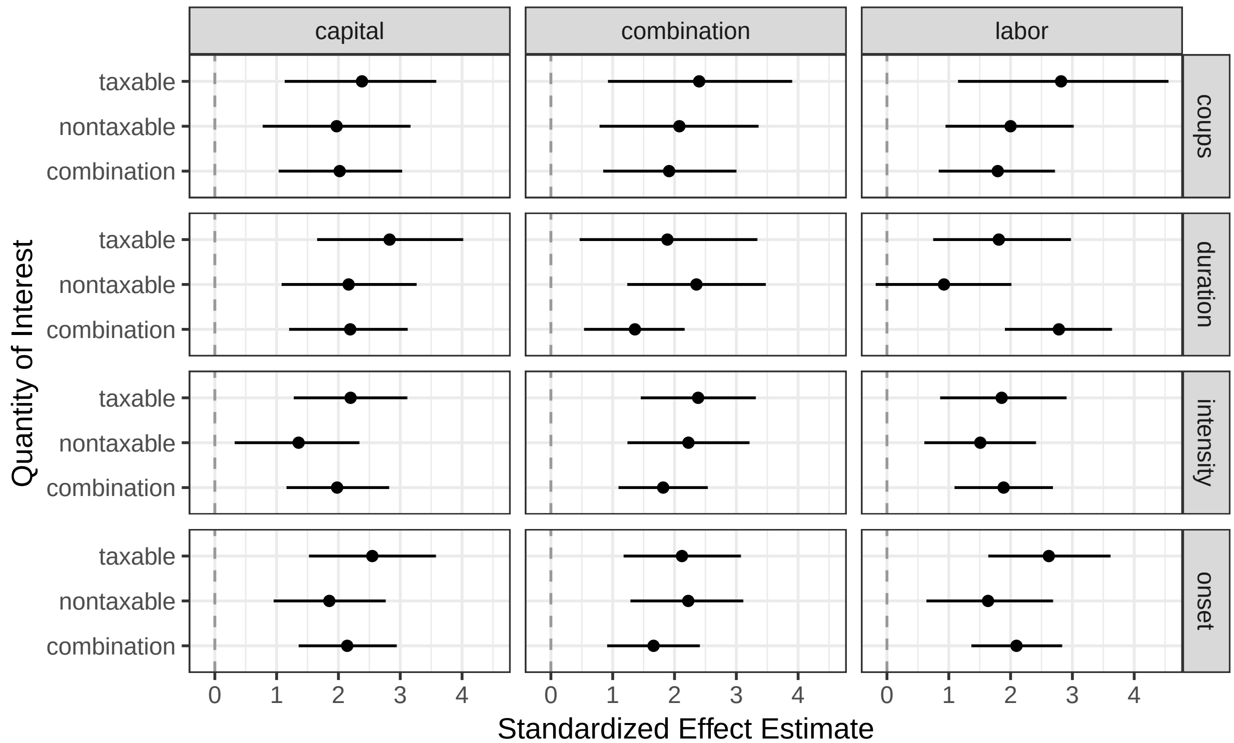

Figure 3: Heterogeneous effects of commodity prices on armed conflict and coups by factor intensity of the production of the commodity.

13.3 Mock Publication Bias Analysis

p-curves. Here we present the p-curve across all outcomes.

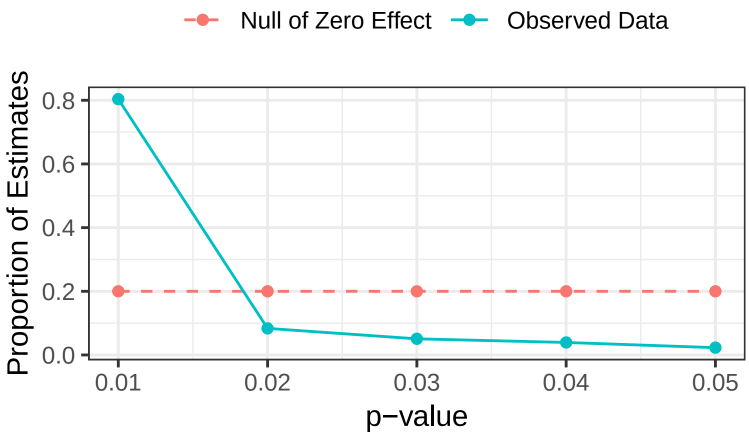

Figure 4: Mock P-Curve Plot

Funnel Plots. Here we present a funnel plot across all outcomes.

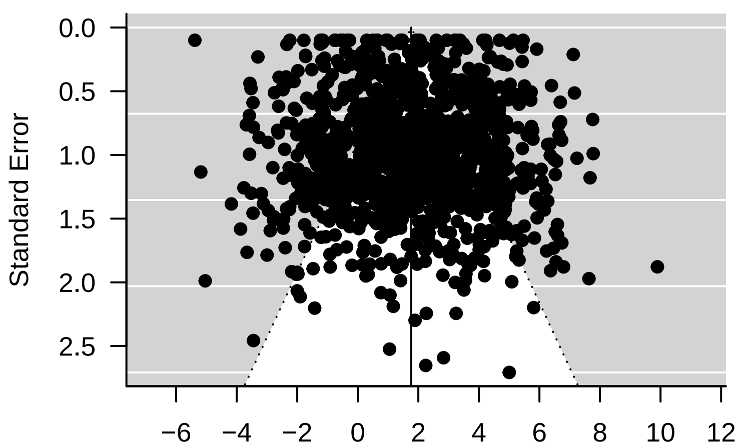

Figure 5: Funnel Plot

14 Appendix

14.1 Coding Rules

Our hypotheses require us to classify models along three dimensions:

1. Are the conflicts over control of the central government, or control of territory?
2. Do prices relate to commodities that are labor- or capital-intensive?; and
3. Do prices relate to commodities that are taxable or non-taxable?

Below, we outline a set of rules for classifying models. Specifying a coding scheme ex ante that includes all cases is difficult; if necessary, we will submit amendments to this PAP to incorporate revisions:

- **Government vs. Territory**: We employ the distinction that UCDP/PRIO codebook makes between incompatibility related to government vs. territory (Themnér 2014). An incompatibility concerning government is "Incompatibility concerning type of political system, the replacement of the central government, or the change of its composition." An incompatibility concerning territory is "Incompatibility concerning the status of a territory, e.g. the change of the state in control of a certain territory (interstate conflict), secession or autonomy (internal conflict)." We will be able to classify along this dimension when models employ a dependent variable that relates to one or the other incompatibility *or* for single-country studies where we can rely on the UCDP coding of the case.
- **Labor vs. Capital Intensive**: We will separate commodities into three groups: (1) oil and natural gas; (2) agricultural commodities; and (3) minerals, including both artisanal and commercial mining. We consider the first group capital-intensive and the latter two groups to be labor-intensive. We will not be able to classify along this dimension when the price measure relates to commodities that are both capital and labor intensive.
- **Taxable vs. Non-taxable**: Taxable is shorthand for commodities the *state* can profit from. We code a commodity as non-taxable if it is illicit (e.g., drugs, illegal mining) or produced on a small scale (e.g., artisanal mining or small-holder agriculture); otherwise, a commodity is considered taxable.

14.2 Mock Results

Table 3: Results for Main Outcomes

Outcome	Coefficient	Estimate	Std. Err.	CI (Lower)	CI (Upper)	Equivalence Prop.
Onset	Non-Taxable, Capital-Intensive Commodities	1.852	0.462	0.953	2.758	0.000
Onset	Non-Taxable, Labor-Intensive Commodities	1.635	0.517	0.640	2.681	0.002
Onset	Non-Taxable, Labor-and-Capital-Intensive Commodities	2.222	0.459	1.289	3.106	0.000
Onset	Overall Effect	2.042	0.292	1.486	2.611	0.000
Onset	Taxable, Capital-Intensive Commodities	2.547	0.525	1.524	3.571	0.000
Onset	Taxable, Labor-Intensive Commodities	2.619	0.508	1.640	3.610	0.000
Onset	Taxable, Labor-and-Capital-Intensive Commodities	2.120	0.481	1.178	3.068	0.000
Onset	Taxable-and-Nontaxable, Capital-Intensive Commodities	2.142	0.401	1.358	2.937	0.000
Onset	Taxable-and-Nontaxable, Labor-Intensive Commodities	2.095	0.381	1.367	2.828	0.000
Onset	Taxable-and-Nontaxable, Labor-and-Capital-Intensive Commodities	1.661	0.380	0.911	2.404	0.000
Intensity	Non-Taxable, Capital-Intensive Commodities	1.356	0.504	0.324	2.333	0.012
Intensity	Non-Taxable, Labor-Intensive Commodities	1.510	0.465	0.606	2.405	0.003
Intensity	Non-Taxable, Labor-and-Capital-Intensive Commodities	2.225	0.494	1.240	3.207	0.000
Intensity	Overall Effect	1.907	0.285	1.329	2.466	0.000
Intensity	Taxable, Capital-Intensive Commodities	2.197	0.465	1.279	3.108	0.000
Intensity	Taxable, Labor-Intensive Commodities	1.856	0.522	0.862	2.899	0.000
Intensity	Taxable, Labor-and-Capital-Intensive Commodities	2.382	0.467	1.456	3.308	0.000
Intensity	Taxable-and-Nontaxable, Capital-Intensive Commodities	1.979	0.420	1.161	2.814	0.000
Intensity	Taxable-and-Nontaxable, Labor-Intensive Commodities	1.887	0.405	1.095	2.677	0.000
Intensity	Taxable-and-Nontaxable, Labor-and-Capital-Intensive Commodities	1.816	0.360	1.096	2.531	0.000
Duration	Non-Taxable, Capital-Intensive Commodities	2.165	0.553	1.081	3.257	0.000
Duration	Non-Taxable, Labor-Intensive Commodities	0.922	0.552	-0.181	2.006	0.072
Duration	Non-Taxable, Labor-and-Capital-Intensive Commodities	2.355	0.560	1.237	3.470	0.000
Duration	Overall Effect	2.003	0.276	1.465	2.550	0.000
Duration	Taxable, Capital-Intensive Commodities	2.827	0.594	1.657	4.012	0.000
Duration	Taxable, Labor-Intensive Commodities	1.809	0.564	0.749	2.970	0.004
Duration	Taxable, Labor-and-Capital-Intensive Commodities	1.885	0.723	0.466	3.334	0.008
Duration	Taxable-and-Nontaxable, Capital-Intensive Commodities	2.191	0.486	1.204	3.114	0.000
Duration	Taxable-and-Nontaxable, Labor-Intensive Commodities	2.781	0.445	1.908	3.634	0.000
Duration	Taxable-and-Nontaxable, Labor-and-Capital-Intensive Commodities	1.359	0.411	0.538	2.156	0.002
Coups	Non-Taxable, Capital-Intensive Commodities	1.971	0.606	0.777	3.159	0.001
Coups	Non-Taxable, Labor-Intensive Commodities	2.000	0.533	0.950	3.014	0.001
Coups	Non-Taxable, Labor-and-Capital-Intensive Commodities	2.079	0.643	0.789	3.353	0.002
Coups	Overall Effect	2.053	0.318	1.435	2.682	0.000
Coups	Taxable, Capital-Intensive Commodities	2.381	0.615	1.132	3.576	0.000
Coups	Taxable, Labor-Intensive Commodities	2.817	0.879	1.152	4.547	0.001
Coups	Taxable, Labor-and-Capital-Intensive Commodities	2.399	0.769	0.925	3.896	0.002
Coups	Taxable-and-Nontaxable, Capital-Intensive Commodities	2.019	0.500	1.035	3.023	0.000
Coups	Taxable-and-Nontaxable, Labor-Intensive Commodities	1.792	0.478	0.838	2.711	0.001
Coups	Taxable-and-Nontaxable, Labor-and-Capital-Intensive Commodities	1.912	0.544	0.846	2.996	0.001
Social	Non-Taxable, Capital-Intensive Commodities	2.132	0.653	0.852	3.443	0.002
Social	Non-Taxable, Labor-Intensive Commodities	1.228	0.622	0.026	2.474	0.037
Social	Non-Taxable, Labor-and-Capital-Intensive Commodities	-0.011	0.907	-1.834	1.713	0.176
Social	Overall Effect	2.050	0.323	1.412	2.714	0.000
Social	Taxable, Capital-Intensive Commodities	2.818	0.550	1.719	3.875	0.000
Social	Taxable, Labor-Intensive Commodities	2.368	0.681	1.028	3.709	0.000
Social	Taxable, Labor-and-Capital-Intensive Commodities	2.904	0.655	1.609	4.176	0.000
Social	Taxable-and-Nontaxable, Capital-Intensive Commodities	2.106	0.504	1.123	3.097	0.000
Social	Taxable-and-Nontaxable, Labor-Intensive Commodities	1.932	0.503	0.949	2.927	0.000
Social	Taxable-and-Nontaxable, Labor-and-Capital-Intensive Commodities	2.026	0.513	1.031	3.048	0.000

Table 4: Results for Territorial Incompatibility Outcomes

Outcome	Coefficient	Estimate	Std. Err.	CI (Lower)	CI (Upper)	Equivalence Prop.
Onset	Non-Taxable, Capital-Intensive Commodities	1.531	0.883	-0.209	3.275	0.035
Onset	Non-Taxable, Labor-Intensive Commodities	2.351	1.130	0.066	4.567	0.018
Onset	Non-Taxable, Labor-and-Capital-Intensive Commodities	2.667	0.854	1.027	4.409	0.002
Onset	Overall Effect	1.957	0.373	1.250	2.728	0.000
Onset	Taxable, Capital-Intensive Commodities	3.457	1.131	1.233	5.790	0.002
Onset	Taxable, Labor-Intensive Commodities	2.395	1.087	0.316	4.568	0.013
Onset	Taxable, Labor-and-Capital-Intensive Commodities	3.844	1.382	1.221	6.666	0.004
Onset	Taxable-and-Nontaxable, Capital-Intensive Commodities	1.651	0.759	0.162	3.139	0.022
Onset	Taxable-and-Nontaxable, Labor-Intensive Commodities	0.837	0.702	-0.550	2.241	0.118
Onset	Taxable-and-Nontaxable, Labor-and-Capital-Intensive Commodities	1.790	0.733	0.373	3.258	0.010
Intensity	Non-Taxable, Capital-Intensive Commodities	1.241	1.335	-1.351	3.849	0.084
Intensity	Non-Taxable, Labor-Intensive Commodities	1.192	0.824	-0.453	2.790	0.066
Intensity	Non-Taxable, Labor-and-Capital-Intensive Commodities	0.562	0.884	-1.182	2.296	0.147
Intensity	Overall Effect	1.854	0.374	1.122	2.598	0.000
Intensity	Taxable, Capital-Intensive Commodities	2.205	0.868	0.513	3.943	0.007
Intensity	Taxable, Labor-Intensive Commodities	3.348	1.442	0.578	6.215	0.005
Intensity	Taxable, Labor-and-Capital-Intensive Commodities	3.062	1.127	0.853	5.340	0.003
Intensity	Taxable-and-Nontaxable, Capital-Intensive Commodities	1.872	0.748	0.390	3.376	0.010
Intensity	Taxable-and-Nontaxable, Labor-Intensive Commodities	1.498	0.718	0.079	2.923	0.026
Intensity	Taxable-and-Nontaxable, Labor-and-Capital-Intensive Commodities	2.513	0.688	1.125	3.862	0.000
Duration	Non-Taxable, Capital-Intensive Commodities	1.765	1.136	-0.375	4.060	0.041
Duration	Non-Taxable, Labor-Intensive Commodities	0.735	1.065	-1.321	2.828	0.119
Duration	Non-Taxable, Labor-and-Capital-Intensive Commodities	0.636	1.101	-1.501	2.765	0.120
Duration	Overall Effect	2.047	0.422	1.216	2.868	0.000
Duration	Taxable, Capital-Intensive Commodities	2.099	0.939	0.269	3.950	0.016
Duration	Taxable, Labor-Intensive Commodities	1.274	1.507	-1.837	4.084	0.070
Duration	Taxable, Labor-and-Capital-Intensive Commodities	3.485	1.468	0.679	6.479	0.006
Duration	Taxable-and-Nontaxable, Capital-Intensive Commodities	2.052	0.856	0.274	3.661	0.013
Duration	Taxable-and-Nontaxable, Labor-Intensive Commodities	2.951	0.755	1.459	4.474	0.000
Duration	Taxable-and-Nontaxable, Labor-and-Capital-Intensive Commodities	2.124	1.053	0.042	4.139	0.018

Table 5: Results for Government Incompatibility Outcomes

Outcome	Coefficient	Estimate	Std. Err.	CI (Lower)	CI (Upper)	Equivalence Prop.
Onset	Non-Taxable, Capital-Intensive Commodities	2.155	0.708	0.785	3.556	0.002
Onset	Non-Taxable, Labor-Intensive Commodities	0.760	0.803	-0.847	2.338	0.119
Onset	Non-Taxable, Labor-and-Capital-Intensive Commodities	2.192	0.742	0.740	3.616	0.003
Onset	Overall Effect	2.034	0.335	1.404	2.695	0.000
Onset	Taxable, Capital-Intensive Commodities	1.489	0.814	-0.160	3.055	0.038
Onset	Taxable, Labor-Intensive Commodities	2.450	0.827	0.882	4.054	0.002
Onset	Taxable, Labor-and-Capital-Intensive Commodities	2.257	0.625	0.982	3.450	0.001
Onset	Taxable-and-Nontaxable, Capital-Intensive Commodities	2.106	0.749	0.622	3.592	0.006
Onset	Taxable-and-Nontaxable, Labor-Intensive Commodities	2.625	0.519	1.601	3.661	0.000
Onset	Taxable-and-Nontaxable, Labor-and-Capital-Intensive Commodities	1.662	0.571	0.508	2.760	0.005
Intensity	Non-Taxable, Capital-Intensive Commodities	1.654	0.822	0.002	3.244	0.031
Intensity	Non-Taxable, Labor-Intensive Commodities	0.820	0.730	-0.570	2.247	0.127
Intensity	Non-Taxable, Labor-and-Capital-Intensive Commodities	3.383	0.853	1.692	5.082	0.000
Intensity	Overall Effect	2.044	0.324	1.404	2.674	0.000
Intensity	Taxable, Capital-Intensive Commodities	2.378	0.693	1.064	3.752	0.001
Intensity	Taxable, Labor-Intensive Commodities	1.169	0.787	-0.375	2.710	0.070
Intensity	Taxable, Labor-and-Capital-Intensive Commodities	2.515	0.626	1.302	3.742	0.000
Intensity	Taxable-and-Nontaxable, Capital-Intensive Commodities	2.094	0.648	0.824	3.360	0.002
Intensity	Taxable-and-Nontaxable, Labor-Intensive Commodities	2.060	0.564	0.944	3.148	0.001
Intensity	Taxable-and-Nontaxable, Labor-and-Capital-Intensive Commodities	2.022	0.568	0.937	3.137	0.000
Duration	Non-Taxable, Capital-Intensive Commodities	4.131	1.079	1.993	6.211	0.001
Duration	Non-Taxable, Labor-Intensive Commodities	1.051	0.912	-0.732	2.833	0.092
Duration	Non-Taxable, Labor-and-Capital-Intensive Commodities	2.056	0.853	0.367	3.708	0.012
Duration	Overall Effect	1.833	0.393	1.060	2.588	0.000
Duration	Taxable, Capital-Intensive Commodities	3.032	1.213	0.611	5.381	0.006
Duration	Taxable, Labor-Intensive Commodities	1.543	0.732	0.059	2.932	0.026
Duration	Taxable, Labor-and-Capital-Intensive Commodities	2.981	2.197	-1.238	7.332	0.028
Duration	Taxable-and-Nontaxable, Capital-Intensive Commodities	2.199	0.853	0.549	3.907	0.008
Duration	Taxable-and-Nontaxable, Labor-Intensive Commodities	2.302	0.757	0.822	3.785	0.002
Duration	Taxable-and-Nontaxable, Labor-and-Capital-Intensive Commodities	0.553	0.788	-0.988	2.116	0.150

14.3 Version History

On **January 30, 2019**, after completing our pre-specified searches but before data analysis, we amended this document to clarify three points.

1. In the section titled "Data Collection Process", we added the following paragraph in order to clarify an ambiguity in our original data collection process:

   > *Some papers report different model link functions. If a paper reports models which otherwise are identical on the above criteria, but differ in the link function they use, we select linear models, including linear probability models, in order to, when possible, compare estimated effect sizes on a common scale.*

2. In the section titled "Data Items", we revised the penultimate paragraph and added the following sentences in order to clarify details about the standardization data we collected:

   > *Where the model selected uses a non-linear link function, we seek replication data sufficient to re-estimate the model on the same data with a linear link function and contact authors if no data is publicly available.*

3. In the section titled "Summary Measures", we revised the final paragraph and added the following sentences in order to clarify what "reasonable measures" to obtain standardization data entailed (non-italicized text in original version):

   > We will make all reasonable efforts, including analyzing replication data and contacting authors, to obtain these statistics *or to estimate them on our own. We discontinue efforts to obtain replication or standardization data when authors request that their paper not be included, or when authors do not respond to good faith attempts to contact them.*

References

Bazzi, Samuel, and Christopher Blattman. 2014. “Economic Shocks and Conflict: Evidence from Commodity Prices.” *American Economic Journal: Macroeconomics* 6 (4).

Becker, Gary S. 1968. “Crime and Punishment: An Economic Approach.” *Journal of Political Economy* 76: 169–217.

Blair, Graeme, Jasper Cooper, Alexander Coppock, and Macartan Humphreys. 2017. “Declaring and Diagnosing Research Designs.” Working Paper. http://declaredesign.org/paper.pdf.

Cederman, Lars-Erik, Kristian Skrede Gleditsch, and Halvard Buhaug. 2013. *Inequality, Grievances, and Civil War.* Cambridge: Cambridge University Press.

Christensen, Darin. 2018. “Concession Stands: How Mining Investments Incite Protest in Africa.” *International Organization* Forthcoming.

Cohen, Jacob. 1988. *Statistical Power Analysis for the Behavioral Sciences.* 2nd ed. Lawrence Erlbaum Associates.

Collier, Paul, and Anke Hoeffler. 2004. “Greed and Grievance in Civil War.” *Oxford Economic Papers* 56 (4). Oxford University Press: 563–95.

Cotet, Anca M., and Kevin K. Tsui. 2013. “Oil and Conflict: What Does the Cross Country Evidence Really Show?” *American Economic Journal: Macroeconomics* 5 (1): 49–80.

Dal Bó, Ernesto, and Pedro Dal Bó. 2011. “Workers, Warriors, and Criminals: Social Conflict in General Equilibrium.” *Journal of European Economic Association* 9 (4).

Dube, Oeindrilla, and Juan Vargas. 2013. “Commodity Price Shocks and Civil Conflict: Evidence from Columbia.” *Review of Economic Studies* 80 (4).

Egger, George Smith, Matthias, and Christoph Minder. 1997. “Bias in Meta-Analysis Detected by a Simple, Graphical Test.” *BMJ* 315.

el-Beblawi, Hazem, and Giacomo Luciani. 1987. *The Rentier State.* Routledge.

Fearon, James D. 2005. “Primary Commodity Exports and Civil War.” *Journal of Conflict Resolution* 49 (4). Sage Publications Sage CA: Thousand Oaks, CA: 483–507.

Fearon, James D., and David Laitin. 2003. “Ethnicity, Insurgency, and Civil War.” *American Political Science Review* 101 (1).

Gelman, Andrew, and Donald Rubin. 1992. “Inference from Iterative Simulation Using Multiple Sequences.” *Statistical Science* 7 (4).

Gelman, John Carlin, Andrew, and Donald Rubin. 2014. *Bayesian Data Analysis.* 3rd ed. CRC Press.

Grossman, Herschell. 1991. “A General Equilibrium Model of Insurrections.” *American Economic Review* 81 (4).

Humphreys, Macartan. 2005. “Natural Resources, Conflict, and Conflict Resolution: Uncovering the

Mechanisms." *Journal of Conflict Resolution* 49 (4): 508–37.

Keen, David. 1998. "The Economic Functions of Violence in Civil Wars." *Adelphi Papers* 38 (320): 1–89.

Light, Richard, and David Pillemer. 1984. *Summing up: The Science of Reviewing Research.* Harvard University Press.

Moher, David, Alessandro Liberati, Jennifer Tetzlaff, Douglas G. Altman, and The PRISMA Group. 2009. "Preferred Reporting Items for Systematic Reviews and Meta-Analyses: The Prisma Statement." *PLoS Med* 6 (7).

Morris, Scott, and Richard DeShon. 2002. "Combining Effect Size Estimates in Meta-Analysis with Repeated Measures and Independent-Groups Designs." *Psychological Methods* 7 (1).

Mummolo, Jonathan, and Erik Peterson. 2018. "Improving the Interpretation of Fixed Effects Regression Results." *Political Science Research and Methods.*

Paine, Jack. 2016. "Rethinking the Conflict 'Resource Curse': How Oil Wealth Prevents Center-Seeking Civil Wars." *International Organization* 70 (4): 727–61.

Ross, Michael L. 2001. "Does Oil Hinder Democracy?" *World Politics* 53 (3).

———. 2004a. "How Do Natural Resources Influence Civil War? Evidence from Thirteen Cases." *International Organization* 58 (1): 35–67.

———. 2004b. "What Do We Know About Natural Resources and Civil War?" *Journal of Peace Research* 41 (3): 337–56.

———. 2012. *The Oil Curse: How Petroleum Wealth Shapes the Development of Nations.* Princeton University Press.

Simonsohn, Leif Nelson, Uri, and Joseph Simmons. 2014. "P-Curve: A Key to the File Drawer." *Journal of Experimental Psychology: General* 143 (2).

Snyder, Richard, and Ravi Bhavnani. 2005. "Diamonds, Blood, and Taxes: A Revenue-Centered Framework for Explaining Political Order." *Journal of Conflict Resolution* 49 (4): 563–97.

Soysa, Indra de. 2002. "Ecoviolence: Shrinking Pie, or Honey Pot?" *Global Environmental Politics* 2 (4): 1–34.

Themnér, Lotta. 2014. "UCDP/Prio Armed Conflict Dataset Codebook." *Uppsala Conflict Data Program (UCDP).*

Wellek, Stefan. 2010. *Testing Statistical Hypotheses of Equivalence and Noninferiority.* Chapman; Hall/CRC.