1 Introduction

A substantial body of scholarship developed over the past two decades argues that, conditional on the presence or magnitude of oil resources, a state is more likely to be autocratic, to have weak bureaucratic institutions, to experience civil conflict, and to suffer from economic misfortunes, especially slower rates of growth. Many hundreds of books and articles have been written about the resource curse. Michael Ross's (2001) landmark study has almost 4,000 citations, and dozens of other prominent works have been cited hundreds or thousands of times. Outside the academy, commentators on global politics and public policy frequently refer to the deleterious consequences of oil (Birdsall and Subramanian 2004; Friedman 2004, 2009).

In this Cambridge Element, we focus on the political resource curse – the claim that oil and democracy are largely antithetical to one another. We do not, and cannot, completely ignore the other elements of the resource curse literature: slower rates of economic growth, higher likelihoods of civil war, and a propensity for corruption and weak state institutions. These diverse phenomena are plausibly related to one another and so we often find scholarship considering multiple outcomes in a single study. Even scholarship that focuses directly on the political resource curse may borrow theories, concepts, methods, models, and data from scholarship that focuses on economic growth or civil war: we cannot always strictly sequester the political resource curse from other hypothesized consequences of oil wealth.

But despite occasionally trespassing into other outcomes, in what follows, we primarily discuss the large literature and derived scholarly consensus that oil wealth creates an environment inhospitable to the flourishing of democracy. The main claim we consider is that oil either degrades democracy and induces autocracy or contributes energetically to autocratic survival; or, perhaps, generates both outcomes. We concede from the beginning that an enormous wealth of scholarly material supports the claim of a political resource curse, and we exert considerable energy documenting the basis of that support. If we drill below the surface consensus that a political resource curse exists, however, we quickly hit multiple strata of debate and dissensus. While giving due weight to the reasons why scholars are convinced of the reality of a political resource curse, much of the first two sections of this Element survey various dimensions along which scholars disagree.

A first layer of debate is conceptual. The political resource curse has been defined and measured in diverse ways, not all of which are mutually consistent with one another. The political resource curse has been understood as causing regimes to become less democratic, as causing autocratic regimes whose origins

were independent of oil to remain autocratic longer than non-oil autocracies, or as causing a broader phenomenon of political stability; and each of these conceptual understandings of the political resource curse can be measured in diverse ways. All of these works can be understood as supporting a claim of a political resource curse, even as they disagree, to a lesser or greater extent, with one another.

A second layer of debate is methodological. Scholars choose from a rich menu of options, starting with concepts and measures, but then moving to data sets, statistical models and assumptions, and research designs to make causal inferences from observational data. Some of these decisions appeared credible when they were first made, but, in retrospect, have been shown to yield non-credible results. Other decisions produce noncomparable findings. Consequently, while there is an enormous literature claiming to find evidence of a resource curse, these results are not necessarily cumulative and do not necessarily replicate and reinforce one another.

A third layer of debate is theoretical. There are multiple theories of the resource curse, often based on incommensurable theoretical premises but still all yielding the same testable empirical hypothesis of a negative relationship between oil wealth, however understood and measured, and democracy, however understood and measured. Even among scholars who agree that a political resource curse exists, therefore, we find substantial disagreement about how to explain it theoretically. Furthermore, several prominent scholars have presented compelling theoretical arguments that the political resource curse is conditional on a broader set of theoretical factors, such that curse-like phenomena will be observed only in specific regions of the overall parameter space.

Returning to the surface after digging through these conceptual, methodological, and theoretical strata, the initial consensus around the reality of the political resource curse appears less well established. While the bulk of published research claims to find evidence for a political resource curse, substantial and very credible analyses find either a null effect, a highly conditional effect, or, in the distinct but still important minority of cases, evidence for a resource blessing. Furthermore, even among those scholars who support the claim of a political resource curse, there are divisions over causal heterogeneity. Perhaps, as we explore further, the political resource curse exists in only specific times – after the great price hike of the 1970s, for example. Or perhaps the political resource curse becomes manifest only in specific places (and perhaps, even then, only in specific times). After all, the United States and Canada industrialized and became major economic powers in significant part because of the late-nineteenth- and early-twentieth-century influx of resource revenues, retaining democratic politics throughout. Within the developing world, decades of

Venezuelan democracy were arguably financed by oil revenues, while more recently, major oil producers and long-standing dictatorships in Indonesia and Mexico both democratized at the end of the last century.

We do not claim that there is a single "best-answer" to all of the disagreements over how to study the resource curse; neither do we claim that the existence of diverse findings and debates over concepts, methods, and theories are grounds for the theory's rejection. We comment on particular theoretical, conceptual, and methodological debates and offer suggestions; but diversity itself is no cause for concern. We do claim, however, that when one looks closely at specific claims made in the literature, and the specific methodological warrants for those claims, the conclusion that a political resource curse exists appears less unassailable than would be the case when all the diverse findings are pooled together without sufficiently discriminating appraisal.

Beyond documenting scholarly dissensus over how to study the resource curse and the findings that result, our primary goal in this Element is to engage with evidence of causal heterogeneity in the relationship of oil and politics across time and space. We also advance several new claims about this spatial and temporal heterogeneity. With reference to spatial heterogeneity, we provide evidence that the treatment effects of oil vary across at least four regions – Latin America, the Middle East, Africa, and Southeast Asia – to a greater degree than most scholarship has recognized.[1] We demonstrate substantial regional heterogeneity across the developing world since the early 1980s, with Latin American oil producers becoming democratic, Middle Eastern oil producers remaining staunchly autocratic, and distinctively mixed patterns in Africa, where most oil producers made some progress toward democracy before stalling and, in some cases, experiencing retrogression, and Southeast Asia, where the cases divide cleanly between democratic transitions and obdurate dictatorships. To explain this inter-regional diversity, we draw on institutional and coalitional theories of political regimes, presenting some evidence that the structure of autocratic institutions and the underlying coalitional basis of autocratic regimes explains why, for example, we observe autocracies thriving among Middle Eastern oil producers to a much greater degree than oil producers elsewhere.

Our two other major claims represent relatively stark deviations from the conventional wisdom about the political resource curse. First, we argue that insofar as there is a political resource curse, it is overwhelmingly a regional phenomenon restricted to the major oil-producing monarchies of the Arabian Peninsula. We argue that the small oil principalities along the eastern shore of the Arabian Peninsula represent instances of survivorship bias; when we correct

[1] In Section 4, we explain why we emphasize diversity in these four regions.

for this source of bias, we fail to find evidence for a political resource curse. Second, with reference to temporal heterogeneity, we advance evidence that during the recent Third Wave of democratization, oil may very well be more of a blessing than a curse, as it appears to aid democratic consolidation in at least some parts of the world.

It is not our claim that oil has absolutely no effect; we are confident that it does, as the new work we present documents. But it is not the uniform effect prevalent in the literature and it is not an effect that, in our opinion, supports a more general idea of a resource curse; if anything, the cumulative evidence we present is more consistent with the idea of relatively circumscribed enclaves of a potential resource curse and a larger region in which oil might be a modest resource blessing.

We write this Cambridge Element with the goal of documenting diversity and dissensus, diagnosing its sources, and directing scholarly attention toward what we believe will be more fruitful avenues of future research and knowledge accumulation. To achieve that goal, we have divided the main body of this Element into four major sections. We begin with a survey of the field, working chronologically to show how the study of the political resource curse has changed – conceptually, theoretically, and methodologically – over time. The first section (2 Exploration and Findings) contains relatively detailed summaries of several dozen studies of the political resource curse because we believe that any effort to evaluate a large body of scholarship, diagnose sources of dissensus, and suggest new paths forward must begin from a firm foundation: we must all have a sturdy grasp on what the field has and has not accomplished if we are to move forward. As we show in this first section, the field has undergone a long-term evolution involving several transitions to better data, models, and research designs. Despite this evolution, much work remains.

The second section (3 Extracting Value) provides our diagnosis of the strengths and weaknesses of existing scholarship. We show that the resource curse is actually a large set of different and incompatible theoretical frameworks, that advances in the conceptualization and measurement of oil wealth have rendered much of the earliest literature invalid, and that very few current studies are based on a credible research design that resolves the problem of endogeneity. We also point to the under-appreciation of two other sources of biased inferences: measurement error and endogenous selection.

The third section (4 Refinement) discusses temporal and regional heterogeneity, demonstrating strikingly different patterns of oil and democratic transition in Latin America, the Middle East, Africa, and Southeast Asia. This section then makes the case that regime outcomes in oil states are mediated by political institutions and coalitions, neither of which are fully endogenous to oil: it is our

belief that focusing on institutions and coalitions can help overcome the pervasive and detrimental theoretical fragmentation that we document in the previous section.

The fourth section (5 The Resource Curse Reconsidered) presents some of our original research. We focus on the Middle East and North Africa (MENA), the global region that gave birth to theories of the rentier state and the resource curse and that provides the strongest evidence of oil-induced political dysfunction. The evidence we present, however, suggests that a set of historically contingent events in the oil-rich countries of MENA have been mistaken for a global relationship of resource wealth to political and economic dysfunction. Once we recognize that the political resource curse may be a historically contingent and highly restrictive outcome, we consider from a fresh perspective whether oil may be, at least in particular historical periods and under some conditions, the source of a pro-democratic resource blessing.

The authors of this text have literally grown up and grown older – and perhaps grown wiser – with this literature. One of us wrote his senior undergraduate thesis on the rentier state in 1984. The other, raised in Alaska and the son of its former State Geologist, began to study rentier states in 1991, and completed his doctoral dissertation on the political economy of oil in 2002. Over this long timespan, we have read the literature, argued about its findings, published our own findings, and thought long and hard about how to best move the research frontier forward. We are wise enough to know that no single study will resolve all existing debates. Our goals in this Element are as follows: to map the territory of the debates, suggest where some wrong turns might have been taken, and suggest fruitful paths forward, noting that the strength of "blessing" findings should warrant re-steering the field in a more open-minded direction. We hope the next generation of scholars will find it to be a valuable resource and guide.

2 Exploration and Findings

Research on the political resource curse stretches back five decades. Looking back over the long term, it appears that a mountain of evidence supports the theory of the political resource curse. In this section, we show that in reality, these studies constitute only the Piedmont, the elevated range of low-lying hills at the base of a mountain range. There is a large amount of evidence, but it takes the form of several piles of evidence, loosely related to one another but still not entirely consistent with each other, for each cluster of studies works with different concepts, measures, theoretical assumptions, statistical models, and research designs.

We believe it an invaluable exercise to explore these clusters of evidence in detail, to understand how much support each cluster lends to the theory but also to understand how the clusters disagree with one another. To facilitate this comparative assessment, we work through the material in three chronologically organized stages: the earliest work on the rentier state in the Middle East (roughly 1970–2000), the first decade (roughly 2000–10) of cross-national statistical research on the political resource curse (an outcome that is related to but distinct from the rentier state), and ending with the last decade of research which has featured a host of innovations – new concepts, new measures, new models, and new designs. We conclude this section by trespassing into the literature claiming a relationship between oil, corruption, and weak state institutions; we do so because corruption and weak state institutions may contribute to the erosion of democracy and because doing so brings this section full circle back to the broader concept of the rentier state with which the literature began.

We think there is an interesting trajectory of change that becomes quite evident in this section; in that sense, this section can be read as a study of how research communities develop and test ideas over time, correcting for past errors. Early cross-national statistical studies use convenient but problematic measures of oil wealth that will be shown later to lack construct validity; estimate statistical models on pooled cross-sectional, time-series data that covers a relatively brief time span and pay little attention to dynamic analysis; and are overly confident that the inclusion of covariates solves problems of endogeneity and hence pay little attention to research design. Over time, new and better measures are adopted, statistical techniques are refined, and research designs confront the problem of endogeneity more directly; but perhaps most importantly, hypotheses are refined, usually moving away from highly general and relatively vague claims to much more precise and narrow claims. Strikingly, though we do not include this material here, we have found the identical trajectory in statistical studies of the association of oil to economic growth or to civil war onset: early claims that are overly broad and lack conceptual and statistical validity are replaced by much more narrow claims based on more credible statistical models and measures. To some extent, then, the resource curse is a moving target.

The Rentier State in the Middle East, 1970–1990s

In the beginning, scholars studied the political and economic impact of oil wealth through the concept of the *rentier state*. As befitting a good narrative, the origins were relatively humble: Hossein Mahdavy's (1970) essay, "The Patterns and Problems of Economic Development in Rentier States: The Case

of Iran," was published in a relatively obscure edited volume on the economic history of the Middle East since the seventh century. Mahdavy defined rentier states based on the volume of external rents, or payments from foreign individuals, companies or governments, accruing directly to the state. Oil rents were of particular interest, because oil production was largely divorced from the rest of the economy and its opportunity costs were effectively zero.

Mahdavy's paper dealt primarily with the effects of massive oil rents on Iran's economic and industrial development, but he ended with the brief but trenchant observation that because rentier states could expand their size and finance their activities without taxing their citizens, their governments enjoyed independence from their citizens, including enhanced capacity to bribe pressure groups and coerce dissidents. This theme became the foundation for a number of conceptual and theoretical amendments over the next two decades, with a handful of works codifying Mahdavy's basic insight into the pithy formulation, "no representation without taxation," inverting the more traditional claim of "no taxation without representation" (Delacroix 1980; Anderson 1987; Beblawi 1987; and Luciani 1987). These works conceptualized rentier states as "distributive or allocation states," or states whose revenues were heavily dependent on oil rents and whose primary function was thus distribution, not extraction.

While these early works were largely conceptual, with limited empirics, detailed monographic studies of Middle Eastern rentier states were being published by the end of the 1980s (Anderson 1987; Crystal 1990; Gause 1994; Vandewalle 1998; and Lowi 2009). These new works did not formally test any hypotheses and did not compare rentier states to non-rentier states, but they illustrated, through detailed case studies, the utility of the concept of the rentier state to shed light on the political implications of oil wealth. Collectively, these works shared four features. First, these works interpreted oil as contributing to *political stability*, very broadly understood. Given their sample of cases, all of which had long histories of autocratic rule, they did not claim that oil *caused* autocracy. Second, these works highlighted the key role of *coalition formation and management* mediating between oil wealth and political stability. Theda Skocpol explained (1982, 269) the 1979 Iranian Revolution, for example, by noting that the given his reliance on oil wealth, the Shah "did not rule through, or in alliance with, any independent social class." A rich vein of evidence supporting the proposition that a distributive state without the necessary social coalition would enjoy only fragile stability runs through the monographic material.

Third, these works also highlighted critical antecedent conditions that appeared to distinguish the rentier state in the Middle East: the absence of

a peasantry or a class of large landlords on the Arabian Peninsula, social structures based on tribes, the "accidental" nature by which sovereign states were reluctantly formed by colonial powers, and the absence of substantial state structures prior to the mid-twentieth century. It was thus not self-evident that claims about the rentier state could be generalized beyond the contextual specificity of the Middle East. Finally, these works considered multiple phenomena to be constituted simultaneously: political stability, states with the capacity for distribution but not for extraction, and failed economic diversification as the basis for sustained growth. The next generation of scholarship would slice these phenomena into discrete dependent variables and research literatures.

The monographic literature did not, however, speak in one voice. Terry Lynn Karl (1997; see also Karl 1987) provocatively argued that Venezuela was a relatively stable democracy for many decades because of, rather than despite, its oil wealth. Venezuela's oil, she posited, "was the single most important factor in shaping the structural conditions for the breakdown of military rule, the subsequent creation of a reformist political space, and the maintenance of a *democracia pactada*" (Karl 1987, 94). Karl's analysis raises the possibility that the resource curse may be sensitive to local context, a point we develop in this Element's fourth section.

Does Oil Hinder Democracy? Studies from the First Decade

By the late 1990s, early explorations into the political and economic ramifications of oil had prepared the ground for a major expansion of research. What had started as a phenomenon generally restricted to the Middle East became – hypothetically at least – a global phenomenon, and the "antiquated" techniques of the case study would be replaced by the proliferation of data sets and quantitative models that treated economic growth, democracy, civil war, and weak state institutions as discrete consequences of oil wealth.

The study that established a discrete political resource curse as a major domain of research was Michael Ross (2001), "Does Oil Hinder Democracy?" Ross was curious whether claims from the rentier state literature could be generalized beyond the Middle East and hence adopted a pooled, cross-national statistical design, in which 113 countries were observed on an annual basis between 1971 and 1997. In a variety of models with different combinations of covariates, Ross found a statistically significant and negative coefficient on his measure of oil reliance.

Yet Ross also differed from the earlier literature in several subtle but crucial ways. First, Ross conceived of oil reliance as the share of fuel exports in GDP,

not as the share of oil rents in government revenue, shifting the key concept from rentier state to rentier economy. Second, this shift from rentier state broadly understood to a narrowly defined political resource curse implied a distinctive causal story, one that relied less on the "no taxation without representation" framework and more on an "anti-modernization" account: accepting the baseline claim that rising incomes render governments more democratic, Ross argued that oil-based income caused this democratizing effect "to shrink or disappear."

Third, while the earlier rentier state literature emphasized the political stability of autocracies, Ross used the Polity scale as his measure of democracy, thereby blurring the distinction between two distinct phenomena: oil can be linked to autocracy because it enhances the survival of preexisting autocracies, whose causal origins may be completely independent of oil, or oil can be linked to autocracy because higher levels of reliance on oil exports over time can make an existing regime less democratic. Ross thus interpreted his regression coefficients as implying that rising reliance on oil exports over time would cause a state to lose points on the Polity scale (i.e., to become less democratic). Ambiguity between oil as the cause of the *type* of regime and oil as the cause of the *survival* of a regime would continue to plague the literature for the next decade, as would debate over how to best isolate cross-sectional variance (differences in oil reliance and regime scores between countries at particular points in time) from within-country variance (differences in oil reliance and regime scores in one country over time). Subsequent research by Ross (especially 2012, 2014) would make important contributions to revising these early models, measures, and methods.

Reasoning that executive discretion over natural resource rents would give incumbents a tremendous advantage in the struggle to consolidate an autocratic regime, Nathan Jensen and Leonard Wantchekon (2004) test the association between their ordinal measure of natural resource reliance (both oils and minerals) and Polity scores for forty-six African countries between 1960 and 1995. In their main finding, relative to the least dependent countries (score = 1), Polity scores in the most highly dependent countries (score = 4) were about 1.59 points lower on the twenty-one-point Polity scale. Prior to 1990, however, almost all African countries had very low Polity scores: natural resource reliance only began to make an appreciable difference in Africa after 1990, when the end of the Cold War triggered a massive movement toward democratic reforms across the continent. Yet between 1994 and 1998, they find, natural resource dependence was clearly associated with a higher probability of democratic backsliding. By distinguishing a pre-1990 and a post-1990 effect of African oil, Jensen and Wantchekon raise the possibility that contextual features

across time and space may produce very specific patterns in the oil-democracy relationship.

To distinguish claims about autocratic *survival* from claims about *levels* of democracy or *types* of political regimes – claims that can sometimes be inadvertently conflated, as we have already noted – requires specific types of data and models. Reasoning that the causes of transitions to autocracy may be distinct from causes of authoritarian survival, Jay Ulfelder (2007) gathers event-history data about authoritarian regimes and tests whether oil wealth affects their rates of survival. Ulfelder codes autocratic regimes using Polity scores, and further codes each autocracy annually as undergoing a transition to democracy if a chief executive exercising effective and not just de jure authority and chosen by elections replaces one who was not. Resource dependence is measured by the share of GDP represented by resource depletion, ranging from 0 to 100 percent. The key finding relates to the hazard rate of democratic transition, the probability of the event occurring in the next time period, $t + 1$, conditional on surviving up to time t. The median autocracy with minimal resource dependence has only a minute probability of a democratic transition in any given year; resource dependence lowers this probability even further, a finding that is robust to multiple model specifications.

Finally, Silje Aslaksen (2010) published some of the earliest work using fixed-effects models that separate within-country variance from cross-sectional variance. Such models control for time-invariant, unit-level sources of heterogeneity and hence permit – at least in principle – the unambiguous interpretation of regression coefficients as unit-specific changes over time in levels of democracy subsequent to changes in levels of oil dependence. Aslaksen estimates that a 10 percent increase in the value of oil extraction as a share of GDP would be associated with a long-run decrease of approximately one-half of a point on the seven-point Freedom House index of Political Rights.

Some of the more fascinating – and surprising – evidence for the resource curse comes from the research of Ellis Goldberg, Erik Wibbels, and Eric Mvukiyehe (2008) into state-level politics in oil- and coal-rich states in the United States as far back as the late 1920s. Much scholarship would not expect to find any evidence of a resource curse in a wealthy democracy in which the oil and coal sectors are privately held and so rents do not accrue directly to the state. The authors counter (2008, 479), however, that nothing in the theory of the resource curse or its purported causal mechanisms imply that no effect should be observed in rich, industrial democracies. Accordingly, they adduce evidence that oil and coal rents are associated with lower rates of taxation and less competitive gubernatorial elections, measured by margins of victory and incumbent vote share.