WORLD POLITICS

A Quarterly Journal of International Relations

Volume 73, Number 2   April 2021

The editors invite submission of research articles and review articles bearing upon problems in international relations and comparative politics. Manuscripts should be double-spaced and submitted through the Web-based submission system, ScholarOne Manuscripts, at http://mc.manuscriptcentral.com/wp. Articles may be up to 12,500 words in length, including notes and references. Tables, figures, appendixes, and supplementary materials need not be included in the word count. Word count should be indicated. Manuscripts that exceed the limit *will not be considered.* Guidelines for articles and special issues, abstracts of forthcoming and recently published articles, and other information can be found on the *World Politics* Web page, at piirs.princeton.edu /wpj; digital versions of articles can be found at cambridge.org/wpo.

Authors can expect to receive decisions on their submissions within four months. Procedures for reviewing manuscripts are based on the anonymity of the author and the confidentiality of readers' and editors' reports; author anonymity is preserved, as well, during the editorial decision-making process. Self-references should therefore be removed. Referees are drawn from Princeton and other institutions; published articles have usually been reviewed by at least two non-Princeton reviewers and often, but not in all instances, one of the editors. Referees for the previous calendar year are acknowledged annually in issue 4 of the journal. In the case of an article deemed to be inappropriate for *World Politics,* the editors strive to notify the author within a month of submission that the article has been withdrawn from consideration.

World Politics does not consider material that has already been published (including in a foreign language), has been concurrently submitted elsewhere, or is already slated for publication even in a somewhat different form, such as a chapter of a book. This policy applies to both print and online formats. For these purposes, an online format that would preclude consideration by the journal refers to a refereed presentation and/or a copyrighted working paper. Examples of pre-published materials that may be considered are print working papers and on-line papers on an author's own homepage or Web site. Certain material already scheduled for publication, such as a chapter of a book, may be considered by *World Politics* if it is to appear no earlier than nine months after the likely date of publication in the journal. Dual submission and dual publication are not permitted while a piece is under consideration at *World Politics*. These restrictions apply to all copyrighted publications (including book chapters, journal articles, and/or working papers). Statements of fact and opinion appearing in the journal are the responsibility of the authors alone and do not imply endorsement of the editors or publisher. The journal does not publish communications to the editor or rejoinders to specific articles. Scholars who believe they have been challenged are encouraged to submit an article that will advance the scholarly debate.

WORLD POLITICS

CONTENTS

MOBILIZATION CAMPAIGNS AND RURAL DEVELOPMENT
The East Asian Model Reconsidered

By KRISTEN E. LOONEY

ABSTRACT

Most accounts of East Asian economic growth have focused on the role of developmental states in successful industrialization. This article expands and challenges that framework by showing that rural policy was different from industrial policy. A key finding is that for more than a century, East Asian states have relied on mass mobilization campaigns rather than on technocratic planning and market-conforming institutions to achieve rural development. Based on case studies of Japan, South Korea, Taiwan, and China, the author argues that three main factors explain the rise of campaign states: revolutionary traditions, rural populism, and policy learning. A brief assessment of outcomes illustrates the payoffs and costs of campaigns and the practical considerations that drive them. The author's analysis offers a new perspective on the East Asian model and disputes the widely held view that campaigns are tragic exercises in social control, demonstrating instead that they were central to the region's rural transformation.

INTRODUCTION

SINCE the early 1980s, numerous studies have linked East Asia's economic success to the emergence of developmental states. According to this theory, state modernizers in Japan, South Korea, and Taiwan determined that catch-up industrialization was essential to national survival and for decades treated that goal as their first priority. The state played an activist role in delivering growth through coordinated industrial policy, which was formulated by a highly professionalized (Weberian) bureaucracy and enforced by a strong (authoritarian) executive.[1] Despite many criticisms of the model, the developmental state has not only remained the dominant framework for understanding the region's political economy, including that of reform-era China (1978–present), but has also informed wider debates about the global financial crisis of the late 2000s and alternatives to neoliberalism.[2]

---

[1] The developmental state was first described by Johnson 1982. Other foundational studies include Amsden 1989; Chang 1994; Evans 1995; Haggard 1990; Wade 1990; Woo 1991.

[2] On the literature's evolution, see Haggard 2018.

Considering the sheer volume of scholarship on the developmental state, surprisingly little attention has been paid to the issues of rural policy and development. Most accounts acknowledge that in East Asia, poor farmers benefited from US-backed land reforms and that compared with other colonial powers, Japan invested heavily in rural infrastructure and institutions, which aided Korean and Taiwanese agricultural recovery after World War II. Beyond that, some economists have detailed agriculture's contribution to food production, export earnings, national income, and tax revenues, but mostly it is assumed that the region conformed to W. Arthur Lewis's dual-sector model in which agriculture's primary function is to supply surplus labor for industrialization.[3] In brief, a deep rural labor pool generates growth by ensuring that industrial wages remain low, and once the labor surplus is exhausted—a milestone known as the Lewis turning point—rising manufacturing costs will trigger economic rebalancing toward services and consumption.[4] The core assumptions are that agriculture is inherently backward, its value lies in how much it supports other sectors, and trickle-down industrialization is what eventually paves the way for rural development. Although some scholars have questioned these ideas, the developmental state literature has largely ignored agriculture or has focused on land reform as the region's most successful rural policy.[5]

There are two main problems with the conventional wisdom. First, it downplays the seriousness of urban bias and rural extraction in East Asia, which eroded the countryside's long-term development prospects and created significant rural-urban disparities. Second, it does not account for how East Asian governments dealt with agricultural adjustment (the declining performance and size of the rural sector) as industrialization deepened. A close examination of rural policy uncovers the limitations of previous research and a curious fact: The Northeast Asian developmental states all attempted to make up for decades of rural neglect and to realize development through mass mobilization campaigns. From "local improvement" in Meiji Japan to "rural revitalization" in contemporary China, campaigns have been integral to the region's rural transformation for over a century. Yet the use of campaigns—defined as policies demanding high levels of mobilization to effect dramatic change—deviates from standard portrayals of the developmental state as technocratic and market-conforming.[6] The timing

---

[3] For an overview, see Francks 1999.

[4] Lewis 1954.

[5] Kay 2002.

[6] The phrase "market-conforming" is closely associated with Johnson 1982, but it should be noted that Amsden 1989 and Wade 1990 rejected that characterization, arguing instead that developmental states deliberately got prices wrong and were market-distorting.

of campaigns is also puzzling. Many occurred at later stages of industrialization, after these regimes had supposedly entered a postrevolutionary phase of governance and the rural problem, as per the Lewis model, had begun to resolve itself.

This article examines the phenomenon of rural modernization campaigns and argues that they stem from three factors—revolutionary traditions, rural populism, and policy learning—each of which presents a challenge to common understandings of the East Asian model. First, in contrast with the classic literature on political development, these regimes did not progress linearly toward a more institutionalized, rational-legal form of rule. Rather, revolutionary traditions continued shaping political discourse, leadership, and practices well into the postwar period (and for China, into the post-Mao Zedong period). Wartime mobilization tactics were adapted to the tasks of state-building, governance, and development, and they were used across the political spectrum—among both left- and right-leaning authoritarian regimes in China, Taiwan, and Korea, and in Japan's more democratic context.[7] Second, these regimes embraced rural populism to address farmers' grievances and, in some cases, to build a political base in the countryside that could counterbalance urban oppositional forces. The populist turn in national politics was marked by pro-rural redistributive commitments, anti-elitist and nationalist ideologies, personalistic leaders, and activation of the peasant masses. Campaigns, as an expression of populism, thus complicate the notion of an East Asian bureaucratic authoritarian regime, which by definition is conservative, impersonal, and repressive.[8] Third, research on the origins of developmental states has stressed factors like colonial history, state-building patterns, and security dilemmas that fostered elite consensus around military readiness and economic development.[9] Without denying the causal role of those variables, this analysis shows that policy borrowing among countries has shaped what the East Asian model means just as much as the structural conditions that gave rise to it in the first place. Attention to policy learning furthermore reveals that campaigns are part of that model.

The broader implication is that East Asian rural development cannot be solely attributed to careful economic planning or to well-designed institutions. Nor is it simply a byproduct of marketization and industrialization. The literature on land reform fits both narratives, insofar as it emphasizes property rights, production incentives, and the transfer

[7] Following Huntington 1965, many scholars have disputed the link between modernization and institutionalization, yet the idea of Leninist regimes transitioning to a postrevolutionary phase has rarely been challenged; see Lowenthal 1970 and Lowenthal 1983.

[8] Cumings 1984; O'Donnell 1973.

[9] Doner, Ritchie, and Slater 2005; Kohli 2004; Vu 2007; Waldner 1999.

of surplus labor and capital to industry.[10] Interestingly, apart from the communist cases, existing Western scholarship hardly acknowledges that land reform was itself a campaign that produced varied results. It led to a short-term boost in agricultural productivity followed by slower growth rates, owing to the displacement of landlords who supported the local economy and a reduction in average farm size, which, combined with discriminatory pricing policies, made it nearly impossible to survive on farm income alone.[11] Indeed, East Asia's experience of rural development has been considerably more mixed than its stellar record of industrial development. The countryside's prospects improved when rural development was treated as a deliberate policy goal. But increased investment, subsidies, and other protection measures (all typical components of an agricultural adjustment program) are only part of the story. A key point of divergence from the developmental state literature and from development economics more generally is the means by which rural development was pursued—mobilization tactics that align more with Maoism or Leninism than with Japanese industrial policy.[12]

In addition to explaining the political logic behind campaigns, a secondary contribution of this article is to assess their efficacy and functional appeal. The campaigns that I examine aimed to achieve breakthroughs in rural economic and infrastructure development: to stimulate production, to raise incomes, and to improve access to quality roads, electricity, water, sanitation, and housing. The effect on production and incomes was ultimately marginal, but in all cases the rural-built environment was totally transformed, sometimes appropriately and voluntarily and sometimes not. Variation within the region and the contingent nature of success imply that campaigns are risky, working only when there is centralized bureaucratic control alongside decentralized rural participation. But even then, the state's vision may be plagued by vague and utopian goals that are unattainable no matter how or where implementation occurs. A review of cases from across the world, including China's Great Leap Forward in the late 1950s, which resulted in tens of millions of famine deaths, underscores the point that campaigns often fail. They are difficult to launch, still harder to sustain, and can easily succumb to excesses, such as economic dysfunction, political persecution, social alienation, and violence.

Nevertheless, East Asian officials and scholars frequently describe

---

[10] See, for example, Studwell 2013.
[11] Bramall 2004.
[12] For Bates 1981 and Lipton 1976, the solution to urban bias is less (not more) state intervention.

campaigns as vehicles for social cooperation and national moderniza-
tion, a view at odds with Western ideas about their illiberal and de-
structive tendencies. What explains this discrepancy and why, despite
so many problems, do campaigns persist? The promise of campaigns is
that by galvanizing and directing the country's resources toward certain
ends, they can deliver greater change to more places in a shorter time
frame than market forces alone or politics as usual would produce. By
demanding immediate and visible results, they can overcome bureau-
cratic inertia and demonstrate state power. Moreover, there is sufficient
historical evidence to suggest these payoffs are real. This is not to say
that skepticism of campaigns is unfounded, but skepticism should not
be grounds for dismissing this important mode of state-society inter-
action. Doing so reinforces a picture of East Asian development that is
not entirely accurate.

The following discussion of East Asia covers Japan, South Korea,
Taiwan, and China. The outcome of interest is mobilization campaigns
and the main unit of analysis is the region rather than individual coun-
tries. Country-level observations are treated as sub-cases or within-case
sources of evidence for the argument. North Korea, while no stranger
to campaigns, is excluded because of the limited availability of informa-
tion and because it is not regarded as a developmental state. The study
draws from archival and documentary materials collected during field-
work as well as from secondary sources published in English, Korean,
and Chinese. The article begins by addressing the limits of existing the-
ory. It next elaborates the central argument about the origins of rural
modernization campaigns and then turns to the payoffs and costs as-
sociated with them. The conclusion reflects on the concept of an East
Asian model.

DEVELOPMENTAL STATES IN RURAL SOCIETIES

Most accounts of East Asia's rural transformation privilege the role of
initial conditions, land reform, and developmental states in providing
an institutional foundation for growth. Although this perspective is not
incorrect, it is incomplete. Less well known is the history of state-spon-
sored campaigns to protect and develop the countryside, where farmers
had long suffered from exploitative policies often perpetuated by the
same state actors subsequently trying to undo them. In other words,
developmental states were not always good for development and the
way they operated did not always accord with institutional explanations
of change.

EXISTING THEORY

To be sure, the spread of similar institutions across the region led to extraordinary gains in agricultural production. In Japan, South Korea, and Taiwan, these included smallholder farming systems grounded in private property rights, comprehensive extension services maintained by technocratic bureaucracies and encompassing farmers' organizations, and extensive state control over rice and staple foods. These circumstances enabled the transfer of resources throughout the economy and facilitated broad-based growth in agriculture. In the 1950s, production quickly returned to and then surpassed prewar peak levels. Agriculture's average annual growth rate was 3 percent in Japan (1955–1970), 3.5 percent in South Korea (1954–1973), and 4.6 percent in Taiwan (1952–1967). In contrast, the global average for the period 1952–1971 was just 2.7 percent.[13]

In China, the move away from socialist agriculture after 1978 brought sweeping institutional reforms and impressive production gains. Collective farms and the grain-rationing system were dismantled, household farming and rural markets were introduced, new economic activities were encouraged (for example, township and village enterprises), and technocrats at all levels of the party-state were promoted.[14] In many ways, China started resembling its neighbors, not least in terms of economic performance. China's average annual growth rate for agriculture was 3.9 percent (1984–2006), again higher than the global average of 2.7 percent for that exact period. Even more noteworthy, during the early transition years (1978–1984), agriculture grew 10 percent annually, causing the share of the rural population living in absolute poverty to be cut in half—from 31 percent to 15 percent, or from 250 million people to 128 million people—a historic achievement in global poverty reduction.[15]

These outcomes notwithstanding, the effects of initial conditions, land reform, and developmental states on the rural sector were far from uniformly positive. Regarding initial conditions, on the one hand, an important legacy of the region's prewar (and pre-reform) rural modernization efforts was effective state penetration of the countryside. In Meiji Japan (1868–1912) as well as in colonial Taiwan (1895–1945) and in Korea (1910–1945), state authorities built advanced systems of transport, taxation, finance, education, and administration. In agriculture, they expanded on grassroots innovations and infrastructure, aiding

[13] Figures from Ban, Moon, and Perkins 1980, 16, 35–39; Francks 1999, 38; Yager 1988, 51–52.
[14] Kelliher 1992; Oi 1999.
[15] Bramall 2009, 340, 447. Global average from World Bank 2019.

the diffusion of technologies through cultivating landlords and state-controlled farmers' organizations. And although many policies were met with resistance, local communities were brought into ever closer contact with the state.[16] Similarly, in China under Mao (1949–1976), the rural population was quickly incorporated into the state administration. By the mid-1950s, over 90 percent of administrative villages had functioning Chinese Communist Party (CCP) committees, and a comparable share of farm households had entered into officially mandated cooperatives (precursors to collective farms or people's communes).[17] These institutions were at times strongly ideological and hostile toward science, but they also oversaw major advances in irrigation and green revolution technologies, such as improved seed varieties and chemical inputs.[18] Thus, in all these cases, later generations of modernizers benefited from preexisting resources for development.

On the other hand, throughout the Japanese empire, farmers suffered from exploitative tenancy relations and excessive state extraction, which offset the advantages of higher yields. In Korea, between 1915 and 1933, increased rice exports to Japan caused a 35 percent drop in per capita consumption and over the next decade, more than a million people affected by food shortages migrated to Japan or Manchuria.[19] The situation for Taiwan's farmers was less precarious due to lower population density, fewer absentee landlords, and greater off-farm employment, but Japan still controlled Taiwan's rural economy, especially rice and sugar, and tenant farmers paid at least half of their main crop in rent. The plight of tenant farmers in Japan was not much better. In the 1920s–1930s, faced with exorbitant rents and volatile prices, they formed thousands of tenant unions to curb the power of landlords and local political elites. But as most groups were suppressed or absorbed by the state, they were largely unsuccessful.[20] Postwar efforts to rehabilitate the region's rural institutions were consequently hindered by a certain level of public distrust. In South Korea, many people welcomed the collapse of the colonial agricultural bureaucracy, and it was not until the 1960s that a national farmers' organization was reestablished (with minimal fanfare).[21] Likewise, in Maoist China, many development achievements were wiped out by excessive extraction, population growth, and political extremism. With the advent of reform, Chinese farmers cele-

[16] Francks 1999; Ho 1978; Keidel 1981.
[17] Bernstein 1967.
[18] Bramall 2009; Schmalzer 2016.
[19] Burmeister 1988, 34; Keidel 1981, 20.
[20] Ladejinsky 1977.
[21] Ban, Moon, and Perkins 1980, 269–75.

brated the demise of institutions that symbolized past hardships. And attempts to adapt old institutions to the new economy, for example, by marketizing rural cooperatives, failed to gain traction.[22]

The contributions of land reform and decollectivization to rural development, while significant, have been exaggerated by scholars. In Japan, South Korea, and Taiwan, land reform turned most tenant farmers into land owners. It stabilized the countryside politically, improved the distribution of income, and compelled former landlords to transfer their wealth to industry, but the effect on agriculture was mixed. Whereas some studies assert that land reform raised agricultural output and incomes, others suggest that these outcomes were not due to land reform per se, and that its main effect was to reduce farm size (typically to less than one hectare). The reforms strengthened producer incentives in the short term, but extreme parcelization of land limited the prospects for long-term growth and forced many families to seek sources of non-farm income. Perhaps the biggest legacy of land reform was to lock in place the position of small farmers. It helped the region to avoid the problem of bimodal development (a wealth gap between large commercial farmers and small subsistence farmers) but also impeded mechanization and even modest economies of scale, causing growth rates to level off. The same was true in China, where decollectivization granted land-use rights to individual households, a change that promised greater freedom despite falling short of privatization. After an initial productivity burst, agricultural growth rates slowed considerably. The 3.9 percent average (1984–2006) mentioned above, though higher than most countries, was just one point above pre-reform levels (the annual average for 1963–1981 was 2.9 percent), indicating that smaller farms were not necessarily more efficient. Additionally, in the four countries, rural income growth was mostly tied to non-farm wages, and the reforms offered no solution to urban bias or to the rural-urban divide.[23]

There is a misconception that East Asia's developmental states did not exploit agriculture, but the notion that they achieved "growth with equity" through an "agriculture-oriented" development strategy has a long history in development economics.[24] Its influence is also evident in contemporary policy discussions in other regions. Ethiopia's government, for example, has promoted an agricultural development-led industrialization policy said to have been inspired by the East Asian (as

---

[22] Clegg 2006.

[23] For opposing views, see Bramall 2004; Griffin, Khan, and Ickowitz 2002. Figures from Bramall 2009, 340.

[24] Fei, Ranis, and Kuo 1979; Kuznets 1988.

well as South and Southeast Asian) cases, where pro-poor, pro-rural policies preceded industrialization.[25] To clarify, East Asia's income distribution was, for a time, remarkably egalitarian. The degree of urban bias—as measured by rural-urban differences in capital stocks, government expenditures, public services, and price and exchange rate distortions—was also less severe than in other parts of the world.[26] But still, urban bias was fundamental to East Asian agricultural policy.

The nominal and relative rates of assistance for agriculture (NRA and RRA, respectively), which compare domestic farm prices to international market prices and non-farm product prices, were negative in South Korea and Taiwan until about 1970 and in China until about 2000, implying that the main concern of these governments before then was to extract a rural surplus for industrialization.[27] Agriculture in postwar Japan was protected from the outset, although prior to the war it was similarly exploited. In the 1870s–1890s, agriculture provided up to 91 percent of total direct taxes, amounting to roughly 15 percent of agricultural GDP, yet it received less than 2 percent of total government subsidies. In the 1900s–1930s, industry shouldered a larger tax burden, albeit usually no more than 2 to 4 percent of industrial GDP, and it continued receiving the lion's share of government investment. The result was a serious deterioration in Japanese farmers' relative income position. In 1885, farm households earned 24 percent less than non-farm households; by 1930, the difference was 68 percent.[28]

This pattern repeats in postwar South Korea and Taiwan, where in addition to direct taxes, farmers paid hidden taxes through an official rice-fertilizer barter system managed by the national farmers' organization, which was essentially the only channel for buying or selling those items. These two countries furthermore imported massive amounts of cheap food from the US under its 1954 Agricultural Trade Development and Assistance Act (Public Law 480, or the Food for Peace program), as did Japan, which depressed domestic farm prices but appeased urban consumers, industrialists, and other groups interested in keeping food and wage costs down.[29] And in China, rural reform resulted in less state control over agriculture, but also in a continuation of the developmental squeeze. For most of the 1980s and 1990s, only about 3

[25] Routley 2014.

[26] Anderson 2009; Bezemer and Headey 2008.

[27] The NRA compares domestic and international farm product prices, while the RRA compares the NRAs of farm and nonfarm products; negative values suggest anti-agricultural bias; estimates available in Anderson 2009, 87–92, 373; Francks 1999, 42.

[28] Anderson 2009, 103–107.

[29] Hsiao 1981.

percent of central investment in China went toward agriculture, with the expectation that local officials would raise their own development funds. Many imposed onerous taxes on farmers and used the money for local industry rather than for public goods. The quality of rural roads, schools, and clinics declined, and the rural-urban gap widened. Heavy "peasant burdens" gradually gave rise to widespread social unrest.[30]

In summary, East Asia's developmental states successfully stimulated agricultural production, but like developing countries elsewhere, they sacrificed farmers' interests for industrialization. The institutions responsible for rural economic growth (local governments and farmers' organizations) were used to implement discriminatory price, trade, and investment policies, leading to uneven progress along different dimensions of development. Growth in agricultural production was generally not matched by similar gains in rural incomes, welfare, and infrastructure. It should be noted that strong rural-urban linkages in Japan, Taiwan, and coastal China provided farmers with more economic opportunities than was the case in South Korea or in the Chinese interior. But in all these places, rising rural-urban inequality fueled popular discontent, and the difficulties of agricultural adjustment were magnified by the pace of economic change. In South Korea, for example, it took less than three decades for agriculture's share of GDP to fall from 40 percent to 7 percent, compared with it taking a full century or longer to fall similarly in early industrializing countries.[31]

A purely technocratic response to the agricultural adjustment problem would have been to create a more favorable policy environment for agriculture akin to that of the industrial sector: a state-led but market-conforming development strategy based on targeted investments, skills promotion, and other sector-based coordination and support measures. Instead, technocratic and mobilizational politics were fused together. The governments of Japan, South Korea, Taiwan, and China did not just offer a plan-rational fix to the rural problem—they harnessed the power of the masses to transform the face of the countryside.

THE RISE OF CAMPAIGN STATES

State-led campaigns are policies that rely on intensive mobilization to achieve dramatic change. Although bureaucratic mobilization is central to all campaigns, the degree of popular mobilization may vary based

---

[30] Bernstein and Lü 2003; O'Brien and Li 2006. Central investment data available in *Zhongguo nongcun tongji nianjian* 2013, 77–78.
[31] OECD 2012, 33.

on campaign targets.[32] Early scholarship on Mao-era China classified campaigns as "ideological," "struggle," or "economic" in nature. Charles Cell accordingly defines the concept as "an organized mobilization of collective action aimed at transforming thought patterns, class/power relationships and/or economic institutions and productivity." Cell further notes that despite having fuzzy parameters, all campaigns involve "increased intensity of activity beyond what is expected in regular work and living routines."[33] Similarly, Gordon Bennett describes the Chinese term *yundong* (campaign) as "a government-sponsored effort to storm and eventually overwhelm strong but vulnerable barriers to the progress of socialism through intensive mass mobilization and active personal commitment."[34]

Writing from a broader historical perspective, Julia Strauss documents the many meanings of yundong—physical exercise, military operations, social movements, and efforts by the CCP and KMT (Kuomintang, or the Chinese Nationalist Party) to mobilize the populations they controlled. She explains that today, yundong, in the sense of campaign, refers to "the state's extraordinary mobilization of people and resources to implement a specific program to accomplish particular goals in a defined period of time." Moreover, campaigns constitute a "modality of policy implementation" that relies on the bureaucracy while also diverging from it in terms of "extraordinary mobilization, sharp focus, compressed time scale, and sidestepping rules."[35] Adding to this insight, campaigns are used to try to circumvent institutional constraints to change by reordering and augmenting existing power structures or by creating new structures that incorporate extra-institutional actors, such as local activists, organizations, and interest groups. Campaigns may therefore resemble social movements, but they differ because of the state's preeminent role in their execution. And contrary to official claims of popular spontaneity, there is usually a coercive element dictating mass participation. This conceptualization can be usefully applied to cases that predate Chinese communism and to cases outside China.[36]

Building on these works, I define rural modernization campaigns as policies aimed at overhauling agrarian society through intensive bureaucratic and popular mobilization. Common in authoritarian states,

---

[32] The campaigns under examination all entailed popular mobilization, but for campaigns generally, the popular element may be absent if the main target is the bureaucracy itself or if officials neglect or fail to induce participation.

[33] Cell 1977, 7–9.

[34] Bennett 1976, 18.

[35] Strauss 2020, 20–25, quotes on 22, 25.

[36] On Chinese late imperial campaigns, see Wong 2001.

though by no means limited to them, examples include land reforms, green revolutions, cooperative movements, collectivization schemes, and mass literacy, health, and sanitation drives. Rural modernization campaigns often exhibit both hard and soft qualities, combining infrastructural and economic goals with spiritual and moral ones. They can foster development by overcoming resource shortages and other barriers to change and promote legitimacy by extending ideological control and state patronage. Mobilization—defined here as an extraordinary, state-sponsored effort to activate and involve a population in the pursuit of certain goals—serves as the dominant mode of policy implementation. During campaigns, the state may employ such tactics as defining core tasks, dispatching work teams, empowering activists, identifying models, spreading propaganda, and exerting pressure to break up administrative routines and to overwhelm ingrained attitudes and behaviors.[37]

Japan provides the region's earliest example of a campaign approach to development—beginning with the Local Improvement Movement at the start of the twentieth century, continuing with the Rural Revitalization Campaign in the wake of the Great Depression, and resuming after World War II with the New Village and New Life Campaigns. In Korea, the colonial government mounted its own Rural Revitalization Campaign in the 1930s. Following the peninsula's division in 1945, the North orchestrated the *Chollima* (flying horse) Movement in the 1950s and 1960s, and the South carried out the *Saemaul* (new village, or community) Movement in the 1970s. In Republican-era China (1912–1949), the CCP and KMT went head-to-head in their efforts to establish cooperatives, each claiming leadership over a national Cooperative Movement in the 1920s and 1930s. The KMT also conducted the New Life Movement during that same period (this name was later borrowed by Japan in the 1950s). After the 1949 revolution, the CCP launched hundreds of campaigns on the mainland (affecting not just the countryside, but virtually all sectors and populations), while the KMT, exiled to Taiwan, directed the People's Livelihood Construction Campaign starting in 1955, followed by the Community Development Campaign beginning in 1965. And in all four countries, the entire bureaucracy was mobilized along with tenant farmers to execute a major land redistribu-

---

[37] Cell 1977, in an attempt to measure campaign mobilization, identified forty-three different informational, organizational, and mass participation indicators. This article's intention is not to measure or delimit a threshold of mobilization for policies to qualify as a campaign, but to highlight common campaign features and to explain the origins of a small subset of campaigns.

tion campaign.[38] As an aside, in Western scholarship the term "land reform" implies something less revolutionary and less communist, which may explain why it is usually applied to cases that received US support (Japan, South Korea, Taiwan), and why so much of the scholarship has ignored land reform's campaign features. As the region became more democratic, the line between campaigns and social movements blurred, and state-sponsored mobilization gave way to less interventionist policy measures, such as rural subsidies.[39] But rural modernization campaigns live on in China, as evidenced by the Building a New Socialist Countryside and Rural Revitalization initiatives, which began in 2005 and 2017, respectively.

From the farmers' perspective, these were campaign states rather than developmental states. Certainly, scholars have documented the region's campaigns in greater detail than is possible here. Two noteworthy studies are Sheldon Garon's work on Japan, which examines how campaigns were used from the early Meiji period onward as a social management tool, and Tyrene White's work on China, which shows how campaign tactics were adapted in the post-Mao era to enforce the one-child policy.[40] Yet campaigns have been conspicuously absent from research on East Asia's political economy. To the extent that mobilization is discussed, it is as a background factor that enabled the rise of developmental states and imbued them with revolutionary legitimacy. According to Chalmers Johnson, who pioneered the developmental state concept, Japan had "an economy mobilized for war that never demobilized during peacetime." He further states, "its leaders are somewhat akin to those of revolutionary mass movements."[41] Still, Johnson is writing generally about the mobilization of industry and does not give the impression that campaigns mattered for the Japanese miracle. While few political economists would make that argument, treating mobilization as ancillary to technocratic norms also misses the point that rural policy, and indeed the history of East Asian agrarian change, is inextricably linked to campaigns.

This is not to suggest that campaigns always succeeded or benefited farmers, but they were not invariably harmful either. Campaigns occur-

---

[38] Of course, participation took different forms with tenants playing a more direct, public, and sometimes violent role in the communist cases versus voting and standing for election to local land commissions in the noncommunist cases; see Ladejinsky 1977; Strauss 2017; Strauss 2020.

[39] The Saemaul Movement, for example, evolved into a civil society organization in the 1980s and remains active in Korea; see Korea Saemaul Undong Center n.d.

[40] Garon 1997; White 2006.

[41] Johnson 1999, 41, 53.

ring at the later stages of industrialization, when the transition from urban bias was underway, were especially influential. In each case—Japan's Rural Revitalization Campaign, South Korea's Saemaul Movement, Taiwan's Community Development Campaign, and China's Building a New Socialist Countryside—poor and isolated communities witnessed the arrival of electricity, paved roads, clean water, and upgraded housing, which led to vast improvements in productivity and quality of life.[42] Those outcomes would not have occurred so quickly, often in just two or three years, without the extraordinary pooling of state and community resources, including villagers' own labor and capital. The use of campaigns at a time of industrial expansion also indicates that East Asian governments did not put much trust in surplus labor absorption or even in guided market forces to solve rural underdevelopment. Campaigns are therefore significant, empirically and theoretically, for their far-reaching effects on the countryside and the style of politics they embody. Delving into the origins of campaigns reveals that at the height of authoritarian developmentalism, there were other, less technocratic forces at work—and those forces were just as consequential for East Asian political and economic development.

SOURCES OF MOBILIZATION IN EAST ASIAN POLITICS

Rural modernization campaigns are a product of revolutionary traditions, rural populism, and policy learning, and these influences contradict several key assumptions about East Asian developmental states. Addressing each factor in turn, the analysis below illuminates the importance of mobilizational politics in the region.

REVOLUTIONARY (AND COUNTERREVOLUTIONARY) TRADITIONS OF TECHNOCRACIES

The bulk of research on developmental states provides a snapshot of these East Asian regimes during high-growth periods, depicting them as (mostly) modern Weberian bureaucracies based on hierarchical authority, legal rules, functional specialization, and meritocratic recruitment.[43] Studies focused on the origins of developmental states take a longer view, although the effects of war and revolution on political development remain undertheorized or at least have not been fully explored beyond demonstrating that security threats contributed to

---

[42] In that order, see Smith 2001; Park 1998; Liu 1991; Looney 2015.
[43] Evans and Rauch 1999. For critiques, see Ang 2016; Ha and Kang 2011.

elite cohesion.[44] From imperialism and fascism in Japan, to Maoism in China, to right-wing authoritarianism in Taiwan and South Korea, revolutionary traditions greatly shaped these countries' development trajectories, even as technocratic rule was being consolidated.

My argument rests on the idea that revolutionary traditions can exist under different types of regimes, including the conservative or reformist regimes of postwar Japan, Taiwan, and South Korea, which did not come to power via social revolution. As such, the meaning of "revolutionary tradition" used here is an inherited political strategy, practice, or impulse that developed in the context of war and was bound up with the pursuit of fundamental societal change, even if the change envisioned was to revive traditional morality, to rejuvenate the nation, or to prevent bottom-up communist revolution. As Arif Dirlik observes of the KMT's 1930s' New Life Movement in mainland China, a campaign that emphasized hygienic and behavioral reforms: "The contradictory claims of revolutionism and conservatism did not merely distort or cancel out one another; they were synthesized into an ideology of counterrevolution."[45] Mass mobilization was intended to transform society so that it would fully support rather than undermine existing bureaucratic structures. Despite this difference with the region's communist revolutionary movements, the mobilization methods embraced by both left- and right-wing regimes were surprisingly similar and continued well into the postwar period.[46]

Of course, nowhere is the durability of revolutionary traditions more obvious than in China. Despite Deng Xiaoping's repudiation of campaigns at the start of reform, they persist and continue to be used for policy enforcement, regional development, ideological work, crisis management, and anticorruption. As White and others note, the continuation of campaigns partly stems from path dependence and the Maoist imprint on political culture. It also reflects the regime's confidence in mobilization as a check against bureaucratic ossification and as a practical method for effecting change. The communists' penchant for mass mobilization developed during the Jiangxi Soviet and Yan'an periods, when they leaned on the peasantry for survival and guerrilla warfare. After 1949, wartime mobilization tactics were deployed for purposes of state-building, governance, and development. Maoist revolutionary politics oscillated between storming (subjecting the party to

[44] One exception is Stubbs 2005.

[45] Dirlik 1975, 945–46.

[46] For more on the blurry distinction between revolution and counterrevolution and the KMT's and CCP's mutual borrowing of mobilization practices, see Perry 2006.

rectification by the masses) and engineering (subjecting the masses to party-controlled mobilization); and while the former impulse faded under reform, the latter survived.[47]

In 2005, Hu Jintao's government unveiled Building a New Socialist Countryside, an ambitious effort to stimulate production, raise living standards, reform rural culture, beautify villages, and improve governance. The policy was never officially called a campaign, but it was understood and implemented that way. Plans that stressed speed and hard targets were drawn up, special committees were formed to realign bureaucratic interests, and cadre work teams were sent to the villages. Propaganda appeared on billboards, buildings, and fences. The phrase "building a new socialist countryside" could be traced to the revolutionary days of the 1950s. This slogan's resurrection, which first gained prominence during the Great Leap Forward, signaled the policy's importance. But a wariness of campaigns also surfaced in central leaders' statements, including warnings against "overnight success," wasteful exercises in "image engineering," and actions "violating the wishes of the masses."[48] This tension between mobilization's appeal and its potential excesses is important to understanding why, in the years since Mao, the CCP has repeatedly conducted campaigns without announcing them as such. A recent example is Xi Jinping's Rural Revitalization Strategy, introduced in 2017, which outlined goals similar to the New Socialist Countryside, but with added emphasis on poverty eradication. An outgrowth of China's war on poverty, the policy highlights how campaign-style governance has only become more pronounced under Xi.[49]

The institutionalization that occurs as a revolutionary regime matures does not mean that mobilization inevitably disappears. Examining the development of European Leninist parties, Kenneth Jowitt finds that even with more systematic social inclusion in politics, mobilization still served to tighten control over society and to address political challenges.[50] Although Jowitt would likely classify the KMT as reformist rather than revolutionary, this insight aptly describes the party's political development in Taiwan. In the early 1950s, Chiang Kai-shek's defeated KMT regime reorganized itself based on the CCP's example of elevating the party above the state and military, becoming more Leninist than it was on the mainland.[51] It also mimicked the communists during

[47] White 1990; White 2006. See also Perry 2011.
[48] Wen Jiabao, quoted in Zhang 2007, 60–61. On the policy's campaign features, see Looney 2015; Perry 2011.
[49] Smith 2018.
[50] Jowitt 1975.
[51] Dickson 1993.

land reform, employing similar propaganda and mobilization methods.[52] That pattern continued during the People's Livelihood Construction Campaign, launched in 1955, which like the New Life Movement emphasized behavioral change, but also focused on mobilizing village labor for infrastructure development.[53] Over the next two decades, as the regime became more institutionalized and inclusive, for example, by expanding elections and recruiting more Taiwanese into politics, it continued using mobilization to accomplish its goals. The Community Development Campaign of the late 1960s and 1970s is a case in point. It aimed to "eliminate dirt, disorder, and poverty, increase production and welfare, and promote a new morality" in the countryside. Ostensibly inspired by international notions of community-based, participatory development, the policy in practice relied on compulsory labor and other tactics that were unmistakably Leninist, such as work teams and the like. The regime's fear of rural unrest and communist infiltration, which earlier provided motivation for land reform, once again triggered a big push in response to agriculture's decline.[54] The campaign thus fit with Chiang's brand of right-wing authoritarianism, applying certain communist techniques to the decades-long fight against communism.

In South Korea, Park Chung-hee's regime was both technocratic and mobilizational, with an increasingly professionalized corps of officials setting aside their normal administrative routines to connect with the masses and become conduits for campaigns. The Saemaul Movement of the 1970s was an all-out drive to overcome rural backwardness and to achieve collective prosperity. Cement and steel rods for village infrastructure projects were distributed, communal kitchens were constructed to free up women's labor, and millions of people attended military-style trainings. Every day at 5:45 a.m., villagers woke to loudspeakers playing the "Song of Saemaul." Supposedly penned by Park, the lyrics urged villagers to "work while fighting" and "fight while working," with the implied enemy being communist North Korea. Photographs from the time show people assembled together, listening to lectures, planting rice seedlings, or building bridges underneath the Saemaul flag. In Park's words, the campaign represented a "great leap forward"—without the tragic ending.[55] It is unsurprising that scholars have drawn comparisons with Maoist China and the interwar fascist regimes of Europe

---

[52] On the similarities and differences (in ethos, implementing organization, etc.) between China and Taiwan, see Strauss 2017; Strauss 2020.

[53] Central Committee of the Kuomintang 1961.

[54] Central Committee of the Kuomintang 1972; Hung 1978.

[55] Park 1979, 152, 167, 204, 222.

and Japan.[56] Like Mao, Park came from a rural, military background and believed in mass mobilization. Born in the colonial period, he was obsessed with turning South Korea into a second Japan. During World War II, he attended the Manchurian Xinjing Officers School and the Japanese Military Academy, after which he enlisted in the Japanese Imperial Army. These experiences informed his perspective on modernization.[57] Building richer and stronger Saemaul villages—and later cities, factories, and schools—was essential to building a richer, stronger nation. The parallels to Japan's "rich nation, strong army" ideal and more concretely to the objectives and methods of Korea's colonial-era Rural Revitalization Campaign are striking.[58] They reflect Japan's dual status as a model and a rival and the enduring relevance of South Korea's revolutionary traditions forged in opposition to colonialism and communism (even if Park was not exactly an anti-colonialist).

In Japan, rural modernization campaigns were directly linked to the rise of imperialism and fascism. The Local Improvement Movement, launched in 1900, sought to strengthen the village political economy and to foster emperor-centered nationalism so that the state could finance industrialization and imperial expansion. Although the campaign fell short of its goals (only about 10 percent of villages participated), it set a precedent for greater state intervention in the countryside.[59] The Rural Revitalization Campaign, which began in 1932, aimed to defuse social tensions after the Great Depression devastated Japan's silk and rice markets, but in 1937, it was repurposed to support the military occupation of China. Despite this controversial history, the campaign marked the beginning of agricultural protection in Japan and at least during the early years, it delivered many benefits to impoverished farmers, including debt relief, increased savings, better infrastructure, and new technologies. More importantly, by combining state and private efforts to realize a number of "lifestyle improvement goals," it left a legacy of mobilization that lasted for decades following World War II.[60]

This brief overview of the traditions underpinning campaigns, like the developmental state literature, confirms that regime type matters for how development happens. But shifting the focus from industrial policy to rural policy shows a different side to these regimes—one that is more revolutionary and also more populist.

---

[56] Han 2004; Kohli 2004.
[57] Moon and Jun 2011.
[58] Shin and Han 1999.
[59] Pyle 1973.
[60] Garon 1997; Partner 2001; Smith 2001.

RURAL POPULISM UNDER AUTHORITARIANISM

Populism is not something typically associated with East Asia. Quite the opposite, the success of these states compared with failed developmental states is said to rest on the exclusion of the popular sector from national politics, which allowed conservative coalitions of the state and big business to make economic decisions without having to consider distributional consequences.[61] Far from being populist, East Asia's development states were "bureaucratic authoritarian industrializing regimes" and "social welfare laggards."[62] To protect central power holders' authority, these developmental states relied on repression and corporatist co-optation, building extensive internal security organizations as well as monopolistic interest- and patronage-based associations for farmers, workers, and the private sector.[63] Yet regardless of how one defines populism—as a movement, an ideology, a political style, or a set of policies—East Asia's campaigns contained many populist elements.[64]

Throughout the region, political leaders embraced a personalistic ruling style and proclaimed anti-elitist and nationalist ideologies. They promised greater redistribution to the poor and called on the masses to remake their communities. But unlike more typical cases of populism, East Asia's campaigns were initiated by incumbents rather than political outsiders, and they were rural in orientation. Borrowing from Samuel Huntington, they were about containing the green uprising. According to Huntington, in response to growing rural-urban disparities and political uncertainty, incumbents try to mobilize the peasantry within the existing system, thereby undercutting their political opponents and the potential for rural revolution.[65]

The South Korean case clearly illustrates this point. There, fears of political turmoil triggered the adoption of pro-rural policies. Dissatisfied with the government's squeeze on agriculture, many farmers voted against Park at the ballot box and with their feet, supporting opposition candidates in the 1967 and 1971 presidential elections, and moving to cities where anti-Park sentiment was rising. Farmers' grievances centered around low grain prices and surging production costs: fertilizer prices jumped 80 percent in 1964 alone, followed by a 44 percent increase in 1965. The biggest fallout occurred in the Cholla region, the

[61] Naseemullah and Arnold 2015; Pempel 1982; Waldner 1999.

[62] Cumings 1984; Wong 2004, 351.

[63] On the region's coercive institutions, see Greitens 2016. Related to the idea of corporatism is embedded autonomy, or a state's ability to be both internally coherent and closely linked to society; see Evans 1995.

[64] On different conceptualizations, see Weyland 2001. On the differences between populism and bureaucratic authoritarianism, see O'Donnell 1973.

[65] Huntington 1968, 72–78, 433–60.

so-called rice bowl of Korea, where Park lost twenty-three of thirty-four counties in 1967, a shocking defeat considering that four years earlier he had won them all. The countryside was supposed to be his stronghold, and although he ultimately carried these elections, losing Cholla was a warning that worsening rural conditions could endanger his rule.[66] After minor policy adjustments failed to shut out the opposition, Park steered the country in a more authoritarian and populist direction—removing all electoral and legal constraints on his authority with the Yushin Reforms (October Restoration, 1972–1979) and mobilizing the entire rural population through Saemaul.

In the context of the Yushin Reforms, which otherwise depended on the state's coercive apparatus, Saemaul employed a mix of propaganda and material rewards to create a rural populist base. Park's son-of-the-soil image was reinforced by an endless stream of rural inspection tours and televised interactions with local Saemaul leaders. During the 1970s, approximately 200,000 activists in total were empowered to spearhead the implementation of rural development projects in their villages. In line with the campaign's egalitarian ethos, these people were younger and better educated than traditional village elites, and roughly half were women. They wore caps emblazoned with the campaign's logo, a green-and-yellow three-leaf bud representing the Saemaul spirit of diligence, self-help, and cooperation. They were also publicly exalted for leading their villages and the nation to prosperity.[67] In stark contrast, Park lamented, "are the city people—it makes one doubt whether they are really citizens of the Republic of Korea—who by their weird and unruly behavior pour cold water on the Saemaul movement and dampen the enthusiasm of the rural people."[68] In addition to such populist rhetoric, selective assistance helped to generate mass support for the campaign. Villages demonstrating a stronger commitment, as measured by residents' labor, cash, and material contributions, obtained more government resources. The dramatic transformation of villages receiving aid induced neighboring villages to participate, and a competitive emulation dynamic took hold.[69]

Through campaigns, East Asia's farmers were activated rather than incorporated into national politics, which meant that instead of being mobilized electorally or granted more formal representation in the central government, they were recruited to carry out the regime's vision of

[66] Lee 2011, 353–68.
[67] Han 2004; Han 2010.
[68] Park 1979, 168.
[69] Brandt and Cheong 1979.

modernization.[70] Popular participation was widespread but controlled, localized, and directed toward certain ends. Strategically, rural populism served to legitimize authoritarian rule (and to counteract urban opposition) while maintaining a narrow governing coalition. The apparent success of that strategy in Korea prompted an expansion of Saemaul from Korea's villages to the whole country, although it became less effective over time. Eventually, unresolved economic problems led to antigovernment protests in the countryside and the ideology of *minjung* (mass democracy) spread across the cities, suggesting an untenable contradiction between state populism on the one hand and popular sector exclusion on the other.[71] Nevertheless, those forces did not become powerful until after Park's death, and still today, Saemaul is commonly remembered as one of the country's greatest achievements.[72]

In the other cases as well, populist mobilization was used to reassert state control and to consolidate political power during periods of uncertain change. As industrialization deepened, rising inequality and structural economic shifts created strong currents of rural discontent, which coincided with major leadership transitions at the top. In Japan, the tenant union and agrarianist movements peaked just as the military was trying to establish order following the 1932 assassination of Prime Minister Inukai Tsuyoshi.[73] In Taiwan, the election of sympathetic legislators after 1969 altered public discourse about the countryside just as Chiang Ching-kuo was preparing to take over from his father.[74] And in China, waves of spontaneous mass protests in the late 1990s and early 2000s raised the specter of rural crisis at the same time Hu was confirmed as Jiang Zemin's successor.[75]

Chiang and Hu, in particular, were more populist than their predecessors in terms of political styles, economic policies, and attitudes toward participation. Chiang Ching-kuo supported the KMT's Taiwanization, developing ties with farmers and encouraging them to join the party. He also appointed Taiwanese politicians to the cabinet, the provincial governorship, and the vice presidency.[76] One of Chiang's first

[70] Another distinction worth noting is between corporatist organizations (farmers' associations), which were arguably made stronger by campaigns, and the political incorporation of a lower- and middle-class constituency that could challenge elite interests at the national level, which did not occur in these cases but was central to classic Latin American populism; again, see O'Donnell 1973.

[71] Han 2004. As this source also notes, the mobilization of urban workers through Factory Saemaul aimed to prevent labor unrest through small-group ideological indoctrination, a strategy that ultimately backfired.

[72] Hong 2008.

[73] Smith 2001.

[74] Liao, Huang, and Hsiao 1986.

[75] Bernstein and Lü 2003.

[76] Gold 1986, 91–92, 113–14. To clarify, Taiwanization was already underway, but the highest

policy actions was to abolish the exploitative rice-fertilizer barter system and declare a new mass "campaign to accelerate rural construction" (the Community Development Campaign).[77] Hu was similarly concerned about the countryside because of his background working in three of China's poorest provinces, Gansu, Guizhou, and Tibet. Premier Wen Jiabao had worked in Gansu, too, and skillfully cultivated a man-of-the-people public image. Under their leadership, agricultural taxes were terminated and numerous pro-rural initiatives were enacted as part of the New Socialist Countryside. Within a few years, the Hu-Wen populist coalition had distinguished itself from the elitist coalition associated with Jiang.[78] Like Park, these leaders also appealed to farmers directly, imploring them to help execute state policies.

There were other reasons for the reversal of urban bias, such as concerns about food security and rural-to-urban migration, not to mention the declining cost of agricultural subsidies as the farm sector contracted. But while those factors may explain the timing or feasibility of intervention, they cannot explain the form. Campaigns were a reaction to domestic instability and proof that the region's strong authoritarian states were not immune to societal pressures. They were a manifestation of populism, which, as others have shown, can arise in competitive electoral systems and in hegemonic party-states.[79] For East Asia, populism represents a divergence from governance practices that have long defined the region: exclusion, repression, and corporatist social management. Populist mobilization coexisted with those practices but was intended to deal more effectively with the agrarian problem, both politically and in a more fundamental sense. As explained below, the idea that campaigns can solve rural backwardness derives from certain stereotypes about the countryside and, relatedly, from these countries' efforts to learn from one another.

POLICY LEARNING AND THE AGRARIAN PROBLEM

The Japanese developmental state evolved from a series of economic and war-related crises that caused the country's leaders and the public to equate growth with national security. As Johnson notes, its success was "the result of a tortuous learning and adaptation process" that un-

---

levels of government were unaffected before the 1970s. Chiang also tolerated the election of nonparty Taiwanese, exchanging patronage for loyalty and eventually allowing deeper democratic reforms; see Taylor 2000.
[77] Central Committee of the Kuomintang 1972.
[78] Li 2005.
[79] Dix 1985; Heydemann 1999.

folded over the course of fifty years (1920s–1970s).[80] Although subsequent studies of the East Asian newly industrialized countries (NICs, including Singapore and Hong Kong) confirmed the developmental state model, they also underscored its limited transferability. South Korea and Taiwan were able to emulate Japan because of their shared colonial history, common state-building patterns, and similar geopolitical concerns. They were furthermore "receptacles for declining Japanese industries," part of an integrated product cycle that resembled "flying geese," with technology moving from leader to follower countries.[81]

Despite many obvious differences between China and its neighbors, early reformers in China looked to Japan and the East Asian NICs for clues about how to introduce markets and build institutions while maintaining political control. Policy learning took off because it aligned with the CCP's traditional emphasis on experimentation and satisfied a hunger for practical knowledge as China emerged from a radical and isolationist period.[82] This is not to say that learning always succeeded. For instance, the central government's plan to follow Japan by supporting domestic industry and limiting foreign investment was thwarted by local governments and diaspora capital, which pushed for greater liberalization.[83] Yet China converged with the developmental state model in other ways: commitment to economic growth, merit-based recruitment of officials, and corporatist relations with industry. Learning was also evident in the rural sector, where in the 1990s, the Ministry of Agriculture launched a pilot program to "learn from the Japanese Central Union of Agricultural Cooperatives" and, more recently, promoted Japan's "one village, one product" method of agricultural specialization.[84] Again, these efforts have met with mixed success, but they show that these countries remain an important "cultural reference group" for China.[85]

The links between China's New Socialist Countryside and South Korea's Saemaul Movement are particularly notable. In 2005, Zheng Xinli, then-deputy director of the CCP's Central Policy Research Office, led a group of high-level officials to Korea. They planned to study the Saemaul Movement "in order to draft the 11th five-year plan, resolve the three rural issues [peasants, villages, agriculture], and promote rural economic and social development." Zheng later submitted

[80] Johnson 1982, 306.
[81] Cumings 1984, 3.
[82] Heilmann 2008.
[83] Ye 2009. For more on China studying Japan, see Leutert 2020.
[84] Ma 2008, 40–41; Smith 2019.
[85] Simmons and Elkins 2004.

two reports to the State Council recommending ways that China might emulate Korea, for example, by focusing on village improvement and using selective rewards to spur mass participation.[86] The next year, China's government announced that it was sending thirty thousand civil servants to Korea for Saemaul training.[87] In fact, during the New Socialist Countryside, millions of officials were required to learn about Saemaul and there was a burst of Chinese publications comparing the two campaigns.[88] In a 2010 interview, Li Shuishan, an ethnic Korean born in China who worked in the Ministry of Agriculture, explained that Chinese observers intuitively understood Korea's experience—that of a powerful state trying to modernize a densely populated countryside that had underdeveloped markets and scarce resources. The campaign approach, he added, was something the Chinese practically invented.[89]

The effect of policy learning on East Asia's political economy has not received much attention due in part to the difficulty of separating learning from structural explanations of change. Was learning the reason that Japan, South Korea, and Taiwan enacted similar land reforms, or was US pressure and the communist threat more pivotal? Stated differently, did policy convergence stem from the circulation of ideas or was it the result of parallel developments rooted in structural circumstances? There is compelling evidence that China's rural policy was modeled after Korea, but how can this influence be disentangled from path-dependent theories of campaign-style governance?

The argument is not that learning itself causes policy change, but rather that learning shapes leaders' understanding of appropriate policy choices in the event of change. Whether those choices are adopted depends on a range of ideational and structural factors, but adoption is more likely if the policies succeeded in a familiar setting. Consistent with studies of policy diffusion in other contexts, East Asian policymakers relied on cognitive shortcuts, such as information availability, to narrow the field of possible responses to various challenges, and they copied one another because of perceived cultural affinities and policy successes.[90] Even a cursory examination of each country's rural policy literature reveals frequent exchanges within the region and a self-conscious identification with a broader East Asian model—an acknowledgement of similar experiences with land reform, rural cooperatives,

[86] Fang and Liu 2006, 51.
[87] Do 2009.
[88] Using China National Knowledge Infrastructure (CNKI), a keyword search for *Hanguo xincun yundong* yields thousands of articles.
[89] Author interview with Li Shuishan, Beijing, April 2010. See also Li 2006.
[90] Simmons and Elkins 2004; Weyland 2005.

the green revolution, and agricultural adjustment, and a deliberate at-
tempt to learn through comparison.[91] East Asian sources, moreover,
describe campaigns in the same way, as policies or movements (*undo*
[Japanese], *undong* [Korean], *yundong* [Chinese]) capable of inspiring
state-society cooperation and altering rural material and cultural con-
ditions.

A key lesson exchanged about campaigns is that they can reshape cul-
ture to be more conducive to development, which differentiates them
from ordinary policies. State modernizers in prewar Japan, drawing
from Buddhism and Confucianism, believed that order and prosperity
rested on instilling proper values in people, and they invested enormous
resources in moral suasion. During the Local Improvement Movement
(1900–1918), peasants were exhorted to be hard-working and frugal, a
message that was conveyed through customs reform groups and other
state-backed associations. The campaign popularized the philosophy
of Ninomiya Sontoku, a nineteenth-century agricultural reformer who
stressed the importance of self-help in overcoming poverty. Officials
claimed that providing rural financial support without moral education
would breed indolence and dependency, and that idea carried over into
later campaigns.[92]

In Korea, the agrarian problem was also defined in cultural terms,
first during the Rural Revitalization Campaign and then during the
Saemaul Movement. As Park remarked, constructing a bridge "must be
viewed from the psychological angle," and "modernization starts with
the spiritual modernization of individual farmers."[93] The Saemaul spirit
was said to reflect Sontoku's teachings and to serve as a kind of Protes-
tant ethic (an analogy found in Japan, too).[94] Taiwan's officials likewise
tried to "promote a new morality" during the Community Develop-
ment Campaign. They organized life-basics courses on civilized be-
havior, for example, standing in line, wearing clean clothes, and eating
at a table, and encouraged people to join activities that would promote
healthy living and a collective ethos. These ideas can be traced to earlier
KMT campaigns as well as to the Japanese and South Korean examples.[95]
Similarly, China's New Socialist Countryside aimed to cultivate a new
type of farmer who was professional and civilized both in habits and ap-
pearance. The campaign's fixation on culture was less pronounced than

[91] See, for example, Wen 2011, 118–60. On emulation, see Moore 1993.
[92] Garon 1997, 7, 9–10, 31–32, 45–47.
[93] Park 1979, 75.
[94] Han 2004, 74.
[95] Lee 1979; Liu 1991.

in the other cases, but throughout the region, spiritual and economic mobilization went hand in hand.

Political elites wanted to change villagers' behavior and restore their sense of community, which capitalism had supposedly eroded. These goals were predicated on stereotypical views of the peasantry as a backward social class and the village as a place where positive communal traditions, if revived, could conquer laziness, individualism, and other cultural tendencies inhibiting development. To clarify, the forward march toward advanced (state-guided) capitalism was not in question, but how agrarian society fit into that process was ambiguous. The countryside was both a victim of modernization and the missing link. While market-based solutions were considered insufficient, if not antithetical, to rural development in a smallholder society, state-led campaigns represented a promising alternative. Policymakers believed they could accelerate change and transform rural culture by transmitting the right combination of traditional and modern values. This view from the top, however flawed, originated in Japan before spreading to the rest of East Asia, and it helps to explain the appeal of campaigns in the region. Policymakers were also able to infer cultural change was happening, even if it was not, because of another attractive feature of campaigns—for the most part, they worked.

Payoffs and Costs: East Asia's Campaigns in Comparative Perspective

Campaigns played a central role in East Asian rural development, overhauling economic and social relations and the appearance and organization of the countryside. Still, as those familiar with Soviet or Chinese collectivization can attest, campaigns sometimes fail miserably. Even among the cases examined here, which were quite successful, campaigns produced a mixed range of outcomes.

The effects on production and incomes were significant but short-lived and no match for larger economic forces. In South Korea, the compulsory rollout of a high-yield variety of rice called *Tong'il* (unification) tarnished Saemaul's image, as a few years of bumper harvests were followed by large-scale crop failures and financial losses. The program's collapse in the late 1970s pushed many farmers back into tenancy or into low-wage urban factory work.[96] In China during the New Socialist Countryside, total grain output rose by about 30 percent (2002–2012),

---

[96] Burmeister 1988.

but so did the country's reliance on grain imports, which grew by 400 percent.[97] Rural income growth was also not strong enough to narrow the urban-to-rural income ratio, which registered 3.1:1 or higher during the entire Hu era.[98] It is not clear how much of China's growth in production and incomes was related to subsidies versus more aggressive campaign measures, such as the drive to scale up agriculture or to launch new cooperative and agribusiness ventures. In Japan and Taiwan, farmers' income positions improved, but also for reasons unrelated to campaigns. In addition to the contributions of subsidies, decentralized patterns of industrialization enabled rural households to farm part-time and to work part-time in local factories.[99]

Yet in all these cases, the impact of campaigns on the village environment was profound. This aspect of rural development—defined as the quality of village infrastructure, sanitation, and housing—is extremely important but is generally underappreciated in scholarly accounts. Of course, the state's involvement in this area can be problematic. Officials may benefit politically from the creation of idyllic Potemkin villages. It is also possible that relocating people into planned communities may cause economic disaster, as occurred during Tanzania's 1970s Ujamaa village campaign. But in East Asia, campaigns were directly responsible for the expansion of rural electricity, roads, bridges, water, sewage, irrigation, and waste disposal systems, in addition to upgraded housing.

Returning to the South Korean example, in 1970, about 80 percent of rural homes had thatched roofs made of rice straw, which needed to be replaced annually, and only 20 percent had electricity. Less than 10 percent of roads were paved and around 40 percent of villages were inaccessible by car. In just these three areas—roofs, electricity, and roads—the Saemaul Movement produced extraordinary changes. By 1978, nearly 100 percent of homes had tiled roofs, 98 percent had electricity, and 100 percent of villages could be reached by road. Electricity consumption and ownership of electrical appliances increased at a staggering rate, enabling all kinds of improvements in working and living conditions. With the expansion of bridges and roads, the traditional wooden A-frame carrier, used to transport heavy loads along narrow pathways, virtually disappeared.[100] Once an iconic image of rural poverty, the A-frame was replaced with hand-drawn carts, bicycles, and motorized vehicles.

[97] *Zhongguo liangshi nianjian* 2013, 553, 589.
[98] *Zhongguo nongye fazhan baogao* 2013, 179.
[99] Francks 1999.
[100] Park 1979, 3–6; Park 1998, 1; Rossmiller 1972, 28; van Gevelt 2014, 182–85.

The other cases turned out similarly, with campaigns affecting all types of infrastructure at the community and household levels, although the quality of implementation varied. On the positive side, Taiwan achieved policy coherence among different levels of government by using campaign-coordination committees, competitive evaluations of local officials, and village-based community development councils. The country's small size and centralized political system facilitated bureaucratic oversight, while the presence of strong rural organizations and norms of participation ensured that villagers' interests would be considered, more so than in Korea.[101] Unfortunately, the same was not true for China, where local officials were subject to relatively weak checks from above and below. Implementation was selective but also carried to extremes. Priority was given to goals favoring the urban-industrial sector, such as housing construction and land consolidation, and the policy evolved into a (local) government-led crusade to demolish and reconstruct villages.[102]

China's experience with village renovation, like Korea's with Tong'il, shows how easily campaigns can spiral out of control. First, campaigns often have too many goals and unrealistic expectations about the technical skills and time required to achieve them. Gaps can therefore arise between central aims and local practices. Second, the highly politicized environment surrounding campaigns can generate false or excessive compliance, for example, the erection of fake models to deceive inspecting officials or an all-out effort to exceed central targets through grandiose applications of the policy. Third, campaigns are meant to bypass normal administrative channels, but they also need strong institutions to ensure bureaucratic monitoring and popular feedback. Without those checks, it may become impossible to curb campaign excesses.[103]

Despite these risks, campaigns can produce stunning achievements. In East Asia, the results were immediate, visible, and extensive. As one report on Korea notes, many "status quo oriented bureaucrats" morphed into "relatively enthusiastic activists dedicated to a transformation of the countryside."[104] Even officials who were less than eager to lend support could not stand idly by. In each instance, the state possessed a strong capacity for mobilization, if not for oversight, and it deployed immense resources to overcome bureaucratic resistance. And although campaigns were not necessarily better than market forces and private

[101] Hung 1978.
[102] Looney 2015.
[103] Looney 2020.
[104] Brandt and Cheong 1979, 17.

investment, they did effect broader change in less time than a more lais-sez-faire approach could have accomplished.

Certainly, the outcomes of campaigns and causes of variation deserve more attention than is possible here, but a few things should be clear. First, the payoffs of campaigns, both real and perceived, are one reason they have persisted in East Asia. Second, the factors that seem important for success—appropriate goals, central control, and grassroots participation—do not often come together, which means that positive outcomes are highly contingent. These arguments allow for the possibility of failure but are not as deterministic as some other theories of campaigns.

For James Scott, most state-led development schemes are doomed to fail because they are based on utopian, high-modernist visions of how nature and society should be ordered, and the consequences of failure are more severe in places with strong authoritarian states and weak civil societies. China's disastrous Great Leap Forward, like the Soviet and Tanzanian cases, may easily be explained by that logic.[105] Similar arguments have been made about food self-sufficiency campaigns and collectivization drives in fascist and communist Europe, which left behind complicated and, in some cases, devastating legacies.[106] On a slightly different note, Gregory Witkowski observes that East Germany's campaign state was as weak as it was strong. Campaigns were a substitute for inadequate rural institutions, and because of their ubiquity, the call to action eventually rang hollow.[107]

The larger point that campaigns are prone to failure is well taken, but as this analysis shows, not all cases are like the Great Leap Forward. Moreover, to dismiss campaigns as tragic experiments in social control is unhelpful for understanding their prominence in East Asia. A better approach is to assess the payoffs and costs of campaigns to grasp the practical and political considerations that drive them.

Conclusion

The developmental state concept was initially proposed to explain the region's high-growth outliers, whose effective industrial policies diverged from socialism and capitalism. As interest in the concept grew, intermediate success cases were identified among sectors, agencies, and subnational units within countries that otherwise did not fit the model.

[105] Scott 1998.
[106] Iordachi and Bauerkämper 2014; Pan-Montojo, Prieto, and Cabo Villaverde 2014.
[107] Witkowski 2017.

But its broader applicability was limited and even within East Asia it increasingly appeared as a relic against the backdrop of democratization and neoliberal reforms. For many observers, the 1997 Asian Financial Crisis confirmed the developmental state's shortcomings, especially its inability to create a responsible financial sector, and the model's appeal waned. This history made it all the more surprising when a decade later, after the 2008 global financial crisis, the concept was revived. Contemporary scholarship has focused on how developmental states can overcome market failures and escape the middle-income trap. It points to the success of countries like China and Vietnam and considers the possibility of democratic developmental states emerging in places like India and South Africa. Proponents assert that although the first generation of developmental states grew from a unique set of circumstances, the strategic use of industrial policy is still replicable.[108]

Despite the developmental state's resilience, the literature has remained silent on agriculture's role in development, indirectly suggesting that the rural sector automatically contributes to and benefits from rapid industrialization. This article addresses that gap and offers a corrective to some core assumptions about the East Asian model. A key finding is that rural development, like industrial development, was a state-led phenomenon in East Asia, but one that embodied a different kind of politics, melding technocratic with mobilizational approaches to economic policy to effect transformative change. Whereas previous scholarship has focused on the institutional foundations of development, my work highlights the important role of rural modernization campaigns.

Exploring the deeper politics of campaigns offers a new perspective on East Asia's political economy that challenges the conventional wisdom. First, it shows the enduring salience of revolutionary traditions as East Asian states confronted the problem of agricultural adjustment with the same mobilization tactics forged decades earlier in the context of war. Second, it reveals the populist leanings of so-called bureaucratic authoritarian states and an approach to state-society relations that went far beyond repression or corporatist co-optation. Third, it exposes a process of policy learning that has largely gone undocumented, with East Asian countries recognizing one another and campaigns as belonging to a larger regional model. Taking these insights seriously means revising common assumptions about the nature of developmental states and the process of rural development in East Asia.

[108] Haggard 2018.

References

Amsden, Alice H. 1989. *Asia's Next Giant: South Korea and Late Industrialization*. New York, N.Y.: Oxford University Press.

Anderson, Kym, ed. 2009. *Distortions to Agricultural Incentives: A Global Perspective, 1955–2007*. Washington, D.C.: Palgrave Macmillan and the World Bank.

Ang, Yuen Yuen. 2016. *How China Escaped the Poverty Trap*. Ithaca, N.Y.: Cornell University Press.

Ban, Sung Hwan, Pal Yong Moon, and Dwight H. Perkins. 1980. *Rural Development*. Studies in the Modernization of the Republic of Korea: 1945–1975. Cambridge, Mass.: Harvard University Press.

Bates, Robert H. 1981. *Markets and States in Tropical Africa: The Political Basis of Agricultural Policies*. Berkeley, Calif.: University of California Press.

Bennett, Gordon A. 1976. *Yundong: Mass Campaigns in Chinese Communist Leadership*. Berkeley, Calif.: University of California, Berkeley, Center for Chinese Studies.

Bernstein, Thomas P. 1967. "Leadership and Mass Mobilisation in the Soviet and Chinese Collectivisation Campaigns of 1929–30 and 1955–56: A Comparison." *China Quarterly* 31: 1–47. At https://doi.org/10.1017/S0305741000028691.

Bernstein, Thomas P., and Xiaobo Lü. 2003. *Taxation without Representation in Contemporary Rural China*. New York, N.Y.: Cambridge University Press.

Bezemer, Dirk, and Derek Headey. 2008. "Agriculture, Development, and Urban Bias." *World Development* 36, no. 8: 1342–64. At https://doi.org/10.1016/j.worlddev.2007.07.001.

Bramall, Chris. 2004. "Chinese Land Reform in Long-Run Perspective and in the Wider East Asian Context." *Journal of Agrarian Change* 4, no. 1-2: 107-41. At https://doi.org/10.1111/j.1471-0366.2004.00074.x.

Bramall, Chris. 2009. *Chinese Economic Development*. New York, N.Y.: Routledge.

Brandt, Vincent S. R., and Ji Woong Cheong. 1979. *Planning from the Bottom Up: Community-Based Integrated Rural Development in South Korea*. Essex, Conn.: International Council for Educational Development.

Burmeister, Larry L. 1988. *Research, Realpolitik, and Development in Korea: The State and the Green Revolution*. Boulder, Colo.: Westview.

Cell, Charles P. 1977. *Revolution at Work: Mobilization Campaigns in China*. New York, N.Y.: Academic Press.

Central Committee of the Kuomintang. 1961. *Jiceng minsheng jianshe shiyan yundong zhuanji, zai ban* [The People's Livelihood Construction Experimental Campaign, 2nd ed.]. Taipei, Taiwan: Fifth Division, Central Committee of the Kuomintang.

Central Committee of the Kuomintang. 1972. *Taiwan nongcun jianshe de xin cuoshi* [New Measures for Rural Construction in Taiwan]. Taipei, Taiwan: Division of Cultural Affairs, Central Committee of the Kuomintang.

Chang, Ha-Joon. 1994. *The Political Economy of Industrial Policy*. New York, N.Y.: St. Martin's.

Clegg, Jenny. 2006. "Rural Cooperatives in China: Policy and Practice." *Journal of Small Business and Enterprise Development* 13, no. 2: 219–34. At https://doi.org/10.1108/14626000610665926.

Cumings, Bruce. 1984. "The Origins and Development of the Northeast Asian

Political Economy: Industrial Sectors, Product Cycles, and Political Consequences." *International Organization* 38, no. 1: 1-40. At https://doi.org/10.1017/S0020818300004264.

Dickson, Bruce J. 1993. "The Lessons of Defeat: The Reorganization of the Kuomintang on Taiwan, 1950–52." *China Quarterly* 133, March: 56–84. At https://doi.org/10.1017/S0305741000018191.

Dirlik, Arif. 1975. "The Ideological Foundations of the New Life Movement: A Study in Counterrevolution." *Journal of Asian Studies* 34, no. 4: 945–80. At https://doi.org/10.2307/2054509.

Dix, Robert H. 1985. "Populism: Authoritarian and Democratic." *Latin American Research Review* 20, no. 2: 29-52. At https://www.jstor.org/stable/2503519.

Do, Je-hae. 2009. "New Community Movement Inspires 72 Countries." *Korea Times*, August 30. At https://www.koreatimes.co.kr/www/nation/2020/02/113_50973.html, accessed November 20, 2020.

Doner, Richard F., Bryan K. Ritchie, and Dan Slater. 2005. "Systemic Vulnerability and the Origins of Developmental States: Northeast and Southeast Asia in Comparative Perspective." *International Organization* 59, no. 2: 327–61. At https://doi.org/10.1017/S0020818305050113.

Evans, Peter. 1995. *Embedded Autonomy: States and Industrial Transformation*. Princeton, N.J.: Princeton University Press.

Evans, Peter, and James E. Rauch. 1999. "Bureaucracy and Growth: A Cross-National Analysis of the Effects of 'Weberian' State Structures on Economic Growth." *American Sociological Review* 64, no. 5: 748–65. At https://doi.org/10.2307/2657374.

Fang Ming, and Liu Jun, eds. 2006. *Xin nongcun jianshe zhengce lilun wenji* [Essays on Building a New Countryside, Policy and Theory]. Beijing, China: Zhongguo jianzhu gongye chubanshe.

Fei, John C. H., Gustav Ranis, and Shirley W. Y. Kuo, with Yu-Yuan Bian and Julia Chang Collins. 1979. *Growth with Equity: The Taiwan Case*. New York, N.Y.: Oxford University Press.

Francks, Penelope, with Johanna Boestel and Choo Hyop Kim. 1999. *Agriculture and Economic Development in East Asia: From Growth to Protectionism in Japan, Korea, and Taiwan*. New York, N.Y.: Routledge.

Garon, Sheldon. 1997. *Molding Japanese Minds: The State in Everyday Life*. Princeton, N.J.: Princeton University Press.

Gold, Thomas B. 1986. *State and Society in the Taiwan Miracle*. Armonk, N.Y.: M. E. Sharpe.

Greitens, Sheena Chestnut. 2016. *Dictators and Their Secret Police: Coercive Institutions and State Violence*. New York, N.Y.: Cambridge University Press.

Griffin, Keith, Azizur Rahman Khan, and Amy Ickowitz. 2002. "Poverty and the Distribution of Land." *Journal of Agrarian Change* 2, no. 3: 279–330. At https://doi.org/10.1111/1471-0366.00036.

Ha, Yong-Chool, and Myung-koo Kang. 2011. "Creating a Capable Bureaucracy with Loyalists: The Internal Dynamics of the South Korean Developmental State, 1948–1979." *Comparative Political Studies* 44, no. 1: 78–108. At https://doi.org/10.1177/0010414010381075.

Haggard, Stephan. 1990. *Pathways from the Periphery: The Politics of Growth in the Newly Industrializing Countries*. Ithaca, N.Y.: Cornell University Press.

Haggard, Stephan. 2018. *Developmental States*. Elements in the Politics of Development. New York, N.Y.: Cambridge University Press.

Han Do-hyun (Han Tohyŏn). 2010. "1970 nyŏndae saemaŭl undong esŏ maŭl chidojadŭl ŭi kyŏnghŏm segye" [Life World of Village Leaders During the Saemaul Movement in the 1970s: Focusing on Male Leaders]. *Sahoe wa yŏksa* 88: 267–305.

Han, Seung-Mi. 2004. "The New Community Movement: Park Chung Hee and the Making of State Populism in Korea." *Pacific Affairs* 77, no. 1: 69–93. At https://www.jstor.org/stable/40022275.

Heilmann, Sebastian. 2008. "From Local Experiments to National Policy: The Origins of China's Distinctive Policy Process." *China Journal* 59, January: 1–30. At https://doi.org/10.1086/tcj.59.20066378.

Heydemann, Steven. 1999. *Authoritarianism in Syria: Institutions and Social Conflict, 1946–1970*. Ithaca, N.Y.: Cornell University Press.

Ho, Samuel P. S. 1978. *Economic Development of Taiwan, 1860–1970*. New Haven, Conn.: Yale University Press.

Hong Yeoung-lim (Hong Yŏngnim). 2008. "Saemaŭl undong, 88 Ollimp'ik, kyŏngje kaebal 5 kaenyŏn kyehoek sun" [Saemaul Movement, 1988 Olympics, Five-Year Economic Development Plans]. *Chosŏn ilbo*. March 5. At http://news.chosun.com/site/data/html_dir/2008/03/05/2008030500037.html, accessed October 27, 2020.

Hsiao, Hsin-huang Michael. 1981. *Government Agricultural Strategies in Taiwan and South Korea: A Macrosociological Assessment*. Taipei, Taiwan: Institute of Ethnology, Academia Sinica.

Hung Yu-kun (Hong Yukun). 1978. *Guofu difang zizhi tizhi yu Taiwan shequ fazhan zhi yanjiu* [Research on Sun Yat-sen's Local Self-Governance and Taiwanese Community Development]. Taipei, Taiwan: Zhongyang wenwu gongyingshe.

Huntington, Samuel P. 1965. "Political Development and Political Decay." *World Politics* 17, no. 3 (April): 386-430. At https://doi.org/10.2307/2009286.

Huntington, Samuel P. 1968. *Political Order in Changing Societies*. New Haven, Conn.: Yale University Press.

Iordachi, Constantin, and Arnd Bauerkämper, eds. 2014. *The Collectivization of Agriculture in Communist Eastern Europe: Comparison and Entanglements*. Budapest, Hungary: CEU Press.

Johnson, Chalmers. 1982. *MITI and the Japanese Miracle: The Growth of Industrial Policy, 1925–1975*. Stanford, Calif.: Stanford University Press.

Johnson, Chalmers. 1999. "The Developmental State: Odyssey of a Concept." In Meredith Woo-Cumings, ed., *The Developmental State*. Ithaca, N.Y.: Cornell University Press.

Jowitt, Kenneth. 1975. "Inclusion and Mobilization in European Leninist Regimes." *World Politics* 28, no. 1 (October): 69–96. At https://doi.org/10.2307/2010030.

Kay, Cristóbal. 2002. "Why East Asia Overtook Latin America: Agrarian Reform, Industrialisation and Development." *Third World Quarterly* 23, no. 6: 1073–102. At https://doi.org/10.1080/0143659022000036649.

Keidel, Albert. 1981. *Korean Regional Farm Product and Income, 1910–1975*. Honolulu, Hawaii: University of Hawaii Press.

Kelliher, Daniel R. 1992. *Peasant Power in China: The Era of Rural Reform, 1979–1989*. New Haven, Conn.: Yale University Press.

Kohli, Atul. 2004. *State-Directed Development: Political Power and Industrialization in the Global Periphery*. New York, N.Y.: Cambridge University Press.

Korea Saemaul Undong Center. n.d. At https://www.saemaul.or.kr/eng/, accessed October 27, 2020.

Kuznets, Paul W. 1988. "An East Asian Model of Economic Development: Japan, Taiwan, and South Korea." *Economic Development and Cultural Change* 36, no. S3: S11–43. At https://doi.org/10.1086/edcc.36.s3.1566537.

Ladejinsky, Wolf Isaac. 1977. *Agrarian Reform as Unfinished Business: The Selected Papers of Wolf Ladejinsky*. Louis J. Walinksy, ed. New York, N.Y.: Oxford University Press.

Lee Chung-yüan (Li Zhongyuan). 1979. "Hanguo de Saemaul Undong xin shequ yundong" [Korea's Saemaul New Community Movement]. *Shequ fazhan jikan* 6: 107-12.

Lee, Young Jo. 2011. "The Countryside." In Byung-Kook Kim and Ezra Vogel, eds., *The Park Chung Hee Era: The Transformation of South Korea*. Cambridge, Mass.: Harvard University Press.

Leutert, Wendy. 2020. "Rethinking the Enterprise of the State Through China-Japan Engagement." Manuscript, Indiana University.

Lewis, W. Arthur. 1954. "Economic Development with Unlimited Supplies of Labor." *Manchester School* 22, no. 2: 139–91. At https://doi.org/10.1111/j.1467-9957.1954.tb00021.x.

Li, Cheng. 2005. "The New Bipartisanship within the Chinese Communist Party." *Orbis* 49, no. 3: 387–400. At https://doi.org/10.1016/j.orbis.2005.04.012.

Li Shuishan. 2006. *Hanguo xincun yundong ji qishi* [Korea's New Village Movement and Its Lessons]. Nanning, China: Guangxi jiaoyu chubanshe.

Liao Cheng-hung, Huang Chun-chieh, and Hsiao Hsin-huang Michael. 1986. *Guangfuhou Taiwan nongye zhengce de yanbian: lishi yu shehui de fenxi* [The Development of Agricultural Policies in Postwar Taiwan: Historical and Sociological Perspectives]. Taipei, Taiwan: Academia Sinica, Institute of Ethnology.

Lipton, Michael. 1976. *Why Poor People Stay Poor: Urban Bias in World Development*. Cambridge, Mass.: Harvard University Press.

Liu Jui-chung (Liu Ruizhong). 1991. *Sanminzhuyi shehui jianshe de linian yu shijian: Taiwan diqu shequ fazhan zhi yanjiu* [Theory and Practice of Building a Society Based on the Three Principles of the People: Research on Community Development in Taiwan]. Taipei, Taiwan: Zhengzhong shuju.

Looney, Kristen E. 2015. "China's Campaign to Build a New Socialist Countryside: Village Modernization, Peasant Councils, and the Ganzhou Model of Rural Development." *China Quarterly* 224, December: 909–32. At https://doi.org/10.1017/S0305741015001204.

Looney, Kristen E. 2020. *Mobilizing for Development: The Modernization of Rural East Asia.* Ithaca, N.Y.: Cornell University Press.

Lowenthal, Richard. 1970. "Development vs. Utopia in Communist Politics." In Chalmers Johnson, ed., *Change in Communist Systems*. Stanford, Calif.: Stanford University Press.

Lowenthal, Richard. 1983. "The Post-Revolutionary Phase in China and Russia." *Studies in Comparative Communism* 16, no. 3: 191–201. At https://doi.org/10.1016/0039-3592(83)90005-4.

Ma Xiaohe, ed. 2008. *Zhongguo de xin nongcun jianshe yu Hanguo de xincun yun-dong: 2006 nian Zhong Han jingji hezuo yanjiu hui wenji.* [China's Building a New Countryside and Korea's New Village Movement, 2006 Sino-Korean Economic Cooperation Research Symposium]. Beijing, China: Zhongguo ji-hua chubanshe.

Moon, Chung-in, and Byung-joon Jun. 2011. "Modernization Strategy: Ideas and Influences." In Byung-Kook Kim and Ezra Vogel, eds., *The Park Chung Hee Era: The Transformation of South Korea*. Cambridge, Mass.: Harvard University Press.

Moore, Mick. 1993. "Economic Structure and the Politics of Sectoral Bias: East Asian and Other Cases." *Journal of Development Studies* 29, no. 4: 79–128. At https://doi.org/10.1080/00220389308422296.

Naseemullah, Adnan, and Caroline E. Arnold. 2015. "The Politics of Developmental State Persistence: Institutional Origins, Industrialization, and Provincial Challenge." *Studies in Comparative International Development* 50, no. 1: 121–42. At https://doi.org/10.1007/s12116-014-9164-0.

O'Brien, Kevin J., and Lianjiang Li. 2006. *Rightful Resistance in Rural China.* New York, N.Y.: Cambridge University Press.

O'Donnell, Guillermo A. 1973. *Modernization and Bureaucratic-Authoritarianism: Studies in South American Politics.* Berkeley, Calif.: University of California Press.

OECD. 2012. *Agricultural Policies for Poverty Reduction: A Synthesis.* Paris, France: OECD Publishing. At https://dx.doi.org/10.1787/9789264167698-en.

Oi, Jean C. 1999. *Rural China Takes Off: Institutional Foundations of Economic Reform.* Berkeley, Calif.: University of California Press.

Pan-Montojo, Juan, Lourenzo Fernández Prieto, and Miguel Cabo Villaverde, eds. 2014. *Agriculture in the Age of Fascism: Authoritarian Technocracy and Rural Modernization, 1922–1945.* Turnhout, Belgium: Brepols.

Park, Chung-Hee. 1979. *Saemaul: Korea's New Community Movement.* Seoul, South Korea: Korea Textbook Co.

Park, Jin-Hwan. 1998. *The Saemaul Movement: Korea's Approach to Rural Modernization in the 1970s.* Seoul, South Korea: Korea Rural Economic Institute.

Partner, Simon. 2001. "Taming the Wilderness: The Lifestyle Improvement Movement in Rural Japan, 1925–1965." *Monumenta Nipponica* 56, no. 4 (Winter): 487–520. At https://doi.org/10.2307/3096671.

Pempel, T. J. 1982. *Policy and Politics in Japan: Creative Conservatism.* Philadelphia, Pa.: Temple University Press.

Perry, Elizabeth J. 2006. *Patrolling the Revolution: Worker Militias, Citizenship, and the Modern Chinese State.* Lanham, Md.: Rowman & Littlefield.

Perry, Elizabeth J. 2011. "From Mass Campaigns to Managed Campaigns: 'Constructing a New Socialist Countryside.'" In Sebastian Heilmann and Elizabeth J. Perry, eds., *Mao's Invisible Hand: The Political Foundations of Adaptive Governance in China.* Cambridge, Mass.: Harvard University Asia Center.

Pyle, Kenneth B. 1973. "The Technology of Japanese Nationalism: The Local Improvement Movement, 1900–1918." *Journal of Asian Studies* 33, no. 1: 51–65. At https://doi.org/10.2307/2052885.

Rossmiller, George E. 1972. *Korean Agricultural Sector Analysis and Recommended Development Strategies, 1971–1985.* East Lansing, Mich.: Department of Agricultural Economics, Michigan State University.

Routley, Laura. 2014. "Developmental States in Africa? A Review of Ongoing De-
    bates and Buzzwords." *Development Policy Review* 32, no. 2: 159–77. At https://
    doi.org/10.1111/dpr.12049.
Schmalzer, Sigrid. 2016. *Red Revolution, Green Revolution: Scientific Farming in
    Socialist China.* Chicago, Ill.: University of Chicago Press.
Scott, James C. 1998. *Seeing Like a State: How Certain Schemes to Improve the Hu-
    man Condition Have Failed.* New Haven, Conn.: Yale University Press.
Shin, Gi-Wook, and Do-Hyun Han. 1999. "Colonial Corporatism: The Rural Re-
    vitalization Campaign, 1932–1940." In Gi-Wook Shin and Michael Robin-
    son, eds., *Colonial Modernity in Korea.* Cambridge, Mass.: Harvard University
    Press.
Simmons, Beth A., and Zachary Elkins. 2004. "The Globalization of Liberaliza-
    tion: Policy Diffusion in the International Political Economy." *American Po-
    litical Science Review* 98, no. 1: 171–89. At https://doi.org/10.1017/S0003055
    404001078.
Smith, Graeme. 2018. "The Campaign Rolls On: Rural Governance in China un-
    der Xi Jinping and the War on Poverty." *China: An International Journal* 16, no.
    3: 163–78. At https://muse.jhu.edu/article/703445, accessed December 4, 2020.
Smith, Kerry Douglas. 2001. *A Time of Crisis: Japan, the Great Depression, and Ru-
    ral Revitalization.* Cambridge, Mass.: Harvard University Asia Center.
Smith, Nick R. 2019. "One Village, One Product: Agro-Industrial Village Cor-
    poratism in Contemporary China." *Journal of Agrarian Change* 19, no. 2: 249–
    69. At https://doi.org/10.1111/joac.12301.
Strauss, Julia C. 2017. "Campaigns of Redistribution: Land Reform and State
    Building in China and Taiwan, 1950–1953." In Miguel A. Centeno, Atul
    Kohli, and Deborah J. Yashar, with Dinsha Mistree, eds., *States in the Develop-
    ing World.* New York, N.Y.: Cambridge University Press.
Strauss, Julia C. 2020. *State Formation in China and Taiwan: Bureaucracy, Cam-
    paign, and Performance.* New York, N.Y.: Cambridge University Press.
Stubbs, Richard. 2005. *Rethinking Asia's Economic Miracle: The Political Economy of
    War, Prosperity, and Crisis.* New York, N.Y.: Palgrave Macmillan.
Studwell, Joe. 2013. *How Asia Works: Success and Failure in the World's Most Dy-
    namic Region.* London, U.K.: Profile Books.
Taylor, Jay. 2000. *The Generalissimo's Son: Chiang Ching-Kuo and the Revolutions in
    China and Taiwan.* Cambridge, Mass.: Harvard University Press.
van Gevelt, Terry. 2014. "Rural Electrification and Development in South Ko-
    rea." *Energy for Sustainable Development* 23, December: 179–87. At https://
    doi.org/10.1016/j.esd.2014.09.004.
Vu, Tuong. 2007. "State Formation and the Origins of Developmental States in
    South Korea and Indonesia." *Studies in Comparative International Development*
    41, no. 4: 27–56. At https://doi.org/10.1007/BF02800470.
Wade, Robert. 1990. *Governing the Market: Economic Theory and the Role of Gov-
    ernment in East Asian Industrialization.* Princeton, N.J.: Princeton University
    Press.
Waldner, David. 1999. *State Building and Late Development.* Ithaca, N.Y.: Cor-
    nell University Press.
Wen Tiejun. 2011. *Zhongguo xin nongcun jianshe baogao* [Report on China's Build-
    ing a New Countryside]. Fuzhou, China: Fujian renmin chubanshe.

Weyland, Kurt. 2001. "Clarifying a Contested Concept: Populism in the Study of Latin American Politics." *Comparative Politics* 34, no. 1: 1–22. At https://doi.org/10.2307/422412.

Weyland, Kurt. 2005. "Theories of Policy Diffusion: Lessons from Latin American Pension Reform." *World Politics* 57, no. 2 (January): 262–95. At https://doi.org/10.1353/wp.2005.0019.

White, Tyrene. 1990. "Postrevolutionary Mobilization in China: The One-Child Policy Reconsidered." *World Politics* 43, no. 1 (October): 53–76. At https://doi.org/10.2307/2010551.

White, Tyrene. 2006. *China's Longest Campaign: Birth Planning in the People's Republic, 1949–2005.* Ithaca, N.Y.: Cornell University Press.

Witkowski, Gregory R. 2017. *The Campaign State: Communist Mobilizations for the East German Countryside, 1945–1990.* DeKalb, Ill: NIU Press.

Wong, Joseph. 2004. "The Adaptive Developmental State in East Asia." *Journal of East Asian Studies* 4, no. 3: 345–62. At https://doi.org/10.1017/S1598240800006007.

Wong, R. Bin. 2001. "Formal and Informal Mechanisms of Rule and Economic Development: The Qing Empire in Comparative Perspective." *Journal of Early Modern History* 5, no. 4: 387–408. At https://doi.org/10.1163/157006501X00168.

Woo, Jung-en. 1991. *Race to the Swift: State and Finance in Korean Industrialization.* New York, N.Y.: Columbia University Press.

World Bank. 2019. "Agriculture, forestry, and fishing, value added (annual % growth)." *World Development Indicators*. At https://data.worldbank.org/indicator/NV.AGR.TOTL.KD.ZG, accessed December 4, 2020.

Yager, Joseph A. 1988. *Transforming Agriculture in Taiwan: The Experience of the Joint Commission on Rural Reconstruction.* Ithaca, N.Y.: Cornell University Press.

Ye, Min. 2009. "Policy Learning or Diffusion: How China Opened to Foreign Direct Investment." *Journal of East Asian Studies* 9, no. 3: 399-432. At https://doi.org/10.1017/S159824080000672X.

Zhang Huimin, ed. 2007. *Shehuizhuyi xin nongcun jianshe jianming duben* [Building a New Socialist Countryside, Concise Reader]. Yangling, China: Xibei nonglin keji daxue chubanshe.

Zhongguo liangshi nianjian [China Grain Yearbook]. 2013. Beijing, China: Jingji guanli chubanshe.

Zhongguo nongcun tongji nianjian [China Rural Statistical Yearbook]. 2013. Beijing, China: Zhongguo tongji chubanshe.

Zhongguo nongye fazhan baogao [China Agricultural Development Report]. 2013. Beijing, China: Zhongguo nongye chubanshe.

Author

Kristen E. Looney is an assistant professor of Asian studies and government at Georgetown University. Her area of specialization is comparative politics and the political economy of China and East Asia, with a focus on rural development and governance. Her book, *Mobilizing for Development: The Modernization of Rural East Asia*, was published in 2020. She can be reached at kl705@georgetown.edu.

TWEETING BEYOND TAHRIR
Ideological Diversity and Political Intolerance in Egyptian Twitter Networks

By ALEXANDRA A. SIEGEL, JONATHAN NAGLER,
RICHARD BONNEAU, and JOSHUA A. TUCKER

ABSTRACT

Do online social networks affect political tolerance in the highly polarized climate of postcoup Egypt? Taking advantage of the real-time networked structure of Twitter data, the authors find that not only is greater network diversity associated with lower levels of intolerance, but also that longer exposure to a diverse network is linked to less expression of intolerance over time. The authors find that this relationship persists in both elite and non-elite diverse networks. Exploring the mechanisms by which network diversity might affect tolerance, the authors offer suggestive evidence that social norms in online networks may shape individuals' propensity to publicly express intolerant attitudes. The findings contribute to the political tolerance literature and enrich the ongoing debate over the relationship between online echo chambers and political attitudes and behavior by providing new insights from a repressive authoritarian context.

I. INTRODUCTION

IN the early days of the Arab Spring, Egyptians flocked to social media to organize, follow, and participate in mass protests erupting in Tahrir Square and across the country.[1] Despite the resurgence of authoritarianism in postrevolutionary Egypt, online mass political communication has continued to grow and evolve at extraordinary rates.[2] Social media data provide new insights into the shifting attitudes and communication networks of politically engaged Egyptians during a period of increasing repression and polarization in the aftermath of the July 2013 coup that ousted Egypt's first democratically elected president, Mohamed Morsi. Exploiting the real-time networked structure of Twitter data, this article investigates the relationship between social-network diversity—in this case, the ideological heterogeneity of an in-

[1] Howard and Hussain 2013; Liu, Kliman-Silver, and Mislove 2014.
[2] Jebril, Stetka, and Loveless 2013.

dividual's Twitter network[3]—and political intolerance[4] in the Egyptian Twittersphere.

Twitter's architecture allows for longitudinal analysis of users' connections to elites, media sources, activists, and everyday citizens on the same platform, unlike traditional survey methods that rely on self-reported descriptions of individuals' networks and attitudes at a single moment in time.[5] The platform's structure provides detailed insight into individuals' communication networks and the sources through which they receive political information. Furthermore, the informality and immediacy of tweets offer real-time measures of engaged individuals' sensitive political attitudes, which can be particularly difficult to evaluate in undemocratic polities.[6]

By using Twitter data to obtain temporally granular behavioral measures of network diversity and intolerance, this article offers new tests of the long-theorized relationship between exposure to ideological diversity and political tolerance. In particular, the longitudinal nature of Twitter data enables us to measure changes in network structures and publicly expressed attitudes over time. Doing so allows us to assess whether any observed relationship between network diversity and tolerance is simply a consequence of the initial self-selection of already tolerant individuals into diverse networks, or whether longer exposure to a diverse network is associated with decreased intolerance over time. Additionally, given that Twitter data enable us to examine mass and elite behavior on the same platform, we can evaluate whether the relationship between network diversity and tolerance is driven by exposure to elite,[7] non-elite,[8] or both elite and non-elite networks, providing new insights into the mechanisms by which network diversity might affect political tolerance.

We find that not only is greater network diversity associated with

---

[3] Individuals who follow both Islamists and secularists or a high proportion of moderates on Twitter have diverse networks, whereas individuals who follow only Islamists or only secularists have low levels of diversity.

[4] We define political intolerance as the degree to which citizens oppose extending civil liberties to rival groups or groups advocating disagreeable viewpoints and ideologies; Gibson 2013; Sullivan and Hendriks 2009.

[5] See, for example, Mutz 2006; Gibson 2013; Ikeda and Richey 2009; Bloom and Bagno-Moldavsky 2015; Gibson and Gouws 2005. These studies often rely on survey questions to develop measures of tolerance and network diversity (or how often people interact with those who hold different political views).

[6] For a discussion of the challenges of conducting survey research in the Arab world, see Lynch 2006.

[7] Elites are defined as well-known political or religious leaders with at least ten thousand Twitter followers.

[8] Non-elites are everyday Twitter users, whom we define as individuals with fewer than ten thousand Twitter followers.

lower levels of intolerance, but also that spending an additional year in a diverse network is associated with a significant decrease in the use of intolerant rhetoric over time. Our results provide suggestive evidence that exposure to both elite and non-elite network diversity—particularly the presence of more moderates in an individual's network—may reduce political intolerance. By examining the relationship between exposure to elite and non-elite ideological diversity and intolerant behavior over time, this article contributes to the political tolerance literature and adds to a growing body of research examining the relationship between online networks and political attitudes and behavior.

II. Political Intolerance in the Egyptian Twittersphere

Since the outbreak of the Arab Spring protests in late 2010, Twitter use among Arab citizens has grown dramatically.[9] Collectively, the Arab world generates 27.4 million tweets per day, with 20 percent of these tweets coming from Egypt.[10] On average, Egyptian Twitter users produce 151 million tweets each month.[11] Although Facebook and WhatsApp are more popular than Twitter in Egypt,[12] Twitter is particularly conducive to political discussions because users follow accounts based on particular interests and tend to focus on sharing information rather than reciprocal social interaction.[13]

Twitter has facilitated the creation of flexible networks of political communication outside of traditional civil society and media centers.[14] In Egypt, it evolved from a tool initially used by educated youth and activists into a platform for diverse individuals, including students, blue-collar workers, and even the elderly.[15] Elites across Egypt's political spectrum—from aged generals to clerics and prominent members of the Muslim Brotherhood—have Twitter accounts. Because Twitter allows connections to elites and non-elites on the same platform, its networks not only transcend offline social connections, but also grant everyday citizens unprecedented access to elites.

Despite early optimism during the Egyptian Revolution that social media could facilitate democratic transition and political engagement, social media has also been used to achieve short-term political goals, to

[9] Liu, Kliman-Silver, and Mislove 2014.
[10] Mourtada 2016.
[11] Salem 2017.
[12] Mourtada 2016.
[13] Halpern, Valenzuela, and Katz 2017; Smith, Fischer, and Yongjian 2012.
[14] Anduiza, Jensen, and Jorba 2012; Tufekci and Wilson 2012.
[15] El-Khalili 2013.

manipulate public opinion, and even to incite violence.[16] In the years following the 2013 coup, Egyptians witnessed widespread civil liberties violations and human rights abuses. The Muslim Brotherhood and most other Islamist organizations were excluded from politics, the media, and civil society. Secular[17] activists were jailed, and many of their organizations were banned.[18] From 2013 to 2017, approximately sixty thousand Egyptians were imprisoned and Egyptian authorities built ten additional prisons to accommodate them. In early 2017, reports estimated that an average of three to four forced disappearances occurred per day, and local human rights groups documented hundreds of extrajudicial killings and thousands of cases of torture.[19]

Following the coup, elite rhetoric became increasingly intolerant, often calling for the total exclusion of opponents from political and public life.[20] In May 2014, for example, on live television, a representative of President Abdul Fatah al-Sisi's election campaign called for the arrest of all Muslim Brotherhood members in Egypt, asserting that those who did not love Egypt "should be hit with shoes." Similarly, in a 2014 interview on a satellite channel, prominent host Mohammed Moussa openly called for the arrest and execution of atheists and secularists as part of a "war against destructive ideas."[21]

In this tense political environment, social media rhetoric by political elites and politically engaged citizens was often inflammatory. Among secular Egyptians, tweets commonly called for death sentences for Muslim Brotherhood leaders, decried Islamists as terrorists, and proclaimed that the Muslim Brotherhood would never again return to power in Egypt. On the Islamist side, tweets inciting violence or proclaiming that Egypt was for Islamists only were also common. Tweets documenting and denouncing the regime's civil liberties violations proliferated as well. Examining these trends, Ingmar Weber, Venkata Garimella, and Alaa Batayneh find increasing polarization among Egyptian Twitter users after the coup, with both Islamist and secular users more likely to retweet and mention users of the same political orientation.[22] Robert Kubinec and John Owen show that the coup led to increased polarization in the Egyptian Twittersphere that diffused transnationally as well.[23] Similarly, Marc Lynch, Dean Freelon, and Sean Aday

---

[16] Morrow and al Omrani 2014.

[17] For ease of interpretation, throughout this article we refer to non-Islamist Egyptians as "secular." Although many of these actors may be religious, they are not supporters of political Islam.

[18] Dunne and Williamson 2014; Hamzawy 2017.

[19] Hamzawy 2017.

[20] Boukhars et al. 2014.

[21] Rollins 2014.

[22] Weber, Garimella, and Batayneh 2013.

[23] Kubinec and Owen 2018.

demonstrate that Egypt's online public became progressively segregated into clusters of like-minded individuals who frequently expressed fear, paranoia, and mistrust.[24]

But what is the relationship between these increasingly polarized Egyptian Twitter networks and political behavior? Does spending time in a more ideologically diverse or homogeneous online network change the content of users' political discourse? Or do people simply self-select into networks that reflect their existing beliefs and continue to update their networks as their attitudes shift? In the tense political climate of postcoup Egypt, the Egyptian Twittersphere provides an ideal setting for studying the relationship between network diversity and political intolerance over time.

III. Theoretical Motivation and Expectations

Social scientists have long posited that diverse social networks contribute to political tolerance by broadening individuals' horizons and exposing them to new ideas.[25] This research was initially motivated by Gordon Allport's contact hypothesis, which suggests that interactions with outgroups alleviate hostility by undermining false stereotypes and drawing attention to individual characteristics rather than to group generalizations.[26] Numerous studies demonstrate that interactions with members of an outgroup create greater awareness of legitimate rationales for opposing views and motivate citizens to care about upholding the civil liberties of those with whom they disagree or whom they dislike.[27] Although the majority of this work has been conducted in the United States,[28] social network diversity has been found to be associated with increased political tolerance in a variety of cultural contexts, including Israel,[29] Japan,[30] and South Africa.[31] But other work suggests that the relationship may be conditional on positive social interactions, threat perception, or other factors.[32]

Exposure to diversity need not come from direct interpersonal con-

---

[24] Lynch, Freelon, and Aday 2016.

[25] Blau 1974; Nunn, Crockett, and Williams 1978.

[26] Allport 1954; Pettigrew 1998.

[27] See Mutz 2002 and Mutz 2006 for an overview of the literature on intergroup contact and political tolerance.

[28] For examples of studies on the positive effect of social network diversity on tolerance in the US context, see Huckfeldt, Mendez, and Osborn 2004; Mutz 2002; Mutz 2006; Peffley, Knigge, and Hurwitz 2001; Robinson 2010; Tadmor, Tetlock, and Peng 2009.

[29] Bloom and Bagno-Moldavsky 2015.

[30] Ikeda and Richey 2009.

[31] Gibson and Gouws 2005.

[32] Haas and Cunningham 2014; Bloom and Bagno-Moldavsky 2015; Roccas and Amit 2011.

tact with outgroup members. Instead it may come from exposure to print and electronic media. Craig Calhoun argues, "In modern societies, most of the information we have about people different from ourselves comes not through any direct relationships, even the casual ones formed constantly in urban streets and shops. Rather it comes through print and electronic media."[33] Similarly, Diana Mutz writes, "The future of communication across lines of political difference lies in technologies that transcend geographic space."[34] Thus, online social network diversity may serve as a form of indirect contact—albeit quite different from the direct contact Allport envisioned in the 1950s.

Indeed, more recent research suggests that online social networks force users to confront political information that they would otherwise avoid.[35] Online social networks have sometimes been optimistically portrayed as a means of spreading information and proliferating points of contact across political and sectarian divides.[36] For example, Ceren Budak and Duncan Watts find that exposure to members of opposition groups online during the 2013 Gezi Park protests in Turkey led to greater support of these groups.[37] But other studies indicate that although social media increases exposure to diverse political information, it also enables individuals to seek out content that reinforces their existing views, creating echo-chamber environments that foster social extremism and political polarization.[38]

Most of the literature on network diversity and tolerance (both online and offline) focuses on the link between non-elite or peer networks and tolerance.[39] We therefore first explore the relationship between non-elite network diversity and tolerance. Motivated by theory and empirical findings in the political tolerance literature, we hypothesize:

—H1a. Non-Elite Network Diversity: Egyptian Twitter users in diverse non-elite networks will express less political intolerance than those in more homogeneous networks.

In line with past studies of the relationship between network diversity and tolerance outlined above, we first expect that greater non-elite

[33] Calhoun 1988.
[34] Mutz and Martin 2001.
[35] Zhang, Fuehres, and Gloor 2011; Wojcieszak and Mutz 2009.
[36] Aday et al. 2010.
[37] Budak and Watts 2015.
[38] Farrell 2012; Garrett 2009; Aday et al. 2012; Bakshy, Messing, and Adamic 2015. Recent studies also suggest that the degree to which social media exposes people to ideologically diverse opinions may be asymmetrical. In the US, for example, liberals may be more likely to engage in cross-ideological online communication than conservatives. Barberá et al. 2015.
[39] But see Dyck and Pearson-Merkowitz 2014 for an exception.

network diversity will be associated with lower levels of intolerance. But merely observing this association tells us nothing about the direction of the relationship. For example, imagine a hypothetical Egyptian Twitter user, Ahmed, who has both a relatively diverse network and a relatively low level of intolerance. Does Ahmed choose to cultivate a Twitter network that reflects his preexisting level of tolerance? Does exposure to diversity through Ahmed's Twitter network cause him to become more tolerant? Do both processes play out at the same time?

Examining how Ahmed's network and level of tolerance change over time helps us to assess the hypothesized association between network diversity and tolerance. Even if Ahmed already expressed low levels of intolerance when he initially chose to create a diverse network, does he become more tolerant after spending more time in his diverse network? We hypothesize that beyond the general association between network diversity and tolerance, individuals who spend more time in diverse networks will become more tolerant over time:

—H1b. Non-Elite Network Diversity over Time: Egyptian Twitter users in diverse non-elite networks will express less political intolerance after spending more time (from t_1 to t_2) in their diverse networks, whereas those in more homogeneous networks will express more intolerance over time.

What might lead Ahmed to express less intolerance after spending more time in a diverse network? Does more time in the network make Ahmed less likely to tweet intolerant content because he has been repeatedly exposed to outgroup members or to a wider range of ideas than he would otherwise encounter? Or do exogenous events lead Ahmed to continue to adapt both the structure of his network and the content of his tweets as his beliefs evolve over time?

Figure 1 illustrates H1b, depicting Ahmed's network diversity and tolerance at t_1 and a year later at t_2. The composition of Ahmed's Twitter network does not change and he becomes more tolerant over the course of spending an additional year in a diverse network.

Beyond descriptively assessing whether exposure to network diversity is associated with increased intolerance over time, we can interpret H1b causally if the following assumptions are met:

—1. Changes in Twitter users' tolerance over time from t_1 to t_2 are not associated with a simultaneous change in the diversity of Twitter users' networks.

—2. Twitter users' decisions to select into either diverse or homogeneous networks at t_1 are independent of a subsequent bidirectional change in Twitter users' behavior, whereby users in diverse networks express less intolerance and users in homogeneous networks express more intolerance
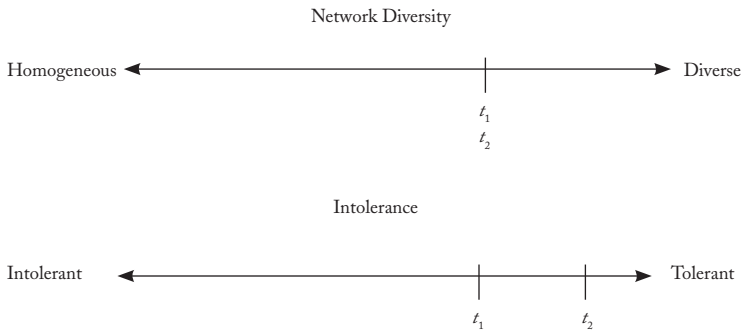
Network Diversity

Homogeneous ◄─────────────────────── ► Diverse

t_1
t_2

Intolerance

Intolerant ◄─────────────────────── ► Tolerant

t_1 t_2

FIGURE 1
NETWORK DIVERSITY AND INTOLERANCE OVER TIME[a]

[a] Figure displays the relationship between network diversity and intolerance over time (from t_1 to t_2) for a hypothetical Egyptian Twitter user, Ahmed.

at t_2 than they did at t_1. Along these lines, the ignorability assumption requires that individuals who self-select into a particular type of network at t_1 are not more likely to have a particular subsequent change in their behavior.

—3. The Stable Unit Treatment Value Assumption (SUTVA) requires that the response of a particular individual depends only on the treatment to which that unit was assigned, not the treatments of others.

The first assumption is met in Figure 1, as the diversity of Ahmed's network does not change over time from t_1 to t_2. But is this a reasonable assumption for Egyptian Twitter users more broadly? Research on the Egyptian Twittersphere suggests that Egyptians became increasingly polarized into Islamist and secular networks in the immediate aftermath of the 2013 coup and subsequent violent repression of Islamists.[40] The period after the coup was marked by a great unfollowing in which anti-Brotherhood social media users stopped following Islamists, and vice versa, en masse, with social media campaigns calling on people to defriend Islamists. This suggests that by the time of our study in 2015, individuals who intended to change the ideological composition of their networks would likely have already done so. We can test empirically whether this first assumption is met by evaluating whether individuals in our study changed the diversity of their network from t_1 to t_2 and the degree to which this change is associated with simultaneous changes in tolerance.

[40] Lynch, Freelon, and Aday 2016.

The second assumption would be violated if (1) individuals like Ahmed who self-select into a particular network were more likely to change their behavior in a particular direction over time, or (2) a confounder both drove individuals like Ahmed to select into a diverse network by t_1 and caused them to express less intolerance at t_2 than they did at t_1, while driving others to both self-select into more homogeneous networks by t_1 and to express more intolerance at t_2 than they did at t_1. For example, if Ahmed held relatively tolerant beliefs at t_1, he might self-select into a diverse network at t_1, and then exogenous events between t_1 and t_2 might lead him—as an already tolerant person—to express lower levels of intolerance over time, independent of his network. At the same time, Twitter users who held relatively intolerant beliefs at t_1 might select into homogeneous networks and then been driven by the same events to express even more intolerance by t_2 than they did at t_1. Although this example is plausible, it would require that the confounder drive initial selection into networks and drive individuals in those networks to change their behavior in different directions.

Regarding the third assumption, in the event that SUTVA were violated, we would need to interpret our results as the effect of an individual's network diversity on their level of tolerance given that other people's behavior is also simultaneously being impacted by the level of network diversity to which they are exposed. Although this interpretation is narrower, it is also more in line with the naturalistic experience of participating in an online social network.

Finding evidence for H1b will, at a minimum, improve our descriptive understanding of how the relationship between network diversity and tolerance evolves over time. If we believe that our three assumptions are reasonable, then we can offer suggestive evidence that exposure to network diversity increases tolerance over time. Because these assumptions are difficult to test, we also conduct sensitivity analysis to understand how strong a confounder would need to be to change our results.

Although studies of social network diversity and political tolerance tend to focus on interactions among ordinary citizens or non-elite network diversity, exposure to elite cues may also play a role.[41] Research in the US and in comparative contexts suggests that citizens rely on simple and reliable cues from elites to make policy judgments.[42] Elites

---

[41] Dyck and Pearson-Merkowitz 2014.
[42] See, for example, Lupia 1994; Lupia and McCubbins 1998; Druckman 2001; Druckman, Peterson, and Slothuus 2013; Brader and Tucker 2008.

often have incentives to shape public opinion toward an outgroup, and when elites perceive a threat to their power, they may advocate punitive policies or work to mobilize public opinion against outgroups.[43]

Exposure to more diverse elite cues—particularly through the opening of the media following democratic transitions—has been shown to move public opinion against policies that violate the civil liberties of an outgroup.[44] Connecting elite cue theory to political tolerance, Allport notes that the effect of intergroup contact on tolerance relies on "authority support."[45] Elites alert the public to social norms within their political ingroup and shape attitudes and behaviors toward outgroups.[46] On Twitter, individuals can choose to follow particular elites, cultivating networks that shape their exposure to elite cues. The literature on the key role of elite actors motivates our second set of hypotheses:

> —H2a. Elite Network Diversity: Egyptian Twitter users who follow more ideologically diverse elite actors will exhibit less political intolerance relative to those who follow more ideologically homogeneous elites.
> —H2b. Elite Network Diversity over Time: Egyptian Twitter users in more diverse elite networks will exhibit less political intolerance over time (from t_1 to t_2), whereas those in more homogeneous elite networks will exhibit more intolerance over time.

We hypothesize that a negative relationship exists between elite network diversity and intolerance generally, and that spending more time in a diverse elite network will be associated with a decrease in an individual's expression of intolerance from t_1 to t_2. Although this hypothesis is descriptive, it can again be interpreted causally if individuals do not change the diversity of their networks from t_1 to t_2 and if we assume that the decision to initially select into a particular network at t_1 is not associated with a future bidirectional change in tweeting behavior—toward less intolerance for those in diverse networks and more intolerance for those in homogeneous networks—between t_1 and t_2.

IV. Measurement and Data

Measuring Political Intolerance on Twitter

Tolerance is traditionally defined in the political science literature as the degree to which citizens support the extension of civil liberties to polit-

---

[43] King and Wheelock 2007; Baumer, Messner, and Rosenfeld 2003.
[44] Stein 2013.
[45] Allport 1954.
[46] Pettigrew 1998; Dyck and Pearson-Merkowitz 2014.

ical rivals, including groups advocating highly disagreeable viewpoints and ideologies.[47] Civil liberties include the right to vote, to participate in political parties, to organize politically, to free speech and assembly, and to a fair and speedy trial, as well as other basic civil rights. Political intolerance, by contrast, is a form of exceptionalism that seeks to limit the civil liberties of particular citizens or groups, and hence, their ability to compete for political power.[48]

The three main approaches to measuring political tolerance use survey questions that ask, respectively, whether the respondent believes political activities should be allowed for members of a specific outgroup (fixed-group approach), whether political activities should be allowed for members of a group that the respondent reports liking least (least-liked approach), and whether the respondent approves of polices that would limit civil liberties generally.[49] A crucial aspect of tolerance is the so-called "objection precondition," which states that tolerance is restraint of the urge to repress one's political enemies or members of an outgroup.[50] For example, in the US context, Democrats cannot be said to be "tolerant" of other Democrats because they are both members of the same political group, but they may or may not be tolerant of Republicans (their primary political competitors). Measuring tolerance is therefore a two-step procedure. First, researchers must establish that an individual dislikes or is in competition with a political group. Then they must measure the extent to which the individual supports or opposes the political rights of that group.[51] Measuring intolerance is a simpler endeavor. Expressing a desire to limit any group's civil liberties implies dislike.

We operationalized Egyptian Twitter users' political intolerance by measuring the frequency and the relative frequency with which they tweeted intolerant content online. We chose to measure intolerance—rather than tolerance—because it is easier to measure, given that the objection precondition is, by definition, satisfied. We measured intolerance through a two-step process. We first classified tweets as relevant or not to civil liberties in Egypt. We next determined whether these relevant tweets were or were not intolerant.

We began with a data set of approximately 130 million tweets related to Egypt or Egyptian politics containing the Arabic key words "Egypt," "Sisi," "Morsi," "Muslim Brotherhood," "coup," "protest," "revolution,"

[47] Gibson 2013; Sullivan, Piereson, and Marcus 1979.
[48] Gibson 2013; Gibson and Gouws 2005.
[49] Gibson and Gouws 2005; Sullivan and Hendriks 2009; Petersen et al. 2011.
[50] Gibson and Gouws 2005.
[51] Sullivan and Hendriks 2009.

or "military rule," as well as a data set of approximately eighteen million tweets with geolocation metadata indicating they were sent from Egypt. Many tweets in these initial data sets were likely irrelevant to Egyptian politics—not to mention civil liberties—because they contained the very broad key word "Egypt" or were included simply because geolocation metadata indicated that they were sent from Egypt. To obtain a training data set that contained a sufficient number of relevant tweets, we first identified the two hundred most common terms in our data and selected those that were plausibly relevant to civil liberties, including "court," "constitution," "prison," "arrest," "rights," "law," "coup," and "military."[52] We next took a sample of five thousand tweets from the data set that contained key words that we might expect to be relevant to civil liberties.

We then used the data enrichment platform Figure Eight to code each of these five thousand tweets according to whether they were relevant to civil liberties in Egypt. The coding was done by three native Arabic-language speakers. Using the human-coded data, we trained a Naive Bayes classifier to classify the roughly 148 million tweets in our data set as relevant (or not) to civil liberties in Egypt.[53] After using our classifier to identify relevant tweets, we used human-coded data to train a second classifier to classify relevant tweets as intolerant or not.[54]

Under our coding scheme for the training data, translated and presented in Appendix A of the supplementary material, intolerant tweets support restricting civil liberties in Egypt. These tweets include those that favor civilian arrests, death sentences, or torture; those that support limiting the right to free speech, protest, or assembly; and those that advocate for banning political parties or excluding certain groups from formal or informal political participation.[55]

An example of an intolerant tweet found in the Twittersphere shows a graphic image portraying former President Morsi and other promi-

---

[52] A full list of these key words and their translations is provided in Appendix A of the supplementary material; Siegel et al. 2021. If we had trained our classifier using random samples of tweets from the Egyptian Twittersphere, the vast majority of the tweets would have been irrelevant to civil liberties and the classifier would have been trained on only the small number of tweets that were relevant. We chose this particular set of key words because they occur most commonly in the data and capture multiple theoretically motivated dimensions of political (in)tolerance.

[53] A detailed description of this process and the performance of our classifier is described in Appendix A of the supplementary material.

[54] Specifically, we used the Figure Eight platform to code fifty thousand relevant tweets as intolerant or tolerant. Because intolerant tweets appear relatively rarely, we needed much more human-coded data to create our training data set for the intolerance classifier than was needed to create the training data set for the relevance classifier. More details on this process are provided in Appendix A of the supplementary material.

[55] Siegel et al. 2021.

nent Muslim Brotherhood members on their way to the gallows. The Arabic text of the tweet states, "God willing, the traitors, spies, and terrorists will be executed. The people want the execution of the Brotherhood." Because this tweet clearly opposes the civil liberties of Islamists, it is classified as intolerant. Translated examples of randomly selected relevant and intolerant tweets are provided in Appendix A. Our intolerance classifier enabled us to obtain longitudinal measures of intolerance for any Egyptian Twitter users who repeatedly tweeted content relevant to civil liberties in Egypt.

MEASURING TWITTER NETWORK DIVERSITY

To measure network diversity, we first developed a measure of political orientation—namely, whether a Twitter user was Islamist, secular, or neither (moderate).[56] Following the 2013 coup, the Islamist-secular divide became more deeply entrenched as the primary fault line in Egyptian politics, and it is therefore the political dimension on which we assessed network diversity.[57]

It is possible to estimate the political orientation of Twitter users based on the political elites whom they choose to follow.[58] Pablo Barberá[59] argues that the decision to follow a politician or political account on Twitter is a "costly signal" that provides information about a Twitter user's ideological position. This is grounded in the assumption that Twitter users prefer to follow accounts that share their political orientations because social networks are homophilic[60] and individuals gravitate toward those with similar leanings. Borrowing from Yosh Halberstam and Brian Knight,[61] we therefore measured the political orientation of Egyptian Twitter users based on the politicians whom they followed on Twitter.

We compiled a list of all Egyptian political elites on Twitter who had more than ten thousand followers and well-known political affiliations. This yielded a list of eighty-five Egyptian political elites (forty-four secular and forty-one Islamist).[62] We used the political affiliation

[56] Moderates, by whom we mean users who are neither Islamist nor secular, may include true moderates (those who fall somewhere between Islamists and secularists on the ideological spectrum) or people whose political orientation cannot be determined because they do not follow any elites.

[57] Of course, there is additional diversity within the Egyptian political spectrum, including the Salafi–Muslim Brotherhood division on the Islamist side and the liberal-promilitary division among secular Egyptians, but here we focus on the Islamist-secular divide, which is the most salient.

[58] Halberstam and Knight 2016.

[59] Barberá 2015.

[60] McPherson, Smith-Lovin, and Cook 2001.

[61] Halberstam and Knight 2016.

[62] See Appendix A of the supplementary material for more details and the list of elites; Siegel et al. 2021.

(Islamist or secular) of these elites to classify the political orientation of each Twitter user as well as the political orientation of each Twitter user's non-elite friends.[63] Users were classified as Islamist if 60 percent or more of the elites they followed were Islamist, secular if 60 percent or more of the elites they followed were secular, and moderate if Islamists and secularists each constituted between 40 and 60 percent of the elites they followed. We chose the 60 percent cutoff point because it appeared to be accurate upon examining the networks of a hundred politically engaged Egyptians on Twitter whose political orientation could be easily determined based on the content of their tweets.[64] But to ensure that our threshold choice did not bias our analysis, we also tried a series of other percentages—ranging from 55 percent to 85 percent—and found that using these alternatives did not substantively change our results.[65] Because the elites whom we chose were either secular or Islamist, our definitions of secular, Islamist, and moderate Egyptian Twitter users were mutually exclusive, allowing us to code the political orientation of each user in our sample and each of their non-elite friends. Users who did not follow any political elites were classified as moderate.

Users in completely homogeneous networks followed only Islamist or only secular elites, whereas those in completely diverse networks followed an even proportion of Islamist and secular elites. Elite network diversity can therefore be measured on a zero to one continuous scale, with those in completely homogeneous networks receiving a score of 0 and those in maximally diverse networks receiving a score of 1 as follows:

$$Elite\ Network\ Diversity = 1 - |\ \%\ Secular\ Elites - \%\ Islamist\ Elites\ |$$

To measure non-elite network diversity, we first classified the political orientation of each non-elite user whom individuals followed on Twitter. Our first measure of non-elite network diversity is analogous to our measure of elite network diversity:

$$Non\text{-}Elite\ Network\ Diversity = 1 - |\ \%\ Secular\ Non\text{-}Elites - \%\ Islamist\ Non\text{-}Elites\ |$$

[63] "Friends" are accounts that a Twitter user follows. By contrast, "followers" are individuals who follow the user.

[64] In a random sample of one hundred accounts, Islamist users frequently had profile metadata and photos that included Brotherhood slogans and symbols, and secular Egyptians often had pro-Sisi or promilitary imagery and text associated with their profiles.

[65] Results using these different thresholds are displayed in Table B7 of the supplementary material.

We also developed a second measure of non-elite network diversity that makes the diversity score higher for users who followed large numbers of moderates, accounting for exposure to secular, Islamist, and moderate users in a network. This measure is also on a zero to one continuous scale:

$$\textit{Non–Elite Network Diversity} =$$
$$(\textit{Non–Elite Network Diversity + Moderate Non–Elites}) \, / \, 2$$

DATA

As we describe above, our initial data set consisted of tweets that included either the Arabic word for Egypt or a reference to Egyptian politics, or geolocation metadata indicating that they were sent from Egypt. This initial data set contained 148,361,405 tweets. Of these, 11,141,302 were sent by non-elite Egyptians.[66] Using our relevance classifier, we determined that 3,050,983 of these tweets were relevant to civil liberties in Egypt. These relevant tweets were produced by 159,619 unique Egyptian Twitter users.

To identify Egyptian Twitter users who were likely to tweet about civil liberties, we next limited this data set to users who had published at least four tweets from March 2014 to April 2015 that were relevant to those liberties. Doing so produced a data set of 9,400 non-elite Egyptian Twitter users. This number is relatively small given the size of our tweet collection because the majority of tweets in the initial data sets were produced by news organizations, verified Twitter accounts, and users who did not appear to be located in Egypt. Further, rather than attempting to capture the full universe of intolerant tweets in the Egyptian Twittersphere, our classifier prioritized precision over recall to ensure that we had accurate data with which to test our hypotheses. Many users deleted their accounts, stopped tweeting, or made their accounts private in the period under study. We also removed users with more than five thousand friends at the time of our initial collection, although this cohort accounted for less than 1 percent of the total data set. To identify bots in our data, we ran all our users through Botometer,[67] an automated bot-detection tool. The results suggested that fewer than 2 percent of the users in our data had above a 50 percent likelihood of being classified as a bot. We also manually examined a random sample of 500 of the 9,400 accounts searching for common indicators of bot activity[68] and did not identify any bots.

---

[66] Tweets were collected using Twitter's streaming application programming interface (API). See Appendix A of the supplementary material for details on determining users' location.

[67] At https://botometer.iuni.iu.edu/#!/faq, accessed February 9, 2021.

[68] Stukal et al. 2017.

We next collected network data for each of the 9,400 Twitter users every day from April 2015 through October 2016 to classify their political orientation, to measure their network diversity, and to evaluate how their networks changed over time. In October 2016, at the end of the data collection period, we scraped the 3,200 most recent tweets of each of these users.[69] For most of these users, these tweets covered several years of Twitter activity and could even date back to their earliest tweet. This collection gave us a data set of almost 29 million tweets produced by 9,400 politically engaged Egyptian Twitter users. As Figure 2 demonstrates, the users in our data set were geographically diverse. Although they were concentrated in Cairo, their locations generally mirrored the distribution of the Egyptian population along the Nile.

Our data set consists of 1,341 Islamist, 3,748 moderate, and 4,311 secular Egyptian Twitter users. As Table A3 of the supplementary material demonstrates, elite and non-elite network sizes varied dramatically, with users following from zero to sixty-two elites and having from zero to 48,800 non-elite friends.[70] The number of relevant tweets and the number of intolerant tweets also differed a great deal across users.[71] Further, as Figure 3 suggests, network diversity varied quite a bit, with concentrations of users in both diverse and homogeneous networks. The large number of users with maximally diverse or maximally homogeneous elite networks was partly driven by the fact that many users followed very few elites, as Table A3 suggests.[72]

V. EMPIRICAL STRATEGY AND RESULTS

We first assessed whether users in more ideologically diverse networks exhibited lower levels of intolerance relative to users in more homogeneous networks (H1a and H2a). When measuring intolerance as a raw count of intolerant tweets, we used a quasi-Poisson model. Quasi-Poisson regression accounts for overdispersion in count data, address-

[69] Twitter's API limits enabled us to scrape the 3,200 most recent tweets and retweets in a given user's timeline. Our results do not change significantly when retweets are excluded from the analysis.

[70] Although some individuals did not follow any elites and some did not follow any non-elites, everyone in the sample followed at least one person. Network statistics are from October 2016, the end of the data collection period.

[71] Although our sample only included users who had sent at least four relevant tweets from March 2014 through April 2015, when we scraped the 3,200 tweets for each of these users in October 2016, the end of our data collection period, some users who tweeted frequently had tweeted fewer than four relevant tweets in their most recent 3,200 tweets. The minimum relevant tweet count in Table A3 is therefore one rather than four; Siegel et al. 2021.

[72] Descriptive statistics of users' tweet content, network size, and network diversity are provided in Appendix A of the supplementary material. The process of collecting data to measure intolerance and network diversity described above is summarized in the flowchart in Figure A2; Siegel et al. 2021.
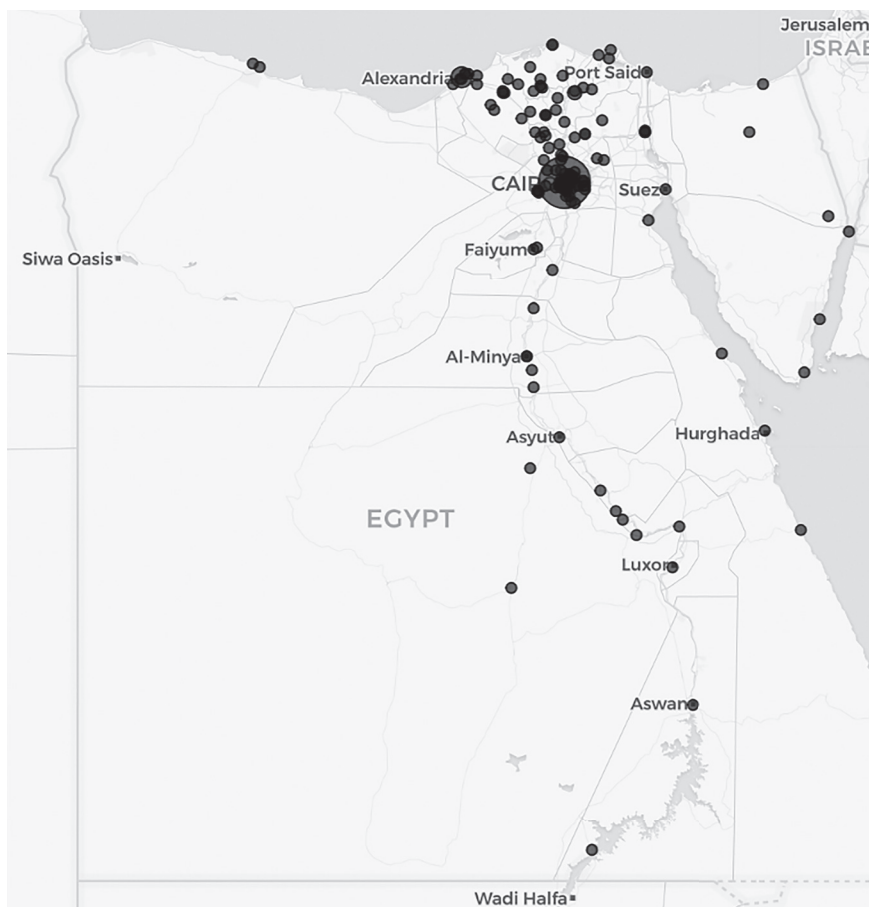
FIGURE 2
GEOGRAPHIC DISTRIBUTION OF TWITTER USERS IN SAMPLE[a]

[a] Figure shows the locations of geolocated Twitter users in our data set. We have geolocation data for 6,431 of the 9,400 users in our sample.

ing the fact that many users in our sample rarely tweeted intolerant tweets in the period under study. Although intolerant tweets are a rare outcome,[73] because thousands of Egyptians produce this rhetoric, a small difference in individual tweet volume represents a meaningful reduction in the volume of intolerant tweets in the Egyptian Twittersphere. When measuring intolerance as a proportion of relevant tweets, we used OLS regressions.

[73] We therefore also replicate this analysis using a negative binomial model, yielding similar results reported in Table B6 of the supplementary material.
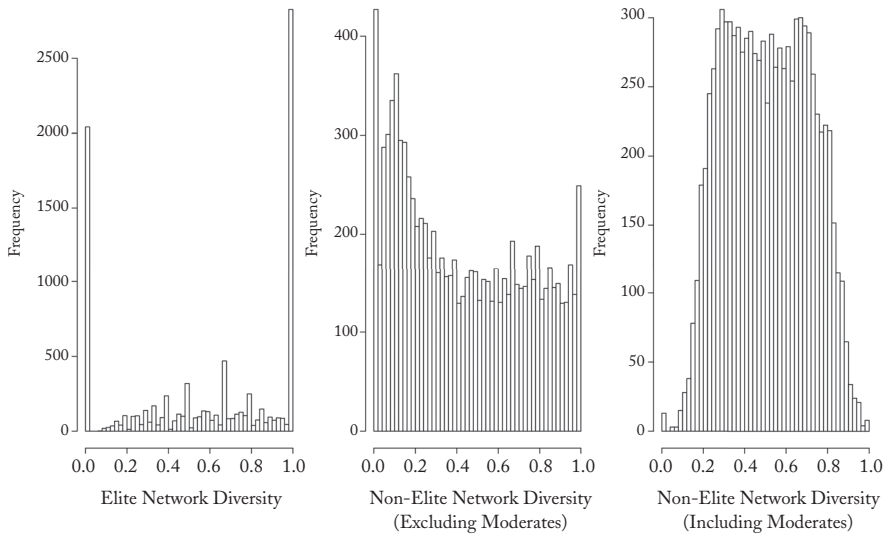
FIGURE 3
DISTRIBUTIONS OF TWITTER NETWORK DIVERSITY[a]

[a] Figure shows variation in elite and non-elite network diversity (including and excluding moderates) of the 9,400 Twitter users in our data set.

Because having a larger political communication network may increase tolerance,[74] we included elite network size and non-elite network size as control variables measured as log(*number of elite friends*) and log(*number of non-elite friends*).[75] We also included an Islamist dummy variable to account for political orientation.[76] Additionally, we controlled for the time a user had been on Twitter, measured as log(*number of days on Twitter*) and—in the count model—included a control for each user's number of relevant tweets.

Figure 4 plots the coefficients from our quasi-Poisson model, demonstrating that both elite and non-elite network diversity are associated with statistically significant lower levels of intolerance whether we measure non-elite network diversity including or excluding moderates. Tables B1 and B2 of the supplementary material demonstrate that these results are consistent whether measuring intolerance as a count or as a proportion, with and without controls. Converting the point esti-

[74] Gibson 2001.
[75] Network size is not significantly correlated with either elite or non-elite network diversity, so we do not have any concerns about collinearity here. Because we exclude moderates from our measure of non-elite diversity, our network size variable captures this information as well.
[76] This is not significantly correlated with elite (or non-elite) network diversity.
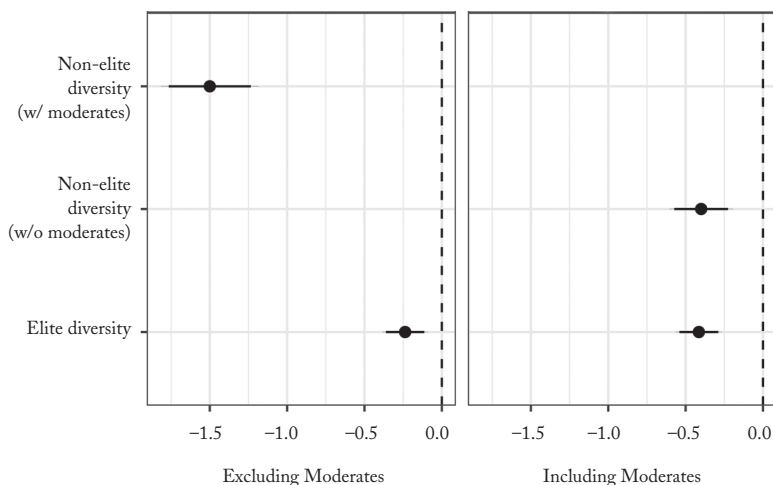
FIGURE 4
NETWORK DIVERSITY AND INTOLERANCE (QUASI-POISSON MODELS)[a]

[a] Figure displays coefficient plots of quasi-Poisson models demonstrating the association between network diversity (with and without moderates) and users' intolerant tweet count. The figure reports nonexponentiated coefficients. The models include controls for users' relevant tweet count, days on Twitter, and political orientation. N = 9,400. The plot shows 95 percent confidence intervals. Full regression tables with and without controls can be found in Appendix B of the supplementary material (Siegel et al. 2021).

mates to exponentiated coefficients to facilitate interpretation, a one-unit increase in elite network diversity from a fully homogeneous to a maximally diverse elite network is associated with a 34 percent decrease in a user's intolerant tweets. Similarly, a one-unit increase in non-elite network diversity (excluding moderates) is associated with a 33 percent decrease in a user's intolerant tweets. But when we measure network diversity including moderates, we see a much stronger relationship between non-elite network diversity and intolerance, with a one-unit increase in non-elite network diversity associated with a 78 percent decrease in intolerance. A larger non-elite network is also significantly associated with lower levels of intolerance. By contrast, a larger elite network is associated with a higher level of intolerance. Demonstrating the validity of our intolerance classifier, we see similar results using only the human-coded data with which we created our training data set (Table B3 of the supplementary material).[77]

[77] We also replicate this analysis looking at tolerant tweets and irrelevant tweets. As we might expect, greater network diversity is significantly associated with higher levels of tolerance. Greater network diversity is also significantly associated with producing more irrelevant tweets. Perhaps more

Next, we tested hypotheses H1b and H2b, that beyond the general association between network diversity and tolerance, individuals who spend more time in diverse networks will exhibit less intolerance over time. The time series structure of Twitter data gave us temporally granular behavioral measures of both network diversity and tolerance that allowed us to better understand the dynamics of the relationship between them. We measured the association between network diversity and the change in a users' level of intolerance from t_1 to t_2. To do so, we used lagged dependent variable models[78] to assess the degree to which spending an additional year in a diverse network was associated with a decrease in public expression of intolerance in the next six-month period, relative to the six months preceding t_1. As in the previous models, we included controls for network size, time on Twitter, political orientation, and—in the case of the count models—the number of relevant tweets.

The results of these analyses—whether we apply proportional or count measures of intolerance, including and excluding controls— suggest that spending a year in a more diverse network is associated with a decrease in intolerance, relative to the previous six-month period. Figure 5 plots the coefficients from our quasi-Poisson lagged dependent variable model. Both elite and non-elite network diversity (with and without moderates) are associated with a decrease in a user's level of intolerance after spending an additional year in a given Twitter network, and the result is statistically significant. Converting the point estimates to exponentiated coefficients for ease of interpretation, we see that spending an additional year in a maximally diverse elite network is associated with a 42 percent decrease in a user's intolerant tweets in the following six-month period, controlling for prior levels of intolerance. Similarly, spending an additional year in a maximally diverse non-elite network excluding moderates is associated with a 45 percent decrease in a user's intolerant tweets in the following six-month period, and spending an additional year in a diverse non-elite network including moderates is associated with an 82 percent decrease in intolerant tweets. As in the aggregate analysis, following a larger number of elites is significantly associated with an increase in intolerance.[79]

---

tolerant individuals also tweet more about topics that are not politically contentious. These results are reported in Tables B4 and B5 of the supplementary material.

[78] Finkel 1995.

[79] Lagged dependent variable models generally underestimate the coefficients of other independent variables besides the lagged dependent variable. This approach thus yields a conservative test of the association between spending an additional year in a diverse network and a user's level of intolerance in the subsequent period (Jorgenson and Burns 2007). These associations are also evident when examin-
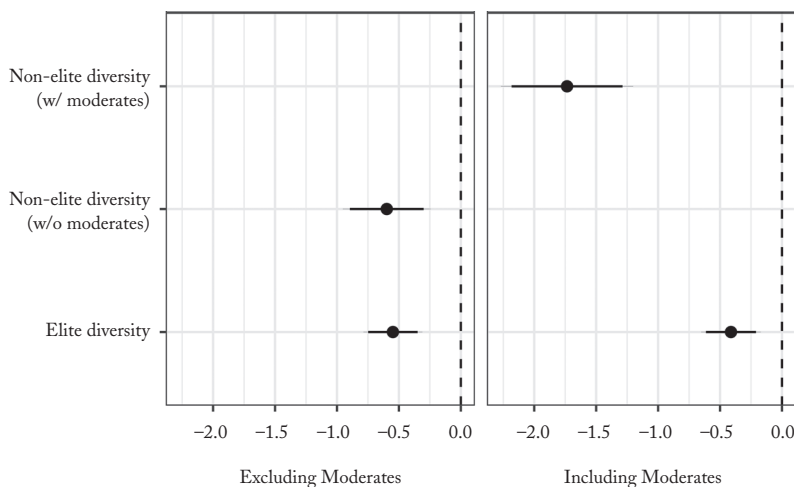
FIGURE 5

NETWORK DIVERSITY AND CHANGE IN INTOLERANCE OVER TIME (QUASI-POISSON LAGGED DEPENDENT VARIABLE MODELS)[a]

[a] Figure displays coefficient plots of lagged dependent variable models evaluating the association between network diversity (with and without moderates) and the change in a user's intolerant tweet count after spending one additional year in a network. The models also include controls for users' relevant tweet count, days on Twitter, and political orientation. N = 7,843. Users in our sample who did not have relevant tweets before May 2015 and after May 2016 were excluded from the analysis. The plot shows 95 percent confidence intervals. Full regression tables with and without controls and models that control for changes in network diversity can be found in the supplementary material (Siegel et al. 2021).

These results provide support for our descriptive hypotheses. As we argue in our discussion of the hypotheses, we can interpret these results causally if changes in Twitter users' levels of intolerance over time are not associated with a simultaneous change in the diversity of their networks and we assume that users' decisions to select into a diverse or homogeneous network at t_1 are independent of their likelihood of bidirectionally changing the frequency with which they tweet intolerant rhetoric a year later at t_2, such that those in more diverse networks express less intolerance and those in more homogeneous networks express more.

The histograms in Figure 6 plot the distribution of users' changes

_____

ing the raw data. The fifty users who exhibited the greatest decrease in intolerance had an average elite network diversity of 0.53 and an average non-elite network diversity of 0.35, whereas those with the greatest increase in intolerance had an average elite network diversity of 0.28 and an average non-elite network diversity of 0.19.
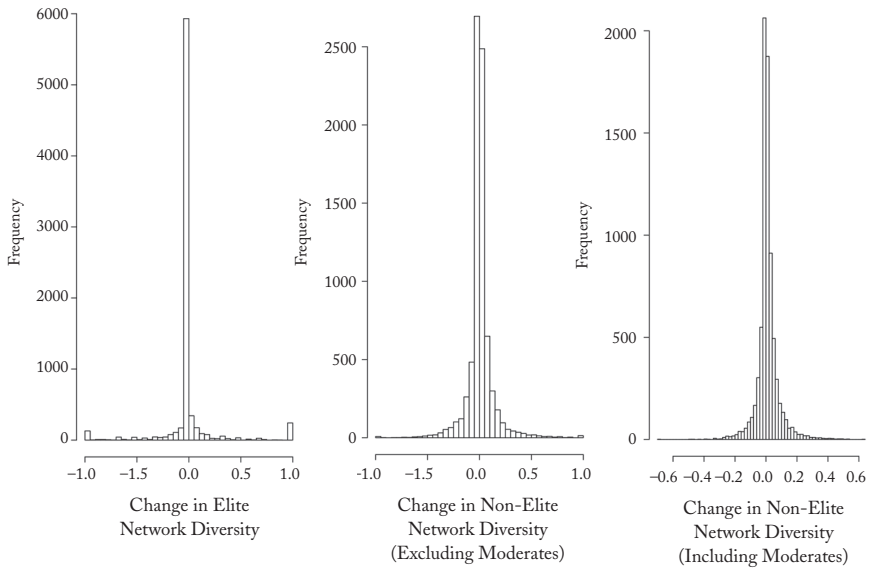
FIGURE 6
CHANGE IN NETWORK DIVERSITY[a]

[a] Figure shows the change in elite and non-elite network diversity (with and without moderates) from May 2015 to May 2016. Most users had minimal change in network diversity in this period, with most of the data clustered around zero.

in network diversity from May 2015 to May 2016, demonstrating that changes in both elite and non-elite network diversity are clustered around zero. Furthermore, the degree to which users in our sample changed the composition of their elite and non-elite networks over the course of our study is not significantly associated with a change in their intolerance levels after spending an additional year on Twitter. This result suggests that our first assumption is valid—the decision to select into a particular network initially is not associated with the likelihood of tolerance changing over the subsequent year.

Of course, we could not randomly assign individuals to initially select into particular networks, which would meet the gold standard for interpreting our results causally. But we could control for cross-sectional variation of initial selection into diverse or homogeneous networks. We reestimated our conditional change models using a fixed effect to mark users in diverse (network diversity greater than 0.5) and non-diverse (network diversity less than 0.5) networks. The results of this analysis are reported in Table B10 of the supplementary material and are consistent with our findings reported in Figure 5, although the effect sizes are

of course much smaller, as they only represent differences among users in diverse (or homogeneous) networks.

Again, without random assignment we cannot rule out the possibility that our findings are biased by omitted variables. We therefore also explored the sensitivity of our results to the inclusion of an omitted variable exogenous to and correlated with both network diversity and a change in intolerance. Sensitivity analysis allows us to assess how strong unobserved confounders would need to be to change our results. The results of this sensitivity analysis, explained in detail in Appendix C of the supplementary material, suggest that it is unlikely that our results are being driven by omitted-variable bias. The effect of elite network diversity and non-elite network diversity on the change in intolerance over time is robust to confounding variables between one and three times as strong as any of the covariates observed in our models. Taken together, we interpret our results as providing support for our descriptive hypotheses and as suggestive evidence that network diversity may cause a reduction in the public expression of intolerance over time.

VI. Discussion and Conclusions

Our findings indicate that individuals in more diverse political communication networks—particularly those that include moderate peers—express lower levels of intolerance toward outgroups than those in homogeneous networks and that the amount of intolerant rhetoric they express decreases after spending more time in the network. By contrast, the amount of intolerance expressed by those in homogeneous networks increases over time. Further, individuals rarely changed the ideological diversity of their networks over the time period under study. Thus, as we assume that the decision to initially select into a given network is independent of a bidirectional change in the expressed level of tolerance in the future—such that those in diverse networks express less intolerance and those in homogeneous networks express more—we interpret our results as providing suggestive evidence that greater exposure to network diversity decreases intolerance over time.

What do our results tell us about the mechanisms by which network diversity might affect tolerance over time? As we outlined above, a diverse elite network may expose individuals to a wider range of viewpoints from well-known, influential political actors. Similarly, spending time in a diverse non-elite network exposes individuals to a wider range of political viewpoints from their peers. As noted previously, decades of social science literature suggest that exposure to members of an out-

group can create greater awareness of the rationale for opposing views.[80] Thus, the associations we observe could be driven by attitude changes from repeated exposure to new information over time.

But spending time in a diverse peer network that contains moderates may influence tolerance over time through a second mechanism: social norms. Recent research suggests that the norms of online communities heavily influence the individuals who belong to them. Experimental evidence demonstrates that being sanctioned by a peer for uncivil or intolerant behavior or being called out for expressing a particular political view can cause individuals to express less extreme and less intolerant rhetoric online.[81] Moreover, individuals in networks where intolerant language is less common are even more affected by social sanctions from peers, who set the bounds of what behavior is acceptable or not within a network.[82] Our finding that diverse networks—particularly those with greater numbers of moderate peers—contain less intolerant speech than more homogeneous networks suggests that using intolerant discourse in these settings may violate their norms. As individuals spend more time in a network where intolerant language is uncommon or is sanctioned by other users, they may reduce their expression of this rhetoric over time.

By exploiting Twitter's architecture to develop a much more detailed measure of elite and non-elite political communication networks than traditional survey methods allow and by describing the relationship between network diversity and tolerance both in the aggregate and over time, this article provides a key contribution to the tolerance and elite cues literatures. Additionally, our analysis offers insight into the potential short- and long-term effects of online social networks on political attitudes and behavior, a topic of great interest in political communication research.

In the vast majority of tolerance studies to date, social scientists have been seriously limited in their ability to comprehensively map and characterize people's networks.[83] The challenge of accurately measuring political communication networks has forced researchers to either rely on self-reported descriptions of individuals' networks or to painstakingly develop detailed pictures of such networks (through anthropological participant observation, for example), which can only be compiled for a small number of people.[84] This article, by contrast, demonstrates the

---

[80] See Mutz 2002; Mutz 2006.
[81] Munger 2017.
[82] Siegel and Badaan 2020.
[83] For discussions of these limitations, see Gibson 2001; Bloom and Bagno-Moldavsky 2015.
[84] Gibson 2001.

tremendous potential of social media data to thoroughly characterize the elite and non-elite political communication networks of large numbers of people in diverse contexts.

Similarly, this article elucidates the opportunity that machine-learning text-analysis techniques afford for using social media data to study shifting behavior over time. By using a relatively small amount of human-coded data to train a classifier, we were able to analyze more than 29 million tweets produced by 9,400 politically engaged Egyptians over two years. Obtaining these types of repeated longitudinal measures of intolerance and networks over time would be impossible using more traditional data sources, particularly in nondemocratic polities.

But before generalizing the findings in this analysis beyond postcoup Egypt, we must consider three caveats. First, this study was conducted using Twitter data because those data are publicly available and allow us to measure the behavior of elites and everyday citizens on the same platform. Twitter users tended to be younger and more politically engaged than users of other social media platforms in Egypt in this period. Moreover, our analysis is further constrained to individuals who regularly talked about politics on Twitter (that is, those who either were geolocated in Egypt or tweeted about Egyptian politics in Arabic and who published at least four tweets relevant to civil liberties from March 2014 to April 2015). Additionally, these individuals did not delete their accounts, have their accounts suspended, or stop tweeting over the course of the study. This means that our analysis is also limited to active Twitter users and likely excludes extremist users whose accounts may have been suspended during the period under study. Although this population is in no way representative of Egyptians or even of Egyptian social media users, we believe it is a reasonable proxy for politically engaged Egyptian Twitter users.

Second, the depth of ideological polarization in postcoup Egypt, exacerbated by the exclusion of Islamists from almost all formal political channels, is extreme. But in other parts of the Arab world, including Jordan, Morocco, and Tunisia, the Islamist-secular divide has also deepened, and we might expect to find similar results.[85] Recent work suggests that polarization and ideological animosity have increased in established democracies, developing democracies, and authoritarian contexts alike over the past decade,[86] perhaps making postcoup Egypt less of an outlier.

Third, our real-world observational data does not enable us to ran-

[85] Kubinec and Owen 2018.
[86] Somer and McCoy 2018.

domly assign individuals to networks to make robust causal inferences. Although our sensitivity analysis suggests that our results are robust to the presence of large confounders in the data, more experimental research is needed to precisely estimate the causal effects of network diversity on tolerance in diverse contexts.

Although evaluating the relationship between network diversity and intolerance in the online sphere may not appear to have immediate consequences on the ground, a growing body of literature suggests that online activity reveals important information about offline attitudes and behavior. Notably, in their study of ideological homophily and segregation on Twitter, Halberstam and Knight[87] find that the degree of ideological segregation in Twitter networks is similar to that in networks of face-to-face interactions with friends and coworkers. This suggests that the effects we observe in the online sphere could be occurring offline as well. The connections between online and offline behavior also deserve further study.

Given that the "Twitter revolutions" failed to bring democratic change and that virulent Islamist-secular polarization continues to threaten Egypt's stability, understanding the structure of political communication among elites and engaged citizens has become particularly pertinent. The results of this article indicate that the online social networks that Egyptian Twitter users cultivate may shape their level of intolerance—or at least their propensity to use intolerant rhetoric—over time. More substantively, they suggest that online social network structures may influence political attitudes and behavior. We hope future research will continue to examine the dynamics of online network diversity and political tolerance in other contexts, paying particular attention to how the patterns we observe in postcoup Egypt might differ in less sharply polarized contexts or other regime types.

SUPPLEMENTARY MATERIAL

Supplementary material for this article can be found at https://doi.org/10.1017/ S0043887120000295.

DATA

Replication files for this article can be found at https://doi.org/10.7910/DVN/KO IKBJ.

REFERENCES

[87] Halberstam and Knight 2016.

Aday, Sean, Henry Farrell, Marc Lynch, John Sides, John Kelly, and Ethan Zuckerman. 2010. "Blogs and Bullets: New Media in Contentious Politics." Washington, D.C.: United States Institute of Peace. At https://www.usip.org/publications/2010/09/blogs-and-bullets-new-media-contentious-politics, accessed February 11, 2021.

Aday, Sean, Henry Farrell, Marc Lynch, John Sides, and Deen Freelon. 2012. "Blogs and Bullets II: New Media and Conflict after the Arab Spring." Washington, D.C.: United States Institute of Peace. At https://www.usip.org/publications/2012/07/blogs-and-bullets-ii-new-media-and-conflict-after-arab-spring, accessed February 11, 2021.

Allport, Gordon Willard. 1954. *The Nature of Prejudice*. Cambridge, Mass.: Addison-Wesley.

Anduiza, Eva, Michael J. Jensen, and Laia Jorba. 2012. *Digital Media and Political Engagement Worldwide: A Comparative Study*. New York, N.Y.: Cambridge University Press.

Bakshy, Eytan, Solomon Messing, and Lada A. Adamic. 2015. "Exposure to Ideologically Diverse News and Opinion on Facebook." *Science*, June 5: 1130–32. At https://doi.org/10.1126/science.aaa1160.

Barberá, Pablo. 2015. "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data." *Political Analysis* 23, no. 1: 76–91. At https://doi.org/10.1093/pan/mpu011.

Barberá, Pablo, John T. Jost, Jonathan Nagler, Joshua A. Tucker, and Richard Bonneau. 2015. "Tweeting from Left to Right: Is Online Political Communication More Than an Echo Chamber?" *Psychological Science* 26, no. 10: 1531–42. At https://doi.org/10.1177%2F0956797615594620.

Baumer, Eric P., Steven F. Messner, and Richard Rosenfeld. 2003. "Explaining Spatial Variation in Support for Capital Punishment: A Multilevel Analysis." *American Journal of Sociology* 108, no. 4: 844–75. At https://doi.org/10.1086/367921.

Blau, Peter M. 1974. "Presidential Address: Parameters of Social Structure." *American Sociological Review* 39, no. 5: 615–35. At https://doi.org/10.2307/2094309.

Bloom, Pazit Ben-Nun, and Olena Bagno-Moldavsky. 2015. "The Conditional Effect of Network Diversity and Values on Tolerance." *Political Behavior* 37, no. 3: 623–51. At https://doi.org/10.1007/s11109-014-9284-2.

Boukhars, Anouar, Nathan J. Brown, Michele Dunne, Raphaël Lefèvre, Marwan Muasher, Frederic Wehrey, Katherine Wilkens, and Scott Williamson. 2014. "The Egypt Effect: Sharpened Tensions, Reshuffled Alliances." Washington, D.C.: Carnegie Endowment for International Peace. At https://carnegieendowment.org/2014/02/13/egypt-effect-sharpened-tensions-reshuffled-alliances-pub-54515, accessed February 11, 2021.

Brader, Ted, and Joshua A. Tucker. 2008. "Pathways to Partisanship: Evidence from Russia." *Post-Soviet Affairs* 24, no. 3: 263–300. At https://doi.org/10.2747/1060-586X.24.3.263.

Budak, Ceren, and Duncan J. Watts. 2015. "Dissecting the Spirit of Gezi: Influence vs. Selection in the Occupy Gezi movement." *Sociological Science* 2: 370–97. At http://doi.org/10.15195/v2.a18.

Calhoun, Craig. 1988. "Populist Politics, Communications Media and Large

Scale Societal Integration." *Sociological Theory* 6, no. 2: 219–41. At https://doi .org/10.2307/202117.

Druckman, James N. 2001. "The Implications of Framing Effects for Citizen Competence." *Political Behavior* 23, no. 3: 225–56. At https://www.jstor.org /stable/1558384.

Druckman, James N., Erik Peterson, and Rune Slothuus. 2013. "How Elite Partisan Polarization Affects Public Opinion Formation." *American Political Science Review* 107, no. 1: 57–79. At https://doi.org/10.1017/S0003055412000500.

Dunne, Michele, and Scott Williamson. 2014. "Egypt, Counterterrorism, and the Politics of Alienation." Washington, D.C.: Carnegie Endowment for International Peace. At https://carnegieendowment.org/2014/08/20/egypt-counter terrorism-and-politics-of-alienation-pub-56428, accessed February 11, 2021.

Dyck, Joshua J., and Shanna Pearson-Merkowitz. 2014. "To Know You Is Not Necessarily to Love You: The Partisan Mediators of Intergroup Contact." *Political Behavior* 36, no. 3: 553–80. At https://doi.org/10.1007/s11109-013-92 40-6.

El-Khalili, Sara. 2013. "Social Media as a Government Propaganda Tool in Post-Revolutionary Egypt." *First Monday* 18, no. 3. At https://doi.org/10.5210/fm .v18i3.4620.

Farrell, Henry. 2012. "The Consequences of the Internet for Politics." *Annual Review of Political Science* 15: 35–52. At https://doi.org/10.1146/annurev -polisci-030810-110815.

Finkel, Steven E. 1995. *Causal Analysis with Panel Data*. Thousand Oaks, Calif.: Sage.

Garrett, R. Kelly. 2009. "Politically Motivated Reinforcement Seeking: Reframing the Selective Exposure Debate." *Journal of Communication* 59, no. 4: 676–99. At https://doi.org/10.1111/j.1460-2466.2009.01452.x.

Gibson, James L. 2001. "Social Networks, Civil Society, and the Prospects for Consolidating Russia's Democratic Transition." *American Journal of Political Science* 45, no. 1: 51–68. At https://doi.org/10.2307/2669359.

Gibson, James L. 2013. "Measuring Political Tolerance and General Support for Pro–Civil Liberties Policies: Notes, Evidence, and Cautions." *Public Opinion Quarterly* 77, no. S1: 45–68. At https://doi.org/10.1093/poq/nfs073.

Gibson, James L., and Amanda Gouws. 2005. *Overcoming Intolerance in South Africa: Experiments in Democratic Persuasion*. New York, N.Y.: Cambridge University Press.

Haas, Ingrid Johnsen, and William A. Cunningham. 2014. "The Uncertainty Paradox: Perceived Threat Moderates the Effect of Uncertainty on Political Tolerance." *Political Psychology* 35, no. 2: 291–302. At https://doi.org/10.1111 /pops.12035.

Halberstam, Yosh, and Brian Knight. 2016. "Homophily, Group Size, and the Diffusion of Political Information in Social Networks: Evidence from Twitter." *Journal of Public Economics* 143, November: 73–88. At https://doi.org/10.1016/j.jpu beco.2016.08.011.

Halpern, Daniel, Sebastián Valenzuela, and James E. Katz. 2017. "We Face, I Tweet: How Different Social Media Influence Political Participation through Collective and Internal Efficacy." *Journal of Computer-Mediated Communication* 22, no. 6: 320–36. At https://doi.org/10.1111/jcc4.12198.

Hamzawy, Amr. 2017. "The Tragedy of Egypt's Stolen Revolution." Washington,

D.C.: Carnegie Endowment for International Peace. At https://carnegieen-dowment.org/2017/01/25/tragedy-of-egypt-s-stolen-revolution-pub-67809, accessed February 11, 2021.

Howard, Philip N., and Muzammil M. Hussain. 2013. *Democracy's Fourth Wave? Digital Media and the Arab Spring*. New York, N.Y.: Oxford University Press on Demand.

Huckfeldt, Robert, Jeanette Morehouse Mendez, and Tracy Osborn. 2004. "Disagreement, Ambivalence, and Engagement: The Political Consequences of Heterogeneous Networks." *Political Psychology* 25, no. 1: 65–95. At https://doi.org/10.1111/j.1467-9221.2004.00357.x.

Ikeda, Ken'ichi, and Sean Richey. 2009. "The Impact of Diversity in Informal Social Networks on Tolerance in Japan." *British Journal of Political Science* 39, no. 3: 655–68. At https://doi.org/10.1017/S0007123408000574.

Jebril, Nael, Václav Stetka, and Matthew Loveless. 2013. *Media and Democratisation: What Is Known about the Role of Mass Media in Transitions to Democracy*. Oxford, UK: Oxford University, Reuters Institute for the Study of Journalism.

Jorgenson, Andrew K., and Thomas J. Burns. 2007. "The Political-Economic Causes of Change in the Ecological Footprints of Nations, 1991–2001: A Quantitative Investigation." *Social Science Research* 36, no. 2: 834–53. At https://doi.org/10.1016/j.ssresearch.2006.06.003.

King, Ryan D., and Darren Wheelock. 2007. "Group Threat and Social Control: Race, Perceptions of Minorities and the Desire to Punish." *Social Forces* 85, no. 3: 1255–80. At https://www.jstor.org/stable/4494972.

Kubinec, Robert, and John Owen. 2018. "When Groups Fall Apart: A Study of Transnational Polarization with Twitter from the Arab Uprisings." New York University, Abu Dhabi, and University of Virginia. At https://doi.org/10.31235/osf.io/wykmj.

Liu, Yabing, Chloe Kliman-Silver, and Alan Mislove. 2014. "The Tweets They Are a-Changin': Evolution of Twitter Users and Behavior." International AAAI Conference on Web and Social Media, Ann Arbor, Mich., June 1–4.

Lupia, Arthur. 1994. "Shortcuts versus Encyclopedias: Information and Voting Behavior in California Insurance Reform Elections." *American Political Science Review* 88, no. 1: 63–76. At https://doi.org/10.2307/2944882.

Lupia, Arthur, and Mathew D. McCubbins. 1998. *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* New York, N.Y.: Cambridge University Press.

Lynch, Marc. 2006. "Public Opinion Survey Research and Public Diplomacy." In Joshua S. Fouts, ed., *Public Diplomacy Practitioners, Policy Makers, and Public Opinion*. Rancho Mirage, Calif.: Annenberg Foundation Trust at Sunnylands, USC Annenberg School for Communication, USC Center on Public Diplomacy, and the Pew Research Center.

Lynch, Marc, Deen Freelon, and Sean Aday. 2016. "Blogs and Bullets IV: How Social Media Undermines Transitions to Democracy." Washington, D.C.: PeaceTech Lab. At https://www.peacetechlab.org/blogs-and-bullets, accessed February 11, 2021.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Annual Review of Sociology* 27, August: 415–44. At https://doi.org/10.1146/annurev.soc.27.1.415.

Morrow, Adam, and Khalid Moussa al-Omrani. 2014. "In Post-Revolution

Egypt, Social Media Shows Dark Side." *IPS News*, May 15. At http://www.ip
snews.net/2013/05/in-post-revolution-egypt-social-media-shows-dark-side/,
accessed February 11, 2021.

Mourtada, Racha. 2014. "Citizen Engagement and Public Services in the Arab
World: The Potential of Social Media." Arab Social Media Report 6. Dubai,
UAE: Mohammed Bin Rashid School of Government. At https://www.mbrsg
.ae/getattachment/e9ea2ac8-13dd-4cd7-9104-b8f1f405cab3/Citizen-En
gagement-and-Public-Services-in-the-Arab.aspx, accessed February 17, 2021.

Munger, Kevin. 2017. "Tweetment Effects on the Tweeted: Experimentally Re-
ducing Racist Harassment." *Political Behavior* 39, no. 3: 629–49. At https://doi
.org/10.1007/s11109-016-9373-5.

Mutz, Diana C. 2002. "Cross-Cutting Social Networks: Testing Democratic
Theory in Practice." *American Political Science Review* 96, no. 1: 111–26. At
https://doi.org/10.1017/S0003055402004264.

Mutz, Diana C. 2006. *Hearing the Other Side: Deliberative versus Participatory De-
mocracy*. New York, N.Y.: Cambridge University Press.

Mutz, Diana C., and Paul S. Martin. 2001. "Facilitating Communication across
Lines of Political Difference: The Role of Mass Media." *American Political
Science Review* 95, no. 1: 97–114. At https://doi.org/10.1017/S000305540100
0223.

Nunn, Clyde Z., Harry J. Crockett, and J. Allen Williams. 1978. *Tolerance for
Nonconformity: A National Survey of Americans' Changing Commitment to Civil
Liberties*. San Francisco, Calif.: Jossey-Bass.

Peffley, Mark, Pia Knigge, and Jon Hurwitz. 2001. "A Multiple Values Model of
Political Tolerance." *Political Research Quarterly* 54, no. 2: 379–406. At https://
doi.org/10.1177%2F106591290105400207.

Petersen, Michael, Rune Slothuus, Rune Stubager, and Lise Togeby. 2011. "Free-
dom for All? The Strength and Limits of Political Tolerance." *British Journal
of Political Science* 41, no. 3: 581–97. At https://doi.org/10.1017/S0007123410
000451.

Pettigrew, Thomas F. 1998. "Intergroup Contact Theory." *Annual Review of Psy-
chology* 49: 65–85. At https://doi.org/10.1146/annurev.psych.49.1.65.

Robinson, Carin. 2010. "Cross-Cutting Messages and Political Tolerance: An Ex-
periment Using Evangelical Protestants." *Political Behavior* 32, no. 4: 495–515.
At https://doi.org/10.1007/s11109-010-9118-9.

Roccas, Sonia, and Adi Amit. 2011. "Group Heterogeneity and Tolerance: The
Moderating Role of Conservation Values." *Journal of Experimental Social Psy-
chology* 47, no. 5: 898–907. At https://doi.org/10.1016/j.jesp.2011.03.011.

Rollins, Tom. 2014. "Egyptian TV Enflames Divisive Politics." *al-Monitor*, May
7. At https://www.al-monitor.com/pulse/originals/2014/05/egypt-media-sisi
-military-journalism-incite.html, accessed February 11, 2021.

Salem, Fadi. 2017. "Social Media and the Internet of Things: Towards Data-
Driven Policymaking in the Arab World: Potential, Limits and Concerns."
Arab Social Media Report 7. Dubai, UAE: Mohammed Bin Rashid School of
Government. At https://www.mbrsg.ae/getattachment/1383b88a-6eb9-476a
-bae4-61903688099b/Arab-Social-Media-Report-2017, accessed February 11,
2021.

Siegel, Alexandra A. 2021. "Replication Files for: Tweeting beyond Tahrir: Ide-

ological Diversity and Political Intolerance in Egyptian Twitter Networks." Harvard Dataverse, V1. At https://doi.org/10.7910/DVN/KOIKBJ.

Siegel, Alexandra A., and Vivienne Badaan. 2020. "#No2Sectarianism: Experimental Approaches to Reducing Sectarian Hate Speech Online." *American Political Science Review* 114, no. 3: 837–55. At https://doi.org/10.1017/S0003055420000283.

Siegel, Alexandra A., Jonathan Nagler, Richard Bonneau, and Joshua A. Tucker. 2021. "Supplementary material for Tweeting beyond Tahrir: Ideological Diversity and Political Intolerance in Egyptian Twitter Networks." At https://doi.org/10.1017/S0043887120000295.

Smith, Andrew N., Eileen Fischer, and Chen Yongjian. 2012. "How Does Brand-Related User-Generated Content Differ across YouTube, Facebook, and Twitter?" *Journal of Interactive Marketing* 26, no. 2: 102–13. At https://doi.org/10.1016/j.intmar.2012.01.002.

Somer, Murat, and Jennifer McCoy. 2018. "Déjà vu? Polarization and Endangered Democracies in the 21st Century." *American Behavioral Scientist* 62, no. 1: 3–15. At https://doi.org/10.1177%2F0002764218760371.

Stein, Elizabeth A. 2013. "The Unraveling of Support for Authoritarianism: The Dynamic Relationship of Media, Elites, and Public Opinion in Brazil, 1972–82." *International Journal of Press/Politics* 18 no. 1: 85–107. At https://doi.org/10.1177%2F1940161212460762.

Stukal, Denis, Sergey Sanovich, Richard Bonneau, and Joshua A. Tucker. 2017. "Detecting Bots on Russian Political Twitter." *Big Data* 5, no. 4: 310–24. At https://doi.org/10.1089/big.2017.0038.

Sullivan, John L., and Henriët Hendriks. 2009. "Public Support for Civil Liberties Pre- and Post-9/11." *Annual Review of Law and Social Science* 5: 375–91. At https://doi.org/10.1146/annurev.lawsocsci.093008.131525.

Sullivan, John L., James Piereson, and George E. Marcus. 1979. "An Alternative Conceptualization of Political Tolerance: Illusory Increases 1950s–1970s." *American Political Science Review* 73, no. 3: 781–94. At https://doi.org/10.2307/1955404.

Tadmor, Carmit T., Philip E. Tetlock, and Kaiping Peng. 2009. "Acculturation Strategies and Integrative Complexity: The Cognitive Implications of Biculturalism." *Journal of Cross-Cultural Psychology* 40, no. 1: 105–39. At https://doi.org/10.1177%2F0022022108326279.

Tufekci, Zeynep, and Christopher Wilson. 2012. "Social Media and the Decision to Participate in Political Protest: Observations from Tahrir Square." *Journal of Communication* 62, no. 2: 363–79. At https://doi.org/10.1111/j.1460-2466.2012.01629.x.

Weber, Ingmar, Venkata Garimella, and Alaa Batayneh. 2013. "Secular vs. Islamist Polarization in Egypt on Twitter." In *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. New York, N.Y.: Association for Computing Machinery: 290–97. At https://doi.org/10.1145/2492517.2492557.

Wojcieszak, Magdalena E., and Diana C. Mutz. 2009. "Online Groups and Political Discourse: Do Online Discussion Spaces Facilitate Exposure to Political Disagreement?" *Journal of Communication* 59, no. 1: 40–56. At https://doi.org/10.1111/j.1460-2466.2008.01403.x.

Zhang, Xue, Hauke Fuehres, and Peter A. Gloor. 2011. "Predicting Stock Market Indicators through Twitter: 'I Hope It Is Not as Bad as I Fear.'" *Procedia: Social and Behavioral Sciences* 26: 55–62. At https://doi.org/10.1016/j.sbspro.2011.10.562.

Authors

ALEXANDRA A. SIEGEL is an assistant professor of political science at the University of Colorado Boulder. Her research interests include comparative political behavior, Middle East politics, social media and politics, networks, and text as data. She can be reached at alexandra.siegel@colorado.edu.

JONATHAN NAGLER is a professor of politics at New York University. His research interests include campaigns and elections, political behavior, econometrics, and social media and politics. He is a coauthor, with Jan E. Leighley, of *Who Votes Now? Demographics, Issues, Inequality, and Turnout in the United States* (2014). He can be reached at jonathan.nagler@nyu.edu.

RICHARD BONNEAU is a professor of biology and computer science at New York University. His research interests include systems biology and protein modeling, and social media and politics. He can be reached at richard.bonneau@nyu.edu.

JOSHUA A. TUCKER is a professor of politics at New York University. His research interests include exploring the effects of social media on political participation and politics, as well as developing methods for using social media data to study politics in new ways. He is a coauthor, with Grigore Pop-Eleches, of *Communism's Shadow: Historical Legacies and Contemporary Political Attitudes* (2017), and the author of *Regional Economic Voting: Russia, Poland, Hungary, Slovakia, and the Czech Republic, 1990–1999* (2006). He can be reached at joshua.tucker@nyu.edu.

Acknowledgements

Funding

Key Words

ELECTORAL MANIPULATION AND REGIME SUPPORT
Survey Evidence from Russia

By ORA JOHN REUTER and DAVID SZAKONYI

ABSTRACT

Does electoral fraud stabilize authoritarian rule or undermine it? The answer to this question rests in part on how voters evaluate regime candidates who engage in fraud. Using a survey experiment conducted after the 2016 elections in Russia, the authors find that voters withdraw their support from ruling party candidates who commit electoral fraud. This effect is especially large among strong supporters of the regime. Core regime supporters are more likely to have ex ante beliefs that elections are free and fair. Revealing that fraud has occurred significantly reduces their propensity to support the regime. The authors' findings illustrate that fraud is costly for autocrats not just because it may ignite protest, but also because it can undermine the regime's core base of electoral support. Because many of its strongest supporters expect free and fair elections, the regime has strong incentives to conceal or otherwise limit its use of electoral fraud.

Does electoral fraud stabilize authoritarian rule or undermine it? On the one hand, electoral fraud may help the regime "win" elections and signal strength to elites.[1] This view suggests that manufacturing dominant electoral victories deters potential challengers. On the other hand, electoral fraud also carries a clear set of risks. Fraud can serve as a focal point around which the opposition can organize mass protests, as the color revolutions clearly demonstrate.[2]

But fraud holds another liability for autocrats, one that is underappreciated: it can undermine popular support for the authorities, even among those who back the regime. In this article, we examine how voters in contemporary Russia respond when they find out that the regime is manipulating elections. We argue that because voters view fraud as morally inappropriate, they disapprove of its use and withdraw support from the candidates who use it.

The effects of increasing awareness of electoral fraud are largest among core regime supporters. In electoral authoritarian regimes, regime partisans are more likely to believe ex ante that elections are con-

[1] Simpser 2013; Rozenas 2016.
[2] Tucker 2007; Bunce and Wolchik 2011.

ducted fairly. This can happen for a number of reasons. Regime support-
ers are more exposed (and possibly susceptible) to regime propaganda,
and partisanship biases may inhibit the internalization of rumors about
fraud. Alternatively, those who support the regime may do so precisely
because they believe it's holding free and fair elections. Given these pre-
conceived notions, core supporters will be most likely to punish regime
incumbents when fraud is revealed to them. By contrast, swing voters
or those who are weakly aligned are already skeptical about electoral in-
tegrity, so the revelation of fraud will do less to affect their vote choice.
Expectations of electoral fraud are already factored in for these voters.

To test these claims, we conducted a framing experiment through the
2016 Russian Election Study (RES), a nationally representative survey
following that year's State Duma election. Our survey experiment ran-
domly prompted respondents to evaluate a hypothetical United Rus-
sia (UR) candidate who was known to have engaged in different types
of fraud, and then asked them to rate their likelihood of voting for the
candidate. We find that all types of electoral fraud—ballot-box fraud,
vote buying, and intimidation—reduce support for the UR candidate.

Using data from the same survey, we report several other findings
that reinforce our main arguments. First, the vast majority of Russians
express moral disapproval of electoral fraud, regardless of their affinity
for the regime in power. Second, a surprisingly large share of Russians
believe that elections are held honestly and—more important for this
study—regime supporters are much more likely to believe that elec-
tions are free and fair. Third, we find that learning about UR candidates'
use of fraud produces a much larger reduction in support among strong
regime backers than it does among weakly aligned voters. We con-
clude that if information on fraud were to become widespread in Rus-
sia, Vladimir Putin's electoral coalition would diminish significantly in
size. We replicate these findings with a second survey experiment con-
ducted in Russia in May 2018, which also examines how an individu-
al's likelihood of voting depends on perceptions of fraud.

Our findings demonstrate that excessive use of fraud can destabilize
autocracy not just because it leads to mass protest, but also because it
erodes the regime's electoral base. Some recent accounts suggest the op-
posite. For instance, Milan Svolik argues that regime supporters in po-
larized societies will endorse illiberal acts if doing so helps their party
defeat the opposition.[3] Our experiments suggest this is not the case in
Russia: polarization isn't strong enough that regime supporters are will-

[3] Svolik 2020.

ing to excuse regime candidates for fraud. Instead, they punish them for it.

Whereas many neo-institutional accounts of autocracy suggest that regimes should publicize fraud to convey strength, our argument helps to explain why autocrats actually go to great lengths to conceal their use of fraud. Indeed, contemporary electoral autocracies like Russia often commit significant resources to improve public perceptions of electoral integrity. More generally, our findings suggest that autocratic regimes maintain a facade of electoral democracy because many voters believe in that facade and express support for free elections. The neo-institutional literature on autocracy has also usefully pointed out that elections can provide dictators with important instrumental benefits, such as information and co-optation. But our findings suggest that scholars of autocracy shouldn't overlook the more prosaic reasons why dictators retain (or introduce) elections. Elections are held simply because voters value them and expect them to be free and fair.

LITERATURE REVIEW

Autocrats turn to electoral manipulation for a number of reasons. Obviously, such tactics can help the regime win elections. Ballot-box fraud adds votes in a straightforward manner, and some studies show that vote buying can also be effective.[4] And although there's less research on intimidation, at least one recent study finds that threatening voters can be effective for turning out the vote.[5] Fraud may have other benefits as well, such as allowing the regime to manufacture large vote margins that convey an image of strength.[6] Fraud can signal to potential challengers that resistance is futile. To regime insiders, it demonstrates that defection will not be rewarded with success. And some scholars argue that fraud can make opposition voters believe their vote is useless, and thereby reduce turnout among these voters.[7]

But fraud is no electoral panacea for autocrats. A major contribution of the new literature on electoral authoritarianism is to point out that these regimes actually use electoral manipulation sparingly. Simply faking an election is rare.[8] Instead, autocratic regimes put considerable effort into securing electoral victories that reflect the revealed preferences of voters. Genuine victories are preferable to manufactured ones

[4] Cantú and García-Ponce 2015; Vicente 2014.
[5] Frye, Reuter, and Szakonyi 2019.
[6] Simpser 2013.
[7] McCann and Domínguez 1998; Simpser 2012.
[8] Magaloni 2006; Levitsky and Way 2010.

because electoral manipulation is costly. Administrative costs are one factor—it's expensive to coordinate and implement nationwide fraud—but most accounts imply that the more important downside of fraud is that the masses react negatively to it. Indeed, that autocrats usually try to hide fraud indicates that they believe they would suffer some consequence for committing it openly.

It has been argued that electoral fraud can undermine the legitimacy of an autocrat's electoral victory[9] and may be used to convey an image of invincibility.[10] More worrying for autocrats is the number of models that link electoral fraud to the eruption of mass protest.[11] These models assume that the opposition detests electoral manipulation and is able to capitalize on that anger to solve collective action problems, mobilize supporters, and overthrow the incumbents deemed responsible.

But fraud has another potential cost. It may reduce levels of political support for the regime, even among those who back it. In almost all countries, electoral manipulation is illegal: those who commit fraud are breaking the law. Moreover, individual acts of electoral manipulation have moral valence. Voter intimidation involves coercion, which in most cultures is viewed as immoral. To the extent that voters prefer virtuous candidates, they should punish those candidates who use coercion to win votes. Moral evaluations of vote buying are more complicated, but the available evidence indicates that most voters view it as inappropriate.[12] And although the moral calculus of ballot-box fraud hasn't been explored in the literature, it's conceivable that voters find this type of fraud inappropriate if they view it as a form of stealing or cheating.

Thus, there are good reasons to think that incumbents may lose votes if voters discover that they manipulated elections. But few studies examine this notion empirically. On the one hand, work by Eric Kramon suggests that vote buying helps candidates demonstrate competence, trustworthiness, and electoral viability to potential voters in places where patronage is pervasive.[13] On the other hand, Rebecca Weitz-Shapiro uses a survey experiment in Argentina to show that middle-class voters withdraw their support from candidates who engage in vote buying.[14] Using vignette experiments, Roxana Gutiérrez-Romero and Adrienne LeBas show that voters in Kenya are less likely to express sup-

[9] Cornelius 1975; Norris 2014; Birch 2011.
[10] Magaloni 2006.
[11] Tucker 2007; Fearon 2011.
[12] Gonzalez Ocantos, Jonge, and Nickerson 2014.
[13] Kramon 2016.
[14] Weitz-Shapiro 2014.

port for candidates rumored to have engaged in preelection violence.[15]

These studies are relevant for our research, but it's hard to directly compare our findings with theirs. Acts of physical violence—Gutiérrez-Romero and LeBas reference murder in their experiment—hold much greater moral valence than do the types of electoral manipulation we study in this article. Isabela Mares and Lauren Young conducted a survey experiment to examine how rumors of vote buying and intimidation affect support for hypothetical candidates in rural Bulgaria.[16] Their research is the closest to ours but, as we discuss below, their main focus is on how sociodemographic factors affect evaluations of manipulation. In addition, we use an expanded definition of electoral manipulation that includes intimidation, vote buying, and ballot-box fraud.

CONSEQUENCES OF ELECTORAL MANIPULATION: A SURVEY EXPERIMENT IN RUSSIA

Our study's main goal is to examine this question: When autocratic regimes commit fraud, which voters do they risk losing? But before turning to that, we seek to determine whether electoral manipulation affects mass support for the authorities at all. We first address this question by investigating how Russians view the moral appropriateness of different types of electoral manipulation. The 2016 RES, the main data source for our study, included a battery of questions that tapped respondents' views on the acceptability of different types of electoral subversion.[17] The wording of the questions and distribution of responses are shown in Table 1.

These specific acts of manipulation were chosen because they're common in Russian elections. The first row in the table is presented as a baseline. Attending ribbon-cutting ceremonies may be perceived by some as an abuse of state resources, but it's unlikely to elicit a strong negative response from most voters. Indeed, as the table shows, 56 percent of voters consider it mostly acceptable. The next two rows assess the acceptability of two common forms of systemic manipulation: restricting the opposition's access to the media and to the ballot. Voters strongly disapprove of both practices.

The last four rows show how Russians view different forms of election-day manipulation. Not surprisingly, most disapprove of vote buy-

---

[15] Gutiérrez-Romero and LeBas 2020.

[16] Mares and Young 2016.

[17] The 2016 RES was a nationally representative survey of 2,010 respondents from forty-eight regions, carried out between November 8 and December 4, 2016, just after the State Duma elections held that year. The survey was conducted by Levada Center using face-to-face interviews.

TABLE 1
ACCEPTABILITY OF ELECTORAL MANIPULATIONS

Parties and politicians use many strategies to attract votes. In your opinion, how acceptable are the following actions?	*Not Acceptable (%)*			*Completely Acceptable (%)*
	1	*2*	*3*	*4*
1. Attend opening ceremonies for cultural or sporting events during the month before elections	27	17	34	22
2. Limit opposition candidates from appearing on television	75	17	6	2
3. Create obstacles for opposition candidates to register	77	17	5	1
4. Hand out food packets to pensioners	37	23	24	16
5. Recruit people to attend political rallies with liquor or food	67	20	10	4
6. Tell workers of a local firm that they will lose their jobs if they don't vote correctly	82	13	4	2
7. Organize karusels by which buses shuttle people to vote at multiple polling stations	75	13	8	4

ing. But interestingly, they evaluate various types of positive inducements differently. Thirty-seven percent of voters approve of distributing food packets to pensioners, but only 14 percent approve of handing out food or alcoholic drinks at rallies.

What's shown in row six is also no surprise. Most voters (82 percent) strongly disapprove of electoral intimidation. The question in row seven taps voters' assessments of *karusels*, a type of ballot-box fraud.[18] Voters are slightly less disapproving of karusels, but the vast majority (88 percent) do disapprove to some degree. On the whole, voters find all types of electoral subversion—with the possible exception of some types of vote buying—to be unacceptable.

These descriptive statistics are informative. Nondemocratic practices aren't supported by the vast majority of the population. But our primary goal is to determine how awareness of electoral manipulation affects regime support. These questions don't tell us whether voters use the bal-

[18] In Russia, the term *karusel* is used for two slightly different electoral practices. It may refer to simple multiple voting, with groups of voters being transported from poll to poll to vote multiple times, usually using assumed names or absentee certificates. It may also refer to a monitoring scheme for facilitating ballot-box fraud.

lot box to punish the authorities for manipulating elections. Voters may view manipulation as unacceptable, but such considerations may not enter into their voting calculus or may be crowded out by other concerns.

One way to approach our question is to ask respondents about their assessments of electoral manipulation and to correlate such attitudes with regime approval ratings. This correlation is informative—and we explore such analyses below—but it suffers from several limitations. For one, the direction of causality is unclear. Perceptions of electoral integrity might increase support for the regime, or support for the regime might increase the likelihood that voters evaluate regime institutions, such as elections, in a positive light. There are other endogeneity concerns as well. It could be that perceptions of electoral integrity have no effect on regime support, but rather that both attitudes are codetermined by some other factor. Additionally, the correlation doesn't tell us how the revelation of information about electoral manipulation might affect those who think elections are free and fair. Those who think elections are honest may still turn against the regime if they discover that elections are manipulated.

To address such shortcomings, we analyze a survey experiment that we embedded in the 2016 RES survey. The experiment was designed to assess the likelihood that respondents would vote for a hypothetical candidate from the ruling party, United Russia, in the next State Duma election conditional on electoral manipulation by that candidate and his professional background. The experiment had a 3×4 factorial design and each respondent was asked the following:

> Imagine that during the next State Duma elections, a [*candidate professional background treatment here*] is nominated by United Russia in your voting district. He is 50 years old and his program focuses on increasing support for local schools and building new roads in the district. During the campaign, some interesting information emerges about the candidate. On the one hand, it becomes known that he adopted two disabled children from a local orphanage. On the other hand, he [*electoral manipulation treatment here*]. How likely is it that you would vote for this candidate?

Respondents were asked to rate their likelihood of voting for this candidate on a five-point scale ranging from "definitely will not vote" to "definitely will vote." Respondents were randomly assigned to one of twelve combinations of candidate professional background and electoral manipulation, as shown in Table 2. Covariate balance checks presented in Section A of the supplementary material indicate that randomiza-

TABLE 2
EXPERIMENT COVERAGE[a]

	Entrepreneur	Head Doctor	Factory Worker
No electoral manipulation	162	167	153
Gave out presents to voters before the elections	124	136	142
Organized karusels to take voters to polls	142	133	153
Threatened several colleagues so they voted	153	145	160

[a]Total number of respondents who received "no electoral manipulation" (control), 535; total number of respondents who received "any fraud" treatment (three treatments), 1,475.

tion was successful.[19] This type of candidate vignette is broadly similar to that used in several recent experimental studies that vary such attributes as gender and policy positions.[20]

We invoke three professional backgrounds in the first experimental arm: an entrepreneur, a doctor, and a worker (*rabochii*). Our experiment was designed with two purposes: to examine voter assessments of workplace mobilization and to examine how electoral manipulation affects regime support. We are interested in the second question, so here we focus on the parts of the experiment that are relevant to it. We collapse the professional background treatments in the subsequent analyses.

Three types of electoral manipulation are included as treatments. The first refers to vote buying. Although middle-income countries like Russia typically see less vote buying than low-income countries, the practice became well known during the 1990s, and poorer segments of the population still report being offered cash or gifts in exchange for their votes. The second is a treatment that references workplace threats against employees. This is by far the most common type of electoral intimidation in Russia[21] and is likely familiar to respondents. The third is a treatment that refers to ballot-box fraud—specifically, a candidate who organized a multiple-voting scheme using buses to ferry voters to multiple precincts. This type of ballot-box fraud is also common in Russia and is one that respondents could envision candidates organizing. As Table 1 shows, respondents easily interpreted and evaluated all

[19] Reuter and Szakonyi 2021b. All respondents received one of the three professional background treatments. One quarter of the respondents did not receive an electoral manipulation treatment, and they constitute the control group.
[20] Schwarz and Coppock forthcoming; Doherty, Dowling, and Miller 2016.
[21] Frye, Reuter, and Szakonyi 2014.

three treatments in terms of their acceptability during elections, with the latter two getting especially low marks.[22]

Several features of the experiment are worth noting. First, we took care to choose wording that maximizes and equalizes the credibility of the fraud information for all respondents. In the real world, information about fraud—whether it comes via mass media, the internet, or rumor—is often discounted. The extent of this discounting can depend on the respondent's disposition or education and the credibility of the source. For this reason, it would be ill-advised to design an experiment that prompts respondents with a specific news story detailing an actual instance of fraud.

Our experiment is designed to hold constant the credibility of the source by prompting respondents to consider a situation in which information on electoral manipulation is internalized with some degree of certainty. We do this through the formulation "it becomes known." This prompts respondents to think that the information about the candidate committing fraud is already public knowledge, rather than being cued to think about whether the information is accurate or who might be disseminating it.[23] In other words, we invite respondents to consider how they would react in this hypothetical scenario if they *knew* that the fraud had occurred.[24] We return to consider how respondents might accept information about fraud in real life in the conclusion.

To make the vignette more realistic, we focus on a specific candidate and hold constant their partisan affiliation (United Russia) and, as such, their proregime status. Fraud by UR candidates reflects poorly on the party and on Putin, who's closely associated with the party.[25] One reason we focus on proregime candidates is because, as we discuss below, we're particularly interested in how proregime voters react to the revelation of electoral fraud.[26] Including a partisan affiliation also reduces the respondents' need for speculation. A common problem with hypothet-

---

[22] Importantly, the response rate for these questions was very high. The vast majority of respondents recognized each practice and felt comfortable passing judgment.

[23] This approach mirrors other work that uses multiple candidate vignettes; Carey et al. 2020.

[24] The Russian language formulation is "*stanovitsiya izvestno, chto*." See Reuter and Szakonyi 2021b for the exact wording of the question in Russian.

[25] Support for United Russia and support for Putin don't fully overlap, but are very highly correlated (r = 0.54). Our data indicate that support for UR is a sufficient but not necessary condition for supporting Putin. Only 1.5 percent of strong UR supporters (UR support > 7) said they didn't support Putin, but 20 percent of UR opponents (UR support < 3) said they did support him.

[26] It's true that we can't directly extrapolate these results to support for Putin. But we feel confident in asserting that a decline in support for UR and its candidates would be politically problematic and dangerous for the regime. If many voters were to abandon UR in a given election, this would fundamentally undermine the regime.

ical survey prompts is that a large proportion of respondents are unable to speculate about their behavior in an imagined situation. Cueing partisanship along with the adoption trait makes it more likely that a large share of respondents can form an opinion about this baseline candidate.

The full results of the experiment are presented in Figure 1.[27] The y-axis shows the mean response on the vote propensity scale. Differences between professional backgrounds are slight, and aren't the focus of this study. The most important result is the difference in mean vote propensity between the three electoral manipulation treatment groups and the control group. As the figure shows, respondents who were told that the hypothetical candidate engaged in some form of electoral manipulation were significantly less likely to express support for the candidate. This effect holds for all types of electoral manipulation, but there are interesting differences across types. Voters are more turned off by ballot-box fraud than by vote buying and threats.[28] Perhaps it's not surprising that vote buying is less offensive. We explore this finding, as well as variation between the treatment arms, in Section C4 of the supplementary material.[29]

Our main interest, however, is the total effect of electoral manipulation on regime support. Therefore, for all subsequent analyses we collapse the manipulation treatment groups. The difference in means between the control group (the set of bars on the left in Figure 1) and the remaining treatment groups (all other bars) is 0.67 (p = .000), which translates into a 13.4 percent decrease in vote propensity.[30] This is a substantively large effect. Since the vote propensity variable is an ordinal scale, this quantity can't be directly interpreted as a 13 percent decrease in the probability of voting for the candidate. Rather, it makes more sense to evaluate effect sizes across the range of the vote propensity variable. Figure 2 compares the distribution of responses on the five-point vote propensity scale for the two groups. We see a sharp in-

[27] Respondents answered this experiment after evaluating acceptability in Table 1. Though it's possible this affected their responses, we note that more than seventy questions were asked in between. The results were replicated on a later survey that didn't include any questions about acceptability.

[28] The difference between the karusel treatment group and across the vote buying and threat treatment groups is 0.21 and statistically significant.

[29] Reuter and Szakonyi 2021b. Surprisingly, we do not find that several expected demographic characteristics (age, education, employment status, etc.) are correlated with positive views of the electoral manipulations listed in Table 1. In addition, these traits don't appear to mediate the treatment effect, nor do they help to explain why the bought votes treatment leads to a less negative response. In Section C4 of the supplementary material, we discuss possible explanations for these null findings and reiterate the call for more research on why some individuals approve or disapprove of certain electoral manipulations (Szakonyi forthcoming).

[30] The mean response for the control group is 3.35. Across the three manipulation treatment groups, the mean is 2.69.
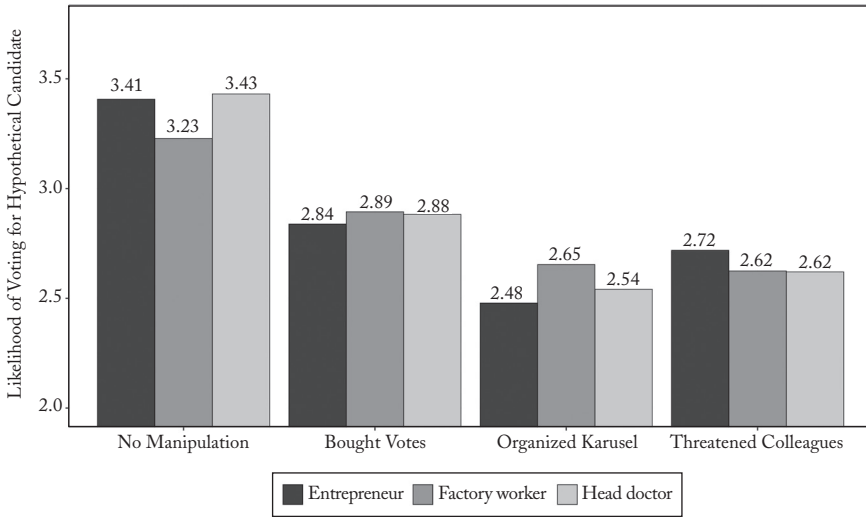
FIGURE 1
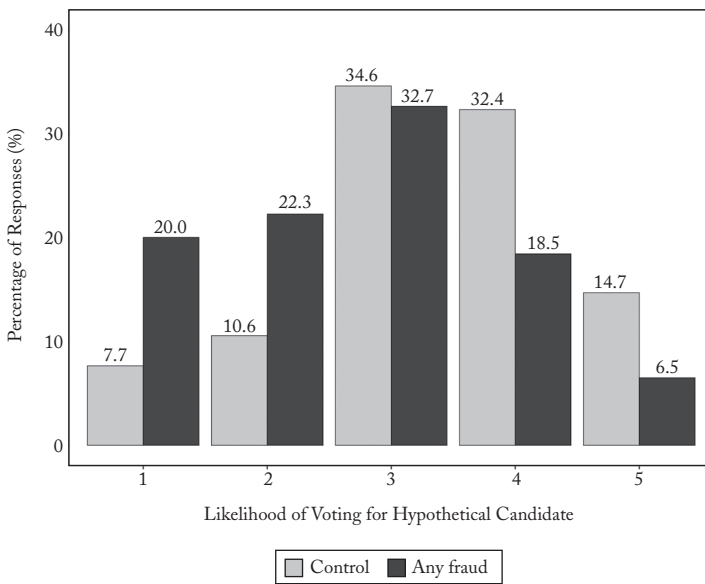SUPPORT FOR HYPOTHETICAL CANDIDATE BY TREATMENT STATUS

FIGURE 2
DISTRIBUTION OF EXPERIMENTAL OUTCOMES BY TREATMENT STATUS

crease in the number of respondents, indicating a very low likelihood of voting for the UR candidate (the values of one and two on the five-point scale) upon receiving any of the manipulation treatments.

MAIN ARGUMENT: ELECTORAL MANIPULATION AND CORE SUPPORTERS

The results described in the section above indicate that information about electoral manipulation committed by proregime candidates reduces support for those candidates. But what types of voters are turned off by electoral manipulation? One conditioning factor that has received little attention is regime affinity. Do so-called swing voters recoil more upon learning about electoral fraud, or would strongly aligned regime supporters be more likely to withdraw their backing? If it's only the swing voters, then electoral manipulation might not be so costly for the regime, since many of these voters wouldn't vote for the regime in any case. But if electoral fraud leads to the loss of core supporters, it could have important consequences for regime stability.

In this section, we argue that strong regime supporters will be just as likely, if not more likely, to punish UR candidates upon learning about electoral manipulation. This will happen if preexisting awareness of fraud varies with regime affinity. If regime partisans have stronger preexisting beliefs that elections are free and fair, they will be more likely to punish incumbents when information about manipulation is revealed. By contrast, if swing or weakly aligned voters are already skeptical about electoral integrity, then the revelation of new information about fraud will do less to affect their vote choice. These voters have already incorporated expectations of significant electoral fraud into their political beliefs, and therefore do not update.

There are a number of reasons to think that on average, strong regime supporters will be less aware of electoral fraud. To the extent that voters disdain electoral fraud—and indeed, the previous section demonstrates that most do—then strong regime supporters might continue to support the regime only because they haven't been exposed to information about electoral manipulation. Fraud is not easy to perceive. It's an illicit activity, and regime officials go to great lengths to cover it up. Regime supporters could be even less attuned to it because they're apolitical or because they're more exposed to proregime media outlets and, therefore, to regime propaganda.[31] Alternatively, strong regime parti-

---

[31] They may self-select proregime media outlets or become regime supporters because of their exposure to such outlets, or both.

sans might be oblivious to fraud for some of the reasons outlined above, including that they have been exposed to rumors in the past, but discounted them because they conflicted with prior notions of the regime's propriety.[32] Indeed, a recent study in Mexico finds evidence of this phenomenon.[33]

The tendencies described above will necessarily be strengthened if propriety is a trait that regime voters value highly. To the extent that regime supporters—or some subset of them—support the ruling party precisely because they perceive it as more trustworthy or honorable than the opposition, they will be more likely to withdraw support when information of malfeasance is revealed. In other words, if new information about fraud erodes a core assumption that they hold about the regime, supporters may punish the regime at the polls.

We can derive two possible hypotheses from these observations. The weak version of the argument suggests that both strongly and weakly aligned regime supporters will withdraw their support from UR candidates when information about electoral manipulation is revealed. The strong version of the argument suggests that strongly aligned regime supporters will be more likely to withdraw their support than will weakly aligned voters. Both arguments contrast with the expectations derived from arguments based on motivated reasoning, as we discuss below.

There are few existing studies on our research question. But a review of work in adjacent literatures suggests there's a strong case to be made that electoral manipulation will only affect vote choice among swing or weakly aligned voters. Strongly aligned voters could be practicing motivated reasoning, and therefore would be more accepting of negative information about United Russia.

Motivated reasoning is a well-established phenomenon in political behavior.[34] One particularly important contributor is partisanship, which in many political settings is as much a determinant of one's worldview as it is a consequence of it.[35] Partisan biases affect public opinion on a huge number of issues, from evaluations of the economy to foreign and public policy,[36] and they can operate via several mechanisms.

One such mechanism is selective exposure to information. Partisans may only seek out information that supports their existing views. This mechanism isn't relevant for our study because our experimental ma-

---

[32] The treatment in our experiments proposes that information on electoral manipulation by the candidate has become widely accepted.

[33] Cantú and García-Ponce 2015.

[34] Kunda 1990; Taber and Lodge 2006.

[35] Campbell et al. 1960; Bartels 2002.

[36] Duch, Palmer, and Anderson 2000; Jerit and Barabas 2012.

nipulation provides subjects with information about fraud. A second possible mechanism is motivated skepticism. Individuals use their reasoning power to downplay or denigrate information that runs counter to their existing beliefs.[37] In our case, motivated skepticism could lead strong UR partisans to discount information about electoral manipulation, perhaps reasoning that the use of fraud is somehow justified or serves a higher purpose.

There's little scholarship on how partisanship affects assessments of electoral manipulation, but related studies suggest that we could expect motivated reasoning to play a role. Several scholars show that voters downplay scandals that afflict leaders of their own party.[38] One recent study from Spain finds that voters are more likely to tolerate corruption if the offending politician is from their own party.[39] Graeme Robertson finds that in Russia, regime supporters are less likely to have knowledge of GOLOS, the country's largest domestic vote-monitoring NGO (sometimes viewed as oppositional), and are less likely to express trust in vote-monitoring organizations.[40] Also, Svolik provocatively argues that political polarization leads voters to tolerate undemocratic policies if it will help their preferred party defeat a detested opponent.[41] In our empirical models below, we seek to adjudicate between this alternative hypothesis and our own. Interestingly, John Carey and colleagues find that partisans in the United States are just as willing to punish copartisans who support undemocratic positions as they are to punish opposition candidates.[42]

Partisanship and Perceptions of Electoral Integrity

The preceding section suggests that the costs of electoral manipulation depend on whether fraud is already common knowledge and who is aware of fraud. We investigate these questions in this section. There are differing views on the integrity of Russian elections. For one, election monitor reports paint a grim picture of opposition candidates being restricted from running, biased media coverage, intimidation, and fraud.[43] Statistical election forensics show a similar picture, demonstrating that ballot-box fraud has become commonplace.[44] All this accords

[37] Taber and Lodge 2006; Lebo and Cassino 2007.
[38] Bhatti, Hansen, and Olsen 2013; Wagner, Tarlov, and Vivyan 2014.
[39] Anduiza, Gallego, and Muñoz 2013.
[40] Robertson 2017.
[41] Svolik 2020.
[42] Carey et al. 2020.
[43] ODIHR 2003; GOLOS 2012; Enikolopov et al. 2013.
[44] Myagkov, Ordeshook, and Shakin 2009; Rundlett and Svolik 2016.

with the Western scholarly consensus that views Russia as an authoritarian regime.

But substantial portions of the Russian electorate hold a much rosier view of how these elections were held. Nationally representative polls find that although citizens detect weaknesses in the electoral process, their perceptions of electoral integrity are generally much more favorable than one might expect from reading the election monitor reports. To demonstrate this, we draw on data from five Russian election surveys held between 2000 and 2016. Each survey included the same set of questions. Respondents were asked whether they agreed with the following statements: elections in Russia are conducted honestly (on a five-point scale, where five indicates that they were completely honest), and voting makes a difference to what happens in the country (also on a five-point scale). Figure 3 plots the averages from respondents following the five national elections.

We see that throughout this period, a large share of the electorate believes that elections are conducted honestly and, to a slightly lesser degree, that electoral outcomes can affect their daily lives (panels [a] and [b]).[45] Interestingly, opinions on these two issues don't shift markedly over time, even as the Russian government took steps to consolidate media ownership in state hands and to limit the ability of opposition parties to contest elections.

The results from the RES polls are by no means unique among work on Russia. Separate opinion polls have found that since 2000, a majority of Russian citizens believe that votes are being counted honestly, media outlets are covering campaigns fairly, and real competition takes place between candidates.[46] Less than 15 percent of respondents felt that electoral results in general could not be trusted.[47]

Russian respondents aren't unique in viewing their elections as relatively clean, even while most outside observers think otherwise. Juliet Pietsch reports that most respondents in Southeast Asian electoral autocracies also believe they're living in democracies.[48] Scott Williamson finds that in most Arab autocracies, a majority of respondents think

---

[45] In 2016, the distribution of responses was as follows: 25 percent responded 5 (honest); 21 percent responded 4; 28 percent responded 3; 14 percent responded 2; and 12 percent responded 1 (dishonest).

[46] McAllister and White 2011; Rose and Mishler 2009.

[47] For comparison, roughly 70 percent of the US electorate were very confident or somewhat confident that their votes were accurately cast and counted in 2004–2016. Justin McCarthy and Jon Clifton, "Update: Americans' Confidence in Voting, Election," Gallup, November 1, 2016. At https://news.gallup.com/poll/196976/update-americans-confidence-voting-election.aspx, accessed November 20, 2020.

[48] Pietsch 2015.
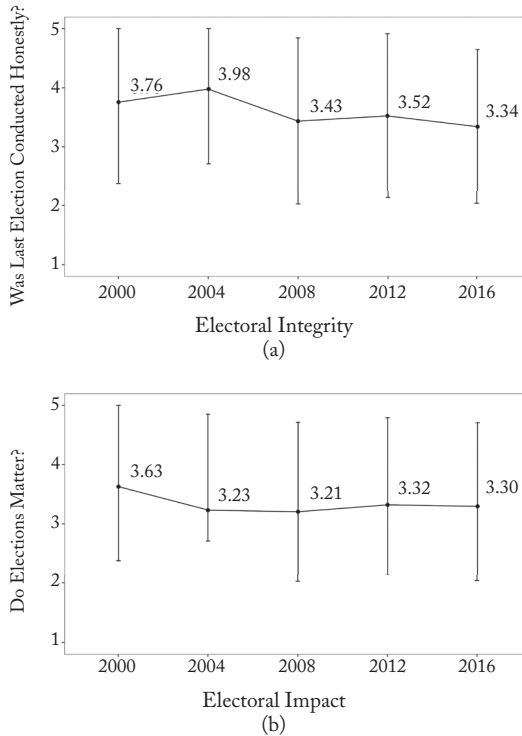
FIGURE 3
BELIEFS ABOUT ELECTORAL INTEGRITY OVER TIME[a]

[a] The panels display the mean agreement among respondents across five RES surveys to the follow-ing statements: (a) elections in Russia are conducted honestly (on a five-point scale, where 5 indicates that they were completely honest); and (b) voting makes a difference to what happens in the country (also on a five-point scale). The error bars show the distribution within one standard deviation above and below the mean.

their elections are free and fair.[49] Table F1 of the supplementary ma-terial presents summary statistics from the latest wave of the World Values Survey (2010–2014) about how respondents living in electoral autocracies (top panel) view the state of elections in their countries.[50] We see that even in regimes generally considered to be unfree, such as Jordan, Singapore, and Zimbabwe, substantial portions of the popula-tion believe that their elections are clean, often to the same degree as in more established democracies.

Digging deeper into the 2016 RES survey results, we find that re-gime supporters are much more likely to have positive perceptions of

[49] Williamson 2018.
[50] Reuter and Szakonyi 2021b. All tables beginning with the letters A–F can be found in the sup-plementary material.

electoral integrity. In Table 3, columns 1 and 2 present multivariate regressions in which the outcome variables are the same measures of election integrity perceptions as those discussed above. Interestingly, basic demographics—gender, age, employment status, and economic situation—explain little of the variation in how people view the quality of elections.[51] What matters most is people's political leanings. Respondents who approve of Putin's performance in office or who support the ruling party United Russia are significantly more likely to believe that elections were held fairly (column 1) and that voting in elections can influence political events in the country (column 2).[52]

These patterns aren't specific to the 2016 Duma elections. In Table B5 of the supplementary material, we show that even during the 2011–2012 electoral cycle, when information about electoral fraud was more widespread across Russian media, regime supporters still held substantially more-positive views of electoral integrity than those who didn't support the regime. In fact, partisanship is the largest predictor of whether respondents viewed the elections as free and fair.

Moreover, as shown in tables B2 and B3, regime supporters are not only more likely to believe that elections are free and fair, but they're also less likely to think that there were electoral violations during either the 2011 or 2016 parliamentary elections. Table B4 also shows that supporters were less likely to have heard of GOLOS. Core UR supporters appear less aware of fraud, and may update more strongly when they find out that fraud has occurred.

But regime supporters may also define electoral integrity differently, and this could be what's driving the correlation between partisanship and views of electoral integrity. By holding electoral processes to a lower standard, supporters may be more likely to believe that the government is administering elections adequately and that a more minimal definition of integrity is being met. We explore this possibility in columns 4–6 of Table 3, which investigates whether respondents believe that common electoral violations are broadly acceptable (the exact statement wordings are shown in Table 1). We find no evidence that support for Putin or the ruling party United Russia is associated with holding a different concept of what electoral integrity means.[53]

[51] For the exact question wordings, please refer to Reuter and Szakonyi 2021b.

[52] These findings hold when either "United Russia support" or "Putin support" are entered into the regression individually. In the supplementary material, we also show similar results using an indicator for whether respondents believe that Russia is a democracy; Reuter and Szakonyi 2021b.

[53] In tables B7 and B8, we also show that regime supporters and opponents who think that violations took place were less likely to view elections as free and fair. We find no difference between regime and opposition supporters in this regard.

TABLE 3
REGIME SUPPORT AND VIEWS OF ELECTORAL INTEGRITY[a]

	Perceptions of Electoral Integrity		Acceptability of Fraud		
	Electoral Integrity 1	*Electoral Impact* 2	*Opp. Blocked* 3	*Karusel Voting* 4	*Media Restrictions* 5
Male	−0.072	−0.113	0.019	−0.022	0.014
	(0.068)	(0.070)	(0.029)	(0.039)	(0.032)
Age (log)	−0.112	−0.125	0.016	−0.097	−0.025
	(0.098)	(0.085)	(0.037)	(0.067)	(0.050)
Education	−0.040	−0.047	−0.015	0.0004	−0.027*
	(0.027)	(0.031)	(0.012)	(0.015)	(0.015)
Town size	−0.049	0.046	−0.002	0.007	0.011
	(0.044)	(0.043)	(0.018)	(0.030)	(0.022)
Economic situation	0.014	0.102***	0.002	−0.017	0.010
	(0.037)	(0.033)	(0.012)	(0.019)	(0.017)
Employed	−0.027	−0.207***	−0.004	0.000	0.031
	(0.063)	(0.07)	(0.028)	(0.034)	(0.033)
CPSU member	0.063	−0.007	−0.027	−0.038	0.025
	(0.117)	(0.091)	(0.054)	(0.073)	(0.050)
Voted	0.116*	0.440***	0.023	0.062*	0.036
	(0.058)	(0.073)	(0.035)	(0.034)	(0.034)
No. civil society orgs.	−0.002	0.003	0.009	0.042*	0.032
	(0.034)	(0.059)	(0.019)	(0.023)	(0.021)
Interest in politics	0.029	0.120***	−0.002	−0.016	−0.012
	(0.049)	(0.041)	(0.019)	(0.022)	(0.024)
Putin support	0.346***	0.338***	−0.029	−0.003	−0.043**
	(0.041)	(0.042)	(0.020)	(0.027)	(0.019)
United Russia support	0.113***	0.114***	0.013*	0.003	0.012
	(0.019)	(0.015)	(0.007)	(0.011)	(0.010)
Region fixed effects	yes	yes	yes	yes	yes
Observations	1,589	1,725	1,641	1,704	1,647
R^2	0.292	0.304	0.122	0.159	0.105

*** $p > 0.01$, ** $p > 0.05$, * $p > 0.1$

[a] This table examines the correlates of perceptions of electoral integrity and the acceptability of different types of electoral fraud using OLS models. The outcomes in the first two columns are measured on five-point scales, with higher values indicating more positive perceptions. The outcomes in columns 3–5 are all measured on four-point scales, with higher values indicating greater acceptance of these activities. All models cluster standard errors at the region level.

Moreover, few predictors are consistently associated with an individual's approval of different electoral malpractices, which might be expected given the tight distribution around unacceptability as shown in Table 1. Cross-national surveys suggest there's remarkable congruence

worldwide among masses and elites regarding the normative standards required to make an election free and fair.[54]

Of course, these correlations don't allow us to identify the direction of causality. Respondents who view elections as free and fair may reward the regime for upholding democratic procedures. Or they may view elections as honest because they themselves are under the influence of partisanship or propaganda. Either way, what's significant for our study is that such a correlation exists. The next section explores some important implications of this finding.

Heterogeneous Effects of Learning about Electoral Manipulation

Our main argument is that core regime supporters should be more sensitive than weakly aligned voters to electoral fraud. We test this by examining the heterogeneous treatment effects from the framing experiment outlined above. We hypothesize that support for a regime-affiliated candidate will fall more among individuals with a stronger preexisting affinity toward United Russia (and Putin) than among those with a weaker commitment to the regime. We use several measures of regime support for these purposes: a five-item scale measuring personal approval of Putin's time in office, a ten-item scale measuring support for UR, and a binary indicator for whether a respondent voted for UR in the 2016 parliamentary elections. The first two indicators capture respondents' self-reported support for the regime; the third measures actual behavior taken in support of UR. For our main analyses, we collapse the three types of electoral manipulations employed in the framing experiment into a binary treatment indicator (any fraud) for whether a respondent received any information about a candidate engaging in this type of behavior.

In Table 4, we show differences in means across both the different treatment conditions and the levels of regime support. In each panel, the three columns divide the sample into strong UR supporters (values nine and ten), weak UR supporters (values four through eight), and opposition supporters (values one through three). We then transform the five-point vote-intention scale into binary indicators that are easier to interpret as reflecting likelihood to vote (or to not vote) for the hypothetical UR candidate. In panel (a), we use a binary indicator coding equal to one if a respondent answered 4 or 5 on the turnout scale, indi-

[54] Norris 2013.

TABLE 4
TREATMENT EFFECTS BROKEN DOWN BY REGIME SUPPORT[a]

	Strong UR Supporters	Weak UR Supporters	Opposition
Control	62.8	46.9	29.2
Treatment (any fraud)	29.5	26.5	15.7

Propensity to Vote for UR Candidates (Choices 4 and 5)
(a)

	Strong UR Supporters	Weak UR Supporters	Opposition
Control	10.5	16.7	34.7
Treatment (any fraud)	43.2	39.3	54.6

Propensity to Vote against UR Candidates (Choices 1 and 2)
(b)

	Strong UR Supporters	Weak UR Supporters	Opposition
Control	26.7	36.5	36.1
Treatment (any fraud)	27.3	34.2	29.6

Propensity to Answer "Maybe Would, Maybe Wouldn't" (Choice 3)
(c)

	Strong UR Supporters	Weak UR Supporters	Opposition
Control	4.4	9.4	8.9
Treatment (any fraud)	9.4	11.6	9.2

Propensity to Answer "Don't Know" (Choice 7)
(d)

[a]All numbers are percentages. The outcome in panel (a) is a binary indicator if a respondent answered 4 or 5 on the five-point scale about whether they would vote for the hypothetical UR candidate (that is, likely voters). The outcome in panel (b) is a binary indicator if a respondent answered 1 or 2 on the five-point scale about whether they would vote for the hypothetical UR candidate (that is, unlikely voters). The outcome in panel (c) is a binary indicator, one if a respondent answered "maybe I would vote for him, maybe not" (response 3) to the question, and zero otherwise. The outcome in panel (d) is a binary indicator, one if a respondent answered "don't know/difficult to respond" (response 7) to the question, and zero otherwise.

cating that they were likely or definitely likely to vote for the candidate. Responses of 3 (50/50), 2 (unlikely), and 1 (definitely unlikely) were all coded as zero. The panel then shows the raw percentages for this variable based on treatment conditions (the rows). The control group didn't receive any of the three treatments indicating the candidate committed fraud, while the treatment group collapsed all respondents assigned to any of the three treatments.

The results show that strong UR supporters in the control group have

a 62.8 percent chance of voting for the hypothetical candidate, which makes sense given their partisanship. But when they're assigned any of the three fraud treatments, their average likelihood of voting for that candidate falls to 29.5 percent, a drop of 33.3 percentage points (or 53 percent, 33.3/62.8). Weak UR supporters don't react as strongly, showing only a 20.4 percentage-point drop (or 43 percent). The treatment effect for opposition supporters is roughly the same.

In panel (b) of Table 4, we code respondents who answered 1 (definitely unlikely) or 2 (unlikely) on the five-point scale as one, and zero otherwise. This measure gives the probability of the respondent voting against the UR candidate. Here again we see larger treatment effects for strong UR supporters, whose propensity to vote against the candidate increases by roughly four times (43.2 percent versus 10.5). It's noteworthy that 43 percent of strong UR supporters affirmed that they would not vote for the UR candidate if that candidate committed fraud. The treatment effects for the other two groups are sizable but not nearly as large. Weak UR supporters are a little more than twice as likely to oppose the UR candidates, while opposition supporters oppose slightly less than twice as often.

In panel (c), we perform the same exercise but with a binary indicator if the respondents answered that they might vote for the UR candidate, but they might not (that is, one if they responded 3 on the scale, and zero otherwise). We don't see any significant differences between treatment and control across the three groups. In addition, respondents in the experiment were allowed to answer "difficult to respond" instead of picking a value on the five-point scale of support for the UR candidate. Since approximately 11 percent of respondents struggled to answer, we might expect that difficulty to reflect differences in the way partisans handle information that conflicts with their priors about their preferred candidates. To explore this, we code a binary indicator in panel (d) for whether a respondent answered "don't know" (response 7). There is a treatment effect among strong UR supporters, although part of that may be because strong UR supporters in the control group were especially likely to have a concrete opinion about the UR candidate (only 4.4 percent answered "don't know"). Alternatively, these supporters could simply be unsure about whether they would vote for their copartisan.[55]

---

[55] In Section C4 of the supplementary material, we show robustness checks that indicate that most of the hard-to-answer treatment effects come from the bought votes treatment. Respondents seem to want more information about this specific practice before making up their minds about the candidate involved. But when all the treatment arms are collapsed, we don't observe that the collapsed any fraud treatment leads to more don't knows.

Taken together, we see that learning about fraud produces the largest effect on strong UR supporters. They become less likely to vote for UR candidates and more likely to affirm that they will not vote for them. We show a series of OLS models in Table 5, in which we regress the likelihood of a respondent voting for the candidate described in the experimental vignette on the any fraud treatment indicator and a range of covariates. In column 1, we exclude the treatment group from the model to examine the benchmark case (the control group). Intuitively, we find that individuals with a stronger affinity for the party are more likely to support its candidates, but that no other demographic characteristics predict support.[56] Adding the any fraud treatment in column 2 confirms the results presented above in Figure 1. Overall support for UR candidates drops when respondents learn about electoral manipulations being committed.

Columns 3–8 present heterogeneous treatment effects along three measures of support for the regime. We find consistent evidence in favor of our main hypothesis. UR candidates who engage in fraud see their electoral support drop more among core supporters than among weakly aligned voters. It makes little difference how the survey population is subset, whether it's by high versus low approval ratings of Putin in office (columns 3 and 4), by high or low levels of UR support more broadly (columns 5 and 6), or by having voted for UR (columns 7 and 8).[57] For the variables measuring Putin and UR approval ratings, the sample is subset among those at the very top of the scale (a rating of five out of five for Putin, or a rating above eight out of ten for UR) and those in the middle (a rating of three or four out of five for Putin, or a rating between four and eight for UR).[58] In each instance, the difference in coefficients on the treatment between regime and opposition supporters is large and statistically significant.

In Table C1 of the supplementary material, we show models including interactions between the treatment dummy and the three measures of regime support; the coefficients on the interaction terms are significant at the 95 percent level or above.[59] Figure 4 is produced on the basis of models 1 and 2 from that table, with the marginal effect of the any

[56] This set of null findings isn't particularly meaningful, given the inclusion of the *UR Support* variable. The results are robust to including or excluding the covariate controls.

[57] All results remain statistically and substantively unchanged when we remove the controls for *Putin support* and *UR support* in the respective models that examine the conditional effects of the other. These two variables are correlated at r = 0.53.

[58] The results are robust to including the bottom part of the distribution for both variables (the opposition) in the "low" category.

[59] Reuter and Szakonyi 2021b.

TABLE 5

HETEROGENEOUS EFFECTS OF LEARNING ABOUT ELECTORAL FRAUD[a]

	Control (1)	Full Sample (2)	Putin Approval		UR Approval		Voted for UR	
			High (3)	Low (4)	High (5)	Low (6)	High (7)	Low (8)
Any Fraud treatment		−0.642***	−0.974***	−0.589***	−0.997***	−0.621***	−0.896***	−0.395***
		(0.06)	(0.124)	(0.072)	(0.173)	(0.073)	(0.109)	(0.142)
Male	−0.204*	−0.113*	−0.161	−0.121*	−0.201	−0.151**	−0.108	−0.101
	(0.116)	(0.059)	(0.132)	(0.07)	(0.19)	(0.069)	(0.111)	(0.133)
Age (log)	0.073	0.019	0.139	0.007	0.266	0.062	0.237	−0.192
	(0.168)	(0.079)	(0.185)	(0.092)	(0.246)	(0.092)	(0.156)	(0.229)
Education	−0.058	0.016	0.049	0.004	−0.022	0.034	0.038	0.04
	(0.048)	(0.023)	(0.045)	(0.029)	(0.066)	(0.028)	(0.04)	(0.054)
Town size	−0.040	0.042*	0.047	0.050*	0.033	0.053*	0.085*	0.063
	(0.046)	(0.023)	(0.052)	(0.029)	(0.074)	(0.028)	(0.045)	(0.058)
Economic situation	0.054	0.003	0.034	−0.019	0.135*	−0.01	0.069	−0.124*
	(0.056)	(0.028)	(0.057)	(0.034)	(0.079)	(0.033)	(0.053)	(0.068)
Employed	−0.034	−0.03	−0.037	−0.049	−0.218	−0.009	0.054	0.020
	(0.113)	(0.06)	(0.133)	(0.071)	(0.181)	(0.072)	(0.116)	(0.145)
CPSU member	0.076	0.145	0.24	0.08	−0.167	0.237*	0.112	0.051
	(0.181)	(0.096)	(0.183)	(0.119)	(0.252)	(0.122)	(0.164)	(0.178)
Voted	0.001	0.162***	0.273**	0.171**	0.145	0.164**		
	(0.116)	(0.061)	(0.135)	(0.072)	(0.192)	(0.073)		
No. civil society orgs.	−0.064	−0.018	−0.030	−0.029	0.080	−0.010	−0.071	−0.069
	(0.057)	(0.029)	(0.053)	(0.036)	(0.095)	(0.033)	(0.058)	(0.059)
Interest in politics	0.017	0.026	−0.016	0.017	−0.068	0.038	−0.059	0.054
	(0.058)	(0.029)	(0.065)	(0.036)	(0.087)	(0.036)	(0.056)	(0.076)
Putin support	0.100	0.064*			−0.052	0.065	0.03	0.340***
	(0.071)	(0.034)			(0.125)	(0.041)	(0.068)	(0.068)
United Russia support	0.111***	0.070***	−0.024	0.084***				
	(0.028)	(0.014)	(0.028)	(0.016)				
Region fixed effects	yes	yes	yes	yes	yes	yes	yes	yes
Observations	436	1,610	404	1,091	258	1,079	533	331
R^2	0.263	0.245	0.391	0.235	0.398	0.249	0.326	0.296

*** $p > 0.01$, ** $p > 0.05$, * $p > 0.1$

[a]This table uses OLS models to examine ordinal responses from the framing experiment. Column 1 restricts the analysis to only the control group, which received no information about fraud. Column 2 estimates the treatment effect graphically depicted in Figure 1 and includes covariates. Columns 3 and 4 use a five-point scale to subset to respondents with high levels of approval of Putin's performance in office (a value of five) and low levels (values of three and four). Columns 5 and 6 use a ten-point scale to subset the sample to respondents with high levels of approval of United Russia (values higher than seven) and low levels (values between four and seven). Columns 7 and 8 subset the sample by whether the respondent voted for United Russia in the 2016 Duma Election.

Putin Approval
(a)
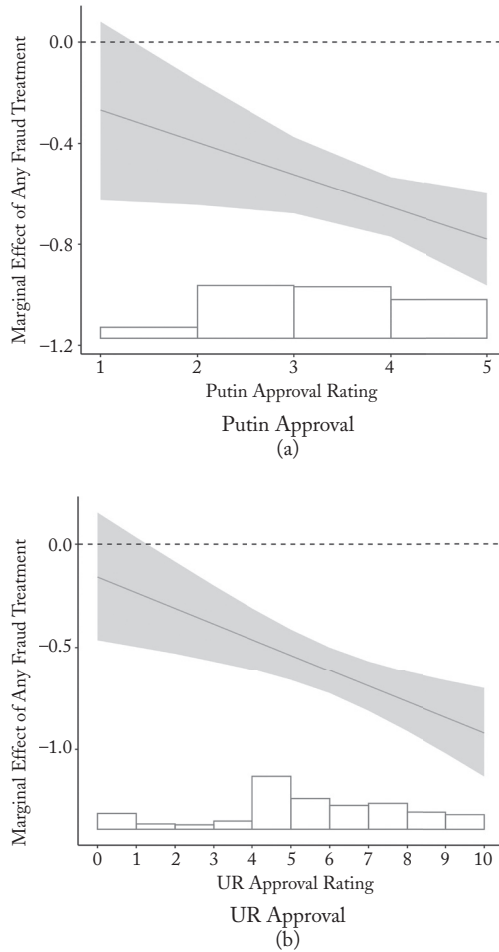
UR Approval
(b)

FIGURE 4
MARGINAL EFFECTS

fraud treatment shown on the y-axis across different values of Putin's approval rating (a) and support for UR (b); the distribution of responses is shown as inlaid histograms. There is a strong negative relationship between the degree of support for the regime and the effect of learning about electoral fraud committed by affiliated candidates.

One concern is that these large differences could be driven by a mechanical feature of our measurement strategy. Since regime supporters are more likely to back the UR candidate ex ante (that is, without any knowledge of their campaign activities), their pretreatment level of candidate support will obviously be higher and therefore, these respon-

dents have farther to fall on the five-point scale. For example, consider an extreme scenario in which the effect of revealing fraud is to induce all respondents to report that they will "definitely not vote" for the candidate (this is equal to 1, the lowest point on the scale). Swing voters, whose pretreatment level of support is 3, exhibit a treatment effect of 2, while core supporters, whose pretreatment level of support is 5, will exhibit a treatment effect of 4. This scenario is patently implausible, but it illustrates the mathematical problem.

But our results are not driven by this floor effect. The percentage drop relative to the group is still higher among strongly aligned regime partisans than it is among those weakly aligned. For example, support for the UR candidate among strong Putin supporters (column 3 in Table 5) falls by 24 percent (relative to their baseline level) when they're informed of electoral fraud; among swing voters (column 4), support drops 18 percent relative to the baseline. Differential effects are still present. Strong regime supporters are more turned off by learning that UR candidates commit electoral violations than are weak ones.

REGIME PERCEPTIONS, INFORMATION, AND THE EFFECT OF ELECTORAL FRAUD

Why does evidence of electoral fraud have such a strong effect on core regime supporters? What are these voters learning that makes them withdraw their support? In Table 6, we explore several explanations. One possibility is that electoral fraud undermines the regime's reputation for propriety. It's hard to measure a respondent's views on the honesty of the regime, especially when different respondents may conceive of the regime in different ways. In Russia, almost all regime supporters also support Putin, and most associate United Russia closely with him. Thus, one adequate proxy could be the respondent's views of Putin's character.

Surveys find that Russians identify several positive traits in Putin. In our survey, 71 percent of respondents agreed with the statement that he is a "strong leader" (24 percent said "mostly yes").[60] Another trait that voters associate with Putin is honesty. In the 2016 RES, 54 percent of respondents agreed that Putin is honest and deserving of trust (33 percent said "mostly yes").

[60] The question asked respondents whether they agreed with certain evaluations of Putin's character, prompting them with a four-point scale with values of no, mostly no, mostly yes, and yes. Voters also view Putin as being competent (77 percent). Interestingly, voters don't simply evaluate Putin highly on every dimension. Only forty-four percent thought that he "really thinks about the interests of people like me."

TABLE 6

How Electoral Fraud Undermines Perceptions of the Regime[a]

	Putin Is Strong		Putin Is Honest		Electoral Integrity		2016 Political Internet	
	High (1)	Low (2)	High (3)	Low (4)	High (5)	Low (6)	Yes (7)	No (8)
Any fraud treatment	−0.701***	−0.466***	−0.763***	−0.492***	−0.918***	−0.579***	−0.451***	−0.696***
	(0.072)	(0.115)	(0.087)	(0.096)	(0.139)	(0.073)	(0.147)	(0.067)
Male	−0.101	−0.066	−0.155*	−0.104	−0.046	−0.125*	−0.366***	−0.027
	(0.071)	(0.107)	(0.087)	(0.092)	(0.137)	(0.070)	(0.136)	(0.066)
Age (log)	0.070	−0.134	0.126	−0.128	0.147	−0.051	−0.124	−0.028
	(0.095)	(0.155)	(0.113)	(0.131)	(0.168)	(0.099)	(0.213)	(0.085)
Education	0.015	−0.007	0.034	0.005	0.047	0.008	0.021	0.017
	(0.028)	(0.044)	(0.034)	(0.037)	(0.051)	(0.028)	(0.060)	(0.026)
Town size	0.064**	0.004	0.053	0.035	−0.009	0.055*	0.105*	0.013
	(0.028)	(0.049)	(0.033)	(0.039)	(0.066)	(0.028)	(0.060)	(0.026)
Economic situation	−0.009	0.00004	0.007	−0.008	−0.037	0.017	−0.029	0.020
	(0.033)	(0.052)	(0.041)	(0.044)	(0.066)	(0.033)	(0.064)	(0.031)
Employed	−0.026	−0.098	0.023	−0.062	−0.157	−0.034	0.010	−0.044
	(0.073)	(0.110)	(0.090)	(0.093)	(0.147)	(0.070)	(0.147)	(0.067)
CPSU member	0.119	0.156	0.180	0.164	0.099	0.153	0.191	0.168
	(0.116)	(0.184)	(0.136)	(0.153)	(0.198)	(0.117)	(0.283)	(0.103)
Voted	0.217***	0.168	0.306***	0.105	0.220	0.202***	0.186	0.217***
	(0.075)	(0.110)	(0.092)	(0.093)	(0.159)	(0.071)	(0.142)	(0.068)
No. civil society orgs.	−0.006	−0.127*	−0.010	−0.017	−0.043	−0.013	−0.020	−0.030
	(0.033)	(0.068)	(0.047)	(0.042)	(0.081)	(0.034)	(0.086)	(0.032)
Interest in politics	0.012	0.058	−0.020	0.058	−0.030	0.016	0.024	0.043
	(0.036)	(0.056)	(0.043)	(0.047)	(0.074)	(0.036)	(0.085)	(0.032)
United Russia support	0.062***	0.128***	0.026	0.107***	0.008	0.095***	0.050*	0.083***
	(0.015)	(0.022)	(0.020)	(0.018)	(0.030)	(0.014)	(0.028)	(0.013)
Region fixed effects	yes	yes	yes	yes	yes	yes	yes	yes
Observations	1,157	436	794	661	366	1,077	335	1,269
R^2	0.228	0.399	0.239	0.305	0.367	0.243	0.314	0.272***

*** $p > 0.01$, ** $p > 0.05$, * $p > 0.1$

[a]This table uses OLS regression analysis to examine additional heterogeneous treatment effects. Columns 1 and 2 use a four-point scale to subset the sample into respondents who rank Putin highly as a strong leader (top value of four) and those who rank him lower (values of less than four). Columns 3 and 4 use a four-point scale to subset the sample into respondents who rank Putin highly as an honest person (top value of four) and those who rank him lower (values of less than four). Columns 5 and 6 use a five-point scale about whether respondents believe elections are conducted honestly to subset the sample by those with positive views (top value of five) and those with less positive views (values of less than five). Columns 7 and 8 subset according to a binary indicator for whether the respondent read Internet news during the 2016 campaign.

Learning about electoral fraud may undermine perceptions of Putin's virtues. Voters who receive objective information about their politicians being corrupt and dishonest are more likely to rescind their electoral support.[61] Despite the appeal of lying to hide undesirable characteristics, honest candidates still enjoy electoral advantages because voters highly value trustworthy candidates whatever their policy promises.[62] Similarly, committing electoral fraud can signal weakness—autocrats need to break the formal electoral rules to ward off challengers.

In Table 6, we construct model specifications that are similar to those used above to examine heterogeneity in treatment effects based on respondents' views of Putin's character and on their appraisal of elections in Russia. We find that those who believe Putin is strong (columns 1 and 2) and honest (columns 3 and 4) are more likely to react negatively when they learn that a UR candidate has committed fraud.[63]

Another possibility is that the effect depends on preconceived notions of fraud. Those who already believe elections are fraudulent should be less likely to update their views on UR candidates when they learn about fraud. Conversely, those who think elections are free and fair should be more likely to update their candidate preferences when they learn about fraud. Columns 5 and 6 examine whether the treatment effect varies according to whether respondents believe that Russian elections are free and fair. The results strongly suggest that they do. The effect of the fraud treatment is much larger among respondents who think Russian elections are free and fair.

Columns 7 and 8 take a different approach to the same question. Specifically, we examine whether respondents who are plausibly less exposed to information about fraud are more affected by the fraud treatment in our experiment. In Russia, information on fraud is rarely encountered on television or radio or in (most) print media. The Internet is practically the only media platform where Russians might learn about fraud. But not all Russians actively use the Internet, and most don't use it for reading political news. In columns 7 and 8, we subset our models according to whether respondents reported that they used the Internet for reading political news during the 2016 election campaign. We find that the treatment had smaller effects for those who reported having read political news on the Internet during the campaign. These individuals are more likely to be preexposed to information on

---

[61] Ferraz and Finan 2008.

[62] Callander and Wilkie 2007.

[63] We divide respondents into two groups: those who said yes when asked to evaluate Putin on these dimensions, and those who gave any other answer.

fraud. The treatment effects were larger for those respondents who accessed Internet political news.[64]

Replication and Extension

One shortcoming of our experiment is that it can't distinguish between two mechanisms that could be driving the observed drop in support among core regime supporters. Electoral fraud may be leading regime supporters to consider other candidates or to consider abstention. Either way, the findings indicate that regime supporters are withdrawing their political support from regime candidates, but it's interesting to separate these potential mechanisms.

In particular, it's possible that fraud might drastically reduce turnout by the opposition, which would offset any decrease in support by regime supporters. Several studies find that fraud deters participation by the opposition.[65] Our findings would have less meaning if fraud produced a drop in regime support that was outweighed by a concomitant decline in opposition turnout. In other words, UR may not fear a slight diminishment in its core support if violations of electoral integrity also cause opposition supporters to disengage from politics and to cease voting against the regime.

To address this, we placed two additional survey experiments on a representative survey of sixteen hundred Russian adults conducted in May 2018, roughly eighteen months after our original survey. The vignette used in both 2018 experiments was nearly identical to that used in September 2016. We gave respondents information about a fictional fifty-year-old businessman from UR who was running for the State Duma in the next elections; this person had also adopted two children.[66] The experimental treatment gave half the sample this additional information: the candidate had organized a multiple-voting scheme, ferrying voters by bus to multiple precincts. Our wording was identical to that used in the "organized karusels" treatment in the 2016 experiment, as shown in Table 2 (see Section E1 of the supplementary material for the exact wording).[67]

The important difference in this second set of experiments is the out-

---

[64] Note that all models here control for regime support. In Section B3 of the supplementary material, we show these heterogeneous treatment effects in just the regime support subset; Reuter and Szakonyi 2021b. Findings are similar. Table B9 of the supplementary material also explores other measures of exposure to prior information about fraud.

[65] McCann and Domínguez 1998; Simpser 2012.

[66] Note that we use a single occupational background, given constraints on sample size.

[67] Reuter and Szakonyi 2021b.

come variable. Our 2018 turnout experiment asked respondents to state their likelihood of turning out to vote on a scale of one to five, and was administered to half the respondents. Our 2018 vote choice experiment was given to the other half of the respondents, who were asked to state their likelihood of voting for this candidate on a scale of one to five. Thus, the outcome in the 2018 vote choice experiment is identical to that asked in the 2016 experiment analyzed above, while the 2018 turnout experiment focuses only on whether respondents would vote at all. Table 7 presents the breakdown of respondents across the different treatment arms and outcome variables. Each respondent was assigned to receive either the turnout or vote choice experiment, and within each one, each respondent had a 50 percent chance of receiving the treatment—that is, learning that the UR candidate had committed fraud.

These experiments accomplish several objectives. First, the 2018 vote choice experiment is essentially a replication check of our initial results from the 2016 experiment, but it uses a simplified set of treatment conditions. This helps to build confidence that the patterns identified above analyzing the 2016 experiment aren't specific to the Russian political climate that year. Second, the 2018 turnout experiment lets us investigate whether learning about fraud decreases turnout or support for the candidate responsible for it.

Figure 5 presents the results. In panel (a), we see that the treatment effect of fraud on turnout is negative. In the control group, the mean turnout propensity on a five-point scale, with 3 indicating uncertainty, is 3.44. When respondents learn that the candidate has committed fraud, that number drops to 2.89, an effect of −0.55 that is statistically significant at the 99 percent level. In general, people are less likely to vote when electoral integrity suffers. In the right panel, we again see that support for the candidate committing the fraud also drops. The treatment effect of −0.62 is roughly the same using a five-point scale measuring candidate support.

In Figure 6, we explore heterogeneity across different levels of support for Putin, measured on a four-point scale. As before, we show the marginal effects of the fraud treatment for each outcome in the panels: turnout (a) and vote choice (b). The point estimates come from models that control for demographics such as age, income, and employment status. First, we see a slightly positive, but not statistically significant interaction effect of fraud and Putin support on turnout. Both regime and opposition supporters are less likely to turn out after learning that a UR candidate committed fraud, and the degree to which fraud dissuades them from voting is relatively small. Just as important, we repli-

TABLE 7
2018 EXPERIMENTS COVERAGE[a]

Experimental Outcome	Turnout	Vote Choice
No electoral manipulation	363	371
Organized karusels to take voters to polls	362	350

[a] Total respondents who received "turnout" outcome, 725; total respondents who received "vote choice" outcome, 721.

cate our findings from the 2016 experiment in the right panel: regime supporters are significantly more turned off than opposition supporters by new information on UR-sponsored fraud. The substantive effect sizes are roughly the same as they were two years prior. In the supplementary material, we show that the effects are robust to interacting the treatment with a ten-point scale of support for United Russia.[68]

Several things about these results are worth noting. First, contrary to some existing accounts, the findings demonstrate that fraud reduces turnout not only among the opposition, but also among regime supporters. Several studies argue that fraud creates the perception that opposition votes will not count.[69] But it stands to reason that fraud could produce a similar effect among regime supporters who, if they realize that electoral outcomes are predetermined, should have less reason to think their vote will matter and less incentive to vote. Consistent with this idea, our experiment shows that fraud reduces turnout across the electorate. Indeed, observational evidence from the 2016 RES shows that perceptions of electoral manipulation reduced self-reported turnout among regime supporters just as much as it did among the opposition and swing voters.[70]

Note that this finding isn't necessarily at odds with the theoretical arguments in previous work. Even if fraud reduces feelings of electoral efficacy among opposition supporters more than it does among regime supporters, it could still be the case that regime supporters would be more disillusioned by the revelation of fraud. In other words, the mechanism we propose in this study could be operating alongside the differential electoral efficacy argument to produce the findings we see in the 2018 turnout experiment.

Taken together, the results suggest that the heterogeneous effects in

[68] See Section E of the supplementary material; Reuter and Szakonyi 2021b.

[69] McCann and Domínguez 1998; Simpser 2012; Nikolayenko 2015.

[70] Models in Table B2 of the supplementary material show a positive and statistically significant relationship between perceptions of electoral integrity and turnout among regime and opposition supporters; Reuter and Szakonyi 2021b.
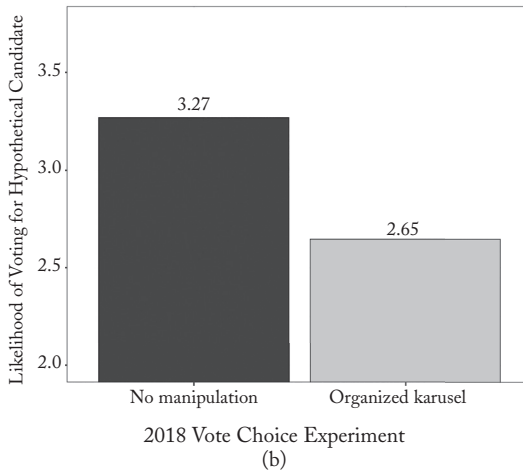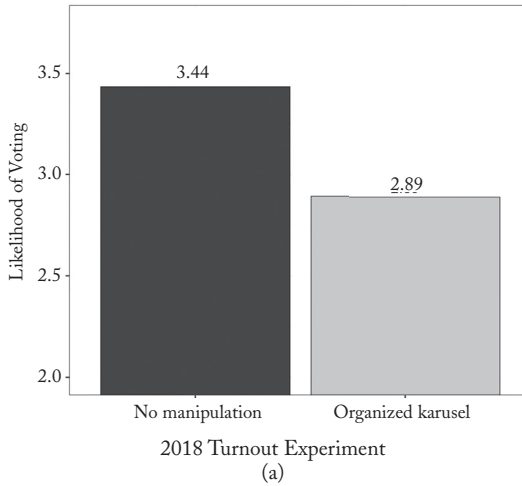
FIGURE 5
FRAUD, TURNOUT, AND VOTE CHOICE

our main 2016 experiment are driven by changes in vote choice rather than changes in turnout. Since fraud appears to reduce turnout equally among regime supporters and the opposition, it stands to reason that the larger treatment effects for regime supporters in the 2016 experiment (and in the 2018 vote choice experiment) are being driven by decisions to withdraw support from regime candidates. Once inside the voting booth, core regime supporters appear to be abandoning ruling party candidates who commit fraud.
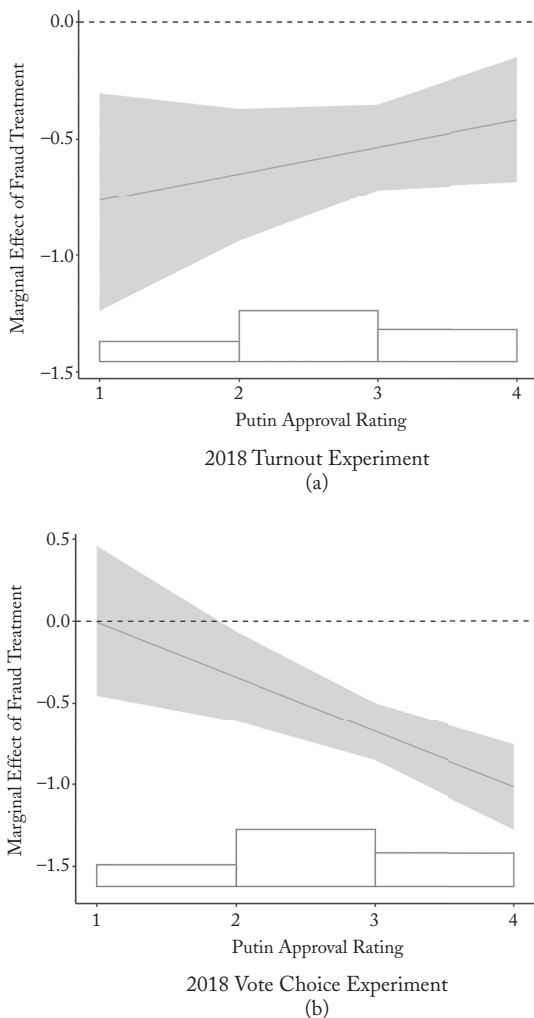
FIGURE 6
HETEROGENEOUS EFFECTS OF FRAUD ON TURNOUT AND VOTE CHOICE

These findings reinforce our contention that fraud is electorally costly for the regime. If fraud reduced opposition turnout to such a degree that it offset any loss of support from regime supporters, then fraud would not undermine the regime's chances of winning. Our findings suggest that this is not the case. The 2018 vote choice experiment shows that fraud reduces turnout for opposition and regime supporters to an equal degree. Moreover, fraud appears to cause regime supporters to withdraw their support from fraudulent United Russia candidates.

DISCUSSION AND IMPLICATIONS

In summary, our findings suggest that voters in Russia punish regime candidates who engage in fraud. This effect is largest among those who are the strongest supporters of the regime. Polarization is not so strong in Russia that regime supporters excuse regime candidates for fraud.[71] Instead, they punish them for it. Most regime supporters believe that elections are free and fair and most believe that is how it should be. Gaining awareness of electoral fraud dispels preconceived notions about the regime and its electoral propriety. When fraud is revealed, many proregime voters withdraw their support, which appears to be conditional on the government maintaining its commitment to clean elections.

These findings have important implications for the comparative literature on autocracy and for the study of contemporary Russian politics. For studies on comparative autocracy, our findings highlight an understudied consequence of electoral fraud. Much of the recent neo-institutional literature on electoral fraud has centered on how fraud sends a signal of strength to elites.[72] One puzzle that emerges from this literature is why autocrats try so hard to conceal fraud. If fraud deters all sorts of subversive and oppositional activity, why don't autocrats publicize it? Scholars of contentious politics suggest that they don't publicize it because it may lead to opposition protest.[73] That seems hard to deny, but we highlight another reason why autocrats disguise fraud— their core supporters are turned off by it and if they learn of it, they will withdraw their support from the regime. That polarization is relatively limited in Russia suggests that findings from this survey experiment reflect actual behavior: strong partisan biases are less likely to outweigh normative concerns in the voting booth in Russia than they might be in polarized countries like Venezuela or the United States.

More generally, our study suggests that scholars of autocracy should pay more attention to the democratic features of nondemocratic elections. The neo-institutional literature on autocracy has made great strides in pointing out the autocratic functions of nominally democratic institutions. But in the midst of the neo-institutional revolution, research continues to show that these elections also serve a democratic function, improving accountability[74] and providing legitimacy to the re-

[71] cf. Svolik 2020.
[72] Rozenas 2016; Simpser 2013; Gehlbach and Simpser 2015.
[73] Tucker 2007; Bunce and Wolchik 2011.
[74] Miller 2015.

gime.[75] Large parts of the electorate expect that elections will be democratic.

These findings also have important implications for how scholars study politics in Russia. This article should serve as a reminder that the demand for democratic institutions remains strong in Russia. In a revealing analysis of Putin's *Pryamaya Liniya* call-in shows, Susanne Wengle and Christine Evans note that Putin frequently touts the role of formal democratic institutions.[76] These authors wonder why Putin seems to frame so much of his political discourse around institutions. Our account demonstrates why the performance is so important. Many voters believe in electoral democracy or at the very least, behave as if they do. Thus, one of the reasons that elections are maintained in Russia is because voters support them.[77] This isn't to say that Russia is a democracy; it is not. But important parts of the electorate behave as if elections are democratic and expect them to be so. Analyses of authoritarian Russia would be remiss to ignore these voters. Understanding their behavior is key to understanding the stability of the regime.

We also provide insight into why the Putin regime goes to such great lengths to both hide and limit electoral fraud. After the 2011–2012 election cycle, regime leaders made it clear to regional subordinates that they wanted future elections to be cleaner—or at least to be perceived as clean. The government spent more than $800 million to install live-streaming cameras in electoral precincts in 2012, and then later appointed the former human rights ombudsman Ella Pamfilova to oversee the Central Election Commission. Available evidence indicates that election cycles since 2011 have been marked by less blatant election-day fraud.[78] The conventional explanation for this new emphasis on electoral legitimacy is that the regime wanted to stem the opposition protest movement that had erupted during the 2011–2012 cycle. But our findings suggest another possibility: regime leaders believed that their electoral base would evaporate if the curtain on fraud was pulled back. The scope of these efforts suggests that fraud could become a salient voting issue if voters were to find out about it.

We believe that such dynamics could also be at play in other electoral autocracies. Our analyses in Table F2 of the supplementary material

[75] Morgenbesser 2017; Gandhi and Lust-Okar 2009.
[76] Wengle and Evans 2018.
[77] Note that this is different from arguing that the authorities hold elections because it's a procedural norm. The regime needs to limit fraud—or limit the spread of information on fraud—because faking elections has real costs in terms of regime support.
[78] GOLOS, "Election Observation Statement 18 September 2016," GOLOS Movement, September 19, 2016. At https://www.golosinfo.org/articles/117564, accessed November 20, 2020.

show that regime supporters in countries as diverse as Kazakhstan, Malaysia, Singapore, and Nigeria all give their governments high marks for upholding democratic practices.[79] Their support for the regime may be contingent on a belief that electoral integrity continues to be respected. This may be especially true in countries where autocrats initially won free elections and then stealthily undermined democratic systems to hold onto power. Voters may believe that electoral results fairly reflect the autocrat's popularity, but they may not be fully aware of the degree of malpractice being committed. Providing information about fraud could change their calculus of support for the regime.

Of course, such conclusions come with caveats. These survey experiments are hypothetical and although they do suggest why new information can cause voter defections, they can't illuminate the conditions under which that happens. But they can illuminate the ways by which fraud impacts voter affect: they suggest that voters have a psychological reaction to fraud. Nevertheless, they can't be extrapolated directly to explain real events. On the one hand, in the real world, information on fraud is contested and subject to perceptual bias. Thus, our experimental estimates may represent a higher bound for the effect of fraud on regime vote totals. On the other hand, the hypothetical nature of survey experiments may mute respondents' reactions. Voters who learn about real candidates committing real fraud could be even more disappointed. Future research could profit by extending these analyses into real-world settings.

In addition, our study can't precisely quantify the net costs of engaging in fraud in the real world. Even if fraud costs autocrats votes by driving away supporters, stuffing ballots or rewriting protocols still adds to a regime's vote totals. The point of our study is not to claim that the former must outweigh the latter, but rather to demonstrate that the loss of votes is a real concern. Factors like the presence of independent media and the competitiveness of the election are likely to affect this calculus. As we show in Section B3 of the supplementary material, regime supporters who rely heavily on state-sponsored news for information are more affected by the experimental treatment.[80] In addition, a strong opposition not only has greater resources and the incentive to inform the public about any electoral fraud committed during election, but it will also attract more support from voters newly disillusioned by the regime's false claims to be upholding free and fair elections.

Our study suggests some other avenues for future research as well.

[79] Reuter and Szakonyi 2021b.
[80] Reuter and Szakonyi 2021b.

For one, little is known about how voters become informed about fraud. Our experimental intervention induced voters to believe that fraud had occurred, but in an autocracy with a partially closed media environment, it's difficult for voters to find out about electoral fraud. Social and independent media clearly play a role here,[81] as do election monitors.[82] Even less is known about how opposition activists can break through partisan biases to broaden the awareness of fraud. The field seems to be moving in the right direction toward answering these questions, but more work is needed.

We also know little about how the vote-depressing effects of fraud compared to other types of unethical and socially undesirable behavior. For example, do voters punish candidates more for committing fraud than for engaging in corruption, committing other crimes, or performing unpatriotic acts? Future research could advance the literature by benchmarking the vote-depressing effects of fraud against other such issues.

SUPPLEMENTARY MATERIAL

Supplementary material for this article can be found at https://doi.org/10.1017/S0043887120000234.

DATA

Replication files for this article can be found at https://doi.org/10.7910/DVN/ITLVMH.

REFERENCES

Anduiza, Eva, Aina Gallego, and Jordi Muñoz. 2013. "Turning a Blind Eye: Experimental Evidence of Partisan Bias in Attitudes toward Corruption." *Comparative Political Studies* 46, no. 12: 1664–92. At https://doi.org/10.1177/0010414013489081.

Bartels, Larry M. 2002. "Beyond the Running Tally: Partisan Bias in Political Perceptions." *Political Behavior* 24, no. 2: 117–50. At https://doi.org/10.1023/A:1021226224601.

Bhatti, Yosef, Kasper M. Hansen, and Asmus Leth Olsen. 2013. "Political Hypocrisy: The Effect of Political Scandals on Candidate Evaluations." *Acta Politica* 48, no. 4: 408–28. At https://doi.org/10.1057/ap.2013.6.

Birch, Sarah. 2011. *Electoral Malpractice*. New York, N.Y.: Oxford University Press.

Bunce, Valerie J., and Sharon L. Wolchik. 2011. *Defeating Authoritarian Leaders in Postcommunist Countries*. New York, N.Y.: Cambridge University Press.

Callander, Steven, and Simon Wilkie. 2007. "Lies, Damned Lies, and Political Campaigns." *Games and Economic Behavior* 60, no. 2: 262–86. At https://doi.org/10.1016/j.geb.2006.12.003.

[81] Reuter and Szakonyi 2015.
[82] Robertson 2017.

Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1960. *The American Voter*. Chicago, Ill.: University of Chicago Press.

Cantú, Francisco, and Omar García-Ponce. 2015. "Partisan Losers' Effects: Perceptions of Electoral Integrity in Mexico." *Electoral Studies* 39: 1–14. At https://doi.org/10.1016/j.electstud.2015.03.004.

Carey, John, Katherine Clayton, Gretchen Helmke, Brendan Nyhan, Mitchell Sanders, and Susan Stokes. 2020. "Who Will Defend Democracy? Evaluating Tradeoffs in Candidate Support among Partisan Donors and Voters." *Journal of Elections, Public Opinion and Parties*. At https://doi.org/10.1080/17457289.2020.1790577.

Cornelius, Wayne A. 1975. *Politics and the Migrant Poor in Mexico City*. Stanford, Calif.: Stanford University Press.

Doherty, David, Conor M. Dowling, and Michael G. Miller. 2016. "When Is Changing Policy Positions Costly for Politicians? Experimental Evidence." *Political Behavior* 38, no. 2: 455–84. At https://doi.org/10.1007/s11109-015-9321-9.

Duch, Raymond M., Harvey D. Palmer, and Christopher J. Anderson. 2000. "Heterogeneity in Perceptions of National Economic Conditions." *American Journal of Political Science* 44, no. 4: 635–52. At https://doi.org/10.2307/2669272.

Enikolopov, Ruben, Vasily Korovkin, Maria Petrova, Konstantin Sonin, and Alexei Zakharov. 2013. "Field Experiment Estimate of Electoral Fraud in Russian Parliamentary Elections." *Proceedings of the National Academy of Sciences* 110, no. 2: 448–52. At https://doi.org/10.1073/pnas.1206770110.

Fearon, James D. 2011. "Self-Enforcing Democracy." *Quarterly Journal of Economics* 126, no. 4: 1661–708. At https://doi.org/10.1093/qje/qjr038.

Ferraz, Claudio, and Frederico Finan. 2008. "Exposing Corrupt Politicians: The Effect of Brazil's Publicly Released Audits on Electoral Outcomes." *Quarterly Journal of Economics* 123, no. 2: 703–45. At https://doi.org/10.1162/qjec.2008.123.2.703.

Frye, Timothy, Ora John Reuter, and David Szakonyi. 2014. "Political Machines at Work: Voter Mobilization and Electoral Subversion in the Workplace." *World Politics* 66, no. 2 (April): 195–228. At https://doi.org/10.1017/S004388711400001X.

Frye, Timothy, Ora John Reuter, and David Szakonyi. 2019. "Hitting Them with Carrots: Voter Intimidation and Vote Buying in Russia." *British Journal of Political Science* 49, no. 3: 857–81. At https://doi.org/10.1017/S0007123416000752.

Gandhi, Jennifer, and Ellen Lust-Okar. 2009. "Elections under Authoritarianism." *Annual Review of Political Science* 12: 403–22. At https://doi.org/10.1146/annurev.polisci.11.060106.095434.

Gehlbach, Scott, and Alberto Simpser. 2015. "Electoral Manipulation as Bureaucratic Control." *American Journal of Political Science* 59, no. 1: 212–24. At https://doi.org/10.1111/ajps.12122.

GOLOS. 2012. *Domestic Monitoring of Elections of the President of Russian Federation*. [In Russian.] Moscow: Association GOLOS. At http://archive.golos.org/asset/5878, accessed November 20, 2020.

Gonzalez Ocantos, Ezequiel, Chad Kiewiet de Jonge, and David W. Nickerson. 2014. "The Conditionality of Vote-Buying Norms: Experimental Evidence

from Latin America." *American Journal of Political Science* 58, no. 1: 197–211. At https://doi.org/10.1111/ajps.12047.

Gutiérrez-Romero, Roxana, and Adrienne LeBas. 2020. "Does Electoral Violence Affect Vote Choice and Willingness to Vote? Conjoint Analysis of a Vignette Experiment." *Journal of Peace Research* 57, no. 1: 77–92. At https://doi.org/10.1177/0022343319892677.

Jerit, Jennifer, and Jason Barabas. 2012. "Partisan Perceptual Bias and the Information Environment." *Journal of Politics* 74, no. 3: 672–84. At https://doi.org/10.1017/S0022381612000187.

Kramon, Eric. 2016. "Electoral Handouts as Information: Explaining Unmonitored Vote Buying." *World Politics* 68, no. 3 (July): 454–98. At https://doi.org/10.1017/S0043887115000453.

Kunda, Ziva. 1990. "The Case for Motivated Reasoning." *Psychological Bulletin* 108, no. 3: 480–98. At https://doi.org/10.1037/0033-2909.108.3.480.

Lebo, Matthew J., and Daniel Cassino. 2007. "The Aggregated Consequences of Motivated Reasoning and the Dynamics of Partisan Presidential Approval." *Political Psychology* 28, no. 6: 719–46. At https://doi.org/10.1111/j.1467-9221.2007.00601.x.

Levitsky, Steven, and Lucan A. Way. 2010. *Competitive Authoritarianism: Hybrid Regimes after the Cold War*. New York, N.Y.: Cambridge University Press.

Magaloni, Beatriz. 2006. *Voting for Autocracy: Hegemonic Party Survival and its Demise in Mexico*. New York, N.Y.: Cambridge University Press.

Mares, Isabela, and Lauren Young. 2016. "Do Voters Punish Electoral Malfeasance?" At https://www.academia.edu/31970095/Do_voters_punish_electoral_malfeasance, accessed November 20, 2020.

McAllister, Ian, and Stephen White. 2011. "Public Perceptions of Electoral Fairness in Russia." *Europe-Asia Studies* 63, no. 4: 663–83. At https://doi.org/10.1080/09668136.2011.566429.

McCann, James A., and Jorge I. Domínguez. 1998. "Mexicans React to Electoral Fraud and Political Corruption: An Assessment of Public Opinion and Voting Behavior." *Electoral Studies* 17, no. 4: 483–503. At https://doi.org/10.1016/S0261-3794(98)00026-2.

Miller, Michael K. 2015. "Elections, Information, and Policy Responsiveness in Autocratic Regimes." *Comparative Political Studies* 48, no. 6: 691–727. At https://doi.org/10.1177/0010414014555443.

Morgenbesser, Lee. 2017. "The Autocratic Mandate: Elections, Legitimacy and Regime Stability in Singapore." *Pacific Review* 30, no. 2: 205–31. At https://doi.org/10.1080/09512748.2016.1201134.

Myagkov, Mikhail, Peter C. Ordeshook, and Dimitri Shakin. 2009. *The Forensics of Election Fraud: Russia and Ukraine*. New York, N.Y.: Cambridge University Press.

Nikolayenko, Olena. 2015. "Do Contentious Elections Depress Turnout?" In Pippa Norris, Richard W. Frank, and Ferran Martínez i Coma, eds., *Contentious Elections: From Ballots to Barricades*. New York, N.Y.: Routledge.

Norris, Pippa. 2013. "Does the World Agree about Standards of Electoral Integrity? Evidence for the Diffusion of Global Norms." *Electoral Studies* 32, no. 4: 576–88. At https://doi.org/10.1016/j.electstud.2013.07.016.

Norris, Pippa. 2014. *Why Electoral Integrity Matters*. New York, N.Y.: Cambridge University Press.

ODIHR (Office for Democratic Institutions and Human Rights). 2003. *Russian Federation: State Duma Elections*. Warsaw, Poland: Organization for Security and Co-operation in Europe. At https://www.osce.org/files/f/documents/1/1/18283.pdf, accessed November 20, 2020.

Pietsch, Juliet. 2015. "Authoritarian Durability: Public Opinion towards Democracy in Southeast Asia." *Journal of Elections, Public Opinion and Parties* 25, no. 1: 31–46. At https://doi.org/10.1080/17457289.2014.933836.

Reuter, Ora John, and David Szakonyi. 2015. "Online Social Media and Political Awareness in Authoritarian Regimes." *British Journal of Political Science* 45, no. 1: 29–51. At https://doi.org/10.1017/S0007123413000203.

Reuter, Ora John, and David Szakonyi. 2021a. "Replication data for Electoral Manipulation and Regime Support: Survey Evidence from Russia." Harvard Dataverse, V1. At https://doi.org/10.7910/DVN/ITLVMH.

Reuter, Ora John, and David Szakonyi. 2021b. "Supplementary material for Electoral Manipulation and Regime Support: Survey Evidence from Russia." At https://doi.org/10.1017/S0043887120000234.

Robertson, Graeme. 2017. "Political Orientation, Information and Perceptions of Election Fraud: Evidence from Russia." *British Journal of Political Science* 47, no. 3: 589–608. At https://doi.org/10.1017/S0007123415000356.

Rose, Richard, and William Mishler. 2009. "How Do Electors Respond to an 'Unfair' Election? The Experience of Russians." *Post-Soviet Affairs* 25, no. 2: 118–36. At https://doi.org/10.2747/1060-586X.24.2.118.

Rozenas, Arturas. 2016. "Office Insecurity and Electoral Manipulation." *Journal of Politics* 78, no. 1: 232–48. At https://doi.org/10.1086/683256.

Rundlett, Ashlea, and Milan W. Svolik. 2016. "Deliver the Vote! Micromotives and Macrobehavior in Electoral Fraud." *American Political Science Review* 110, no. 1: 180–97. At https://doi.org/10.1017/S0003055415000635.

Schwarz, Susanne, and Alexander Coppock. Forthcoming. "What Have We Learned about Gender from Candidate Choice Experiments? A Meta-Analysis of 67 Factorial Survey Experiments." *Journal of Politics*. At https://alexandercoppock.com/schwarz_coppock_2020.html, accessed November 20, 2020.

Simpser, Alberto. 2012. "Does Electoral Manipulation Discourage Voter Turnout? Evidence from Mexico." *Journal of Politics* 74, no. 3: 782–95. At https://doi.org/10.1017/S0022381612000333.

Simpser, Alberto. 2013. *Why Governments and Parties Manipulate Elections: Theory, Practice, and Implications*. New York, N.Y.: Cambridge University Press.

Svolik, Milan W. 2020. "When Polarization Trumps Civic Virtue: Partisan Conflict and the Subversion of Democracy by Incumbents." *Quarterly Journal of Political Science* 15, no. 1: 3–31. At https://doi.org/10.1561/100.00018132.

Szakonyi, David. Forthcoming. "Candidate Filtering: The Strategic Use of Electoral Manipulations in Russia." *British Journal of Political Science*. At https://doi.org/10.2139/ssrn.3438270.

Taber, Charles S., and Milton Lodge. 2006. "Motivated Skepticism in the Evaluation of Political Beliefs." *American Journal of Political Science* 50, no. 3: 755–69. At https://doi.org/10.1111/j.1540-5907.2006.00214.x.

Tucker, Joshua A. 2007. "Enough! Electoral Fraud, Collective Action Problems, and Post-Communist Colored Revolutions." *Perspectives on Politics* 5, no. 3: 535–51. At https://doi.org/10.1017/S1537592707071538.

Vicente, Pedro C. 2014. "Is Vote Buying Effective? Evidence from a Field Experi-

ment in West Africa." *Economic Journal* 124, no. 574: F356–87. At https://doi
.org/10.1111/ecoj.12086.

Wagner, Markus, Jessica Tarlov, and Nick Vivyan. 2014. "Partisan Bias in Opin-
ion Formation on Episodes of Political Controversy: Evidence from Great
Britain." *Political Studies* 62, no. 1: 136–58. At https://doi.org/10.1111/j.1467
-9248.2012.01002.x.

Weitz-Shapiro, Rebecca. 2014. *Curbing Clientelism in Argentina: Politics, Poverty,
and Social Policy*. New York, N.Y.: Cambridge University Press.

Wengle, Susanne, and Christine Evans. 2018. "Symbolic State-Building in Con-
temporary Russia." *Post-Soviet Affairs* 34, no. 6: 384–411. At https://doi.org/10
.1080/1060586X.2018.1507409.

Williamson, Scott. 2018. "Elections, Legitimacy, and Compliance in Authoritar-
ian Regimes: Evidence from the Arab World." At https://www.scott-william-
son.com/s/Elections-Autocracy-Legitimacy-Final3-dt9e.pdf, accessed Novem-
ber 20, 2020.

Authors

Ora John Reuter is an associate professor of political science at the University of
Wisconsin-Milwaukee and a senior research fellow at the National Research Uni-
versity Higher School of Economics, Russian Federation. He is the author of *The
Origins of Dominant Parties: Building Authoritarian Institutions in Post-Soviet Rus-
sia* (2017). His research focuses on authoritarianism and Russian politics. He can
be reached at reutero@uwm.edu.

David Szakonyi is an assistant professor of political science at George Washing-
ton University and a research fellow at the National Research University Higher
School of Economics, Russian Federation. He is the author of *Politics for Profit:
Business, Elections, and Policymaking in Russia* (2020). His research interests include
autocracy, corruption, and business-government relations. He can be reached at
dszakonyi@gwu.edu.

Acknowledgments

Funding

Key Words

FOREIGN AID AND STATE LEGITIMACY

Evidence on Chinese and US Aid to Africa from Surveys, Survey Experiments, and Behavioral Games

By ROBERT A. BLAIR and PHILIP ROESSLER

ABSTRACT

What are the effects of foreign aid on the perceived legitimacy of recipient states? Different donors adhere to different rules, principles, and operating procedures. The authors theorize that variation in these aid regimes may generate variation in the effects of aid on state legitimacy. To test their theory, they compare aid from the United States to aid from China, its most prominent geopolitical rival. Their research design combines within-country analysis of original surveys, survey experiments, and behavioral games in Liberia with cross-country analysis of existing administrative and Afrobarometer data from six African countries. They exploit multiple proxies for state legitimacy, but focus in particular on tax compliance and morale. Contrary to expectations, the authors find little evidence to suggest that exposure to aid diminishes the legitimacy of African states. If anything, the opposite appears to be true. Their results are consistent across multiple settings, multiple levels of analysis, and multiple measurement and identification strategies, and are unlikely to be artifacts of sample selection, statistical power, or the strength or weakness of particular experimental treatments. The authors conclude that the effects of aid on state legitimacy at the microlevel are largely benign.

INTRODUCTION

WHAT are the effects of foreign aid on the perceived legitimacy of recipient states? According to prominent accounts of state legitimacy in political science, the more citizens credit the state with providing services equitably and effectively, the more legitimate they will perceive the state to be. Legitimacy reflects the belief that the state has a "right to rule," and thus to impose restrictions and extractions on citizens.[1] The more legitimate citizens perceive the state to be, the more willing they will be to obey state laws and to pay state taxes, thus increasing the state's capacity to provide services. The result is a virtu-

[1] Gilley 2009; Levi, Sacks, and Tyler 2009.

ous circle in which service provision increases state legitimacy, and state legitimacy increases service provision.[2]

Scholars have long argued that aid and other forms of third-party service provision rupture this virtuous circle.[3] A large and established cross-national literature finds that aid has adverse effects on the behavior of government officials and the performance of government institutions, exacerbating corruption and diminishing the strength of the bureaucracy at the macrolevel.[4] This may decrease state legitimacy if citizens conclude that the state is too weak or corrupt to provide services equitably and effectively. But subnationally, recent research finds little evidence of similarly deleterious effects on state legitimacy as perceived by the citizens who experience aid most directly, such as by using or otherwise learning about donor-provided services.[5] At the microlevel, the effects of aid on state legitimacy in the eyes of citizens appear to be benign.

We hypothesize that one possible explanation for these null results lies in the particular types of donors that previous scholars have studied at the microlevel. Different donors adhere to different rules, principles, and operating procedures. We refer to these rules, principles, and procedures as aid regimes. Donors within the Organization for Economic Cooperation and Development, for example, typically follow guidelines issued by the OECD's Development Assistance Committee (DAC), of which the United States, the European Union, and twenty-eight other countries are members. These guidelines influence the types of projects funded by OECD donors; the regulations, restrictions, and good governance conditionalities that OECD donors impose on recipient states; and the transparency with which aid from OECD donors is disbursed. Although compliance with these guidelines is imperfect, their existence nonetheless ensures a degree of homogeneity across OECD aid regimes.

To our knowledge, all existing studies of aid and state legitimacy at the microlevel focus on OECD donors.[6] But in recent years, the OECD has faced increasingly fierce competition from "nontraditional" donors like Qatar, Saudi Arabia, and, most important, China. Over the past two decades, China has dramatically increased the amount of aid it provides to developing countries, especially in Africa.[7] China's rise as a donor and

---

[2] Schmelzle and Stollenwerk 2018.
[3] Bratton 1989; Bueno de Mesquita and Smith 2009; Gubser 2002; Fowler 1991; Moyo 2010; see Blair and Winters 2020 for a review.
[4] Bräutigam and Knack 2004; Knack 2004; Svensson 2000.
[5] Cruz and Schneider 2017; Dietrich, Mahmud, and Winters 2018; Dietrich and Winters 2015.
[6] Baldwin and Winters 2018; Dietrich, Mahmud, and Winters 2018; Dietrich and Winters 2015.
[7] For concision we use the word "aid" throughout, although there's some ambiguity in defining Chi-

investor has proven highly controversial. Critics argue that China is underwriting "a world that is more corrupt, chaotic, and authoritarian."[8] This perspective is shared by officials at the highest levels of the US government, who accuse China of undermining democracy and human rights in recipient countries.[9] In contrast, advocates—including some African heads of state—praise China for delivering aid more quickly, at lower cost, and with fewer strings attached than its OECD rivals do.[10]

The Chinese aid regime differs dramatically from that of most OECD donors. OECD donors tend to impose antibribery and anticorruption regulations on recipient countries, but China typically does not. Many OECD donors condition aid on the adoption of political or economic reforms; again, China typically does not. And although most OECD donors participate in international aid transparency initiatives, China usually doesn't. We argue that these features of the Chinese aid regime may have important implications for the legitimacy of recipient states, especially those in Africa. But these implications remain understudied and poorly understood, even as Chinese aid to Africa continues to grow and as research on the nature and effects of that aid is expanding.[11]

In a preanalysis plan (PAP),[12] we predicted that Chinese aid would diminish the legitimacy of recipient states, both in absolute terms and relative to aid from OECD donors. We theorized that the effects of aid on state legitimacy would depend on whether citizens give credit to donors or to recipient states for the services donors provide (attribution), whether they believe donor-provided services are distributed fairly and transparently (procedural legitimacy), and whether they perceive donor-provided services to be of high quality (outcome-based legitimacy). We hypothesized that Chinese aid would have adverse effects on state legitimacy because citizens would be less likely to credit recipient states with the services provided by China, less likely to believe Chinese-provided services are distributed fairly and transparently, and less likely

nese assistance to African countries, as we discuss in our description of the AidData data set.

[8] Naím 2007.

[9] E.g., David Smith, "Hillary Clinton Launches African Tour with Veiled Attack on China," *Guardian*, August 1, 2012, at https://www.theguardian.com/world/2012/aug/01/hillary-clinton-africa-china (accessed January 4, 2021).

[10] E.g., "John Magufuli: Tanzania Prefers 'Condition-Free' Chinese Aid," *BBC News*, November 27, 2018, at https://www.bbc.com/news/world-europe-46364342 (accessed December 28, 2020); Wade 2008.

[11] Bräutigam 2009; Dreher and Fuchs 2015; Dreher et al. 2018; Isaksson and Kotsadam 2018a; Isaksson and Kotsadam 2018b; Kelly, Brazys, and Elkink 2016.

[12] Our PAP was registered with the Evidence in Governance and Politics (EGAP) network prior to data collection, and is available at https://osf.io/qek4v. The PAP refers specifically to our behavioral games in Liberia, but our surveys and survey experiments are similar in approach, and our hypotheses are the same, regardless.

to perceive Chinese-provided services as high quality. Drawing on theories of retrospective voting,[13] we predicted that Chinese-provided services would generate disaffection not only with China, but also with the local government officials and institutions responsible for regulating China's actions—even when those actions are largely or entirely out of the government's control. We motivate these predictions in further detail below.

We test our hypotheses by combining multiple data sources and multiple measurement and identification strategies at multiple levels of analysis. We integrate cross-country analysis of data from Afrobarometer, AidData, and the Aid Information Management Systems of African finance and planning ministries with within-country analysis of original surveys, survey experiments, and behavioral games in Liberia, one of the world's most aid-dependent states. Each of these approaches complements and compensates for the limitations of the others. We triangulate between them. Although we test the effects of aid on multiple indicators of state legitimacy, we focus in particular on tax compliance—the extent to which citizens actually pay taxes—and tax morale—the extent to which citizens believe they have a duty to pay taxes.[14]

Tax *morale* captures the "value-based" component of legitimacy, defined as citizens' "sense of obligation or willingness to obey authorities." Tax *compliance* captures the behavioral component, defined as "actual compliance with governmental regulations and laws."[15] We focus on taxation because it's central to the virtuous circle of service provision and state legitimacy described above, and is also central to the more general challenge of "building and sustaining the power of states, and shaping their ties to society."[16] Not coincidentally, tax compliance and morale are among the most prominent and widely used proxies for state legitimacy in the social sciences.[17] Indeed, in many respects the "general problem of legitimacy" is best understood as a problem "exhibited in the authority of the tax collector."[18] To explore the robustness of our results, we also test the effects of aid on several additional proxies for state legitimacy.

Contrary to our expectations, we find little evidence that Chinese

[13] Achen and Bartels 2004; Achen and Bartels 2012; Gasper and Reeves 2011; Healy and Malhotra 2010; Healy, Malhotra, and Mo 2010.
[14] Luttmer and Singhal 2014.
[15] Levi, Sacks, and Tyler 2009, 356.
[16] Bräutigam, Fjeldstad, and Moore 2008, 1.
[17] Gilley 2009; Levi, Sacks, and Tyler 2009; Tyler 2006.
[18] Fain 1972, 20.

aid has diminished the legitimacy of recipient states in Africa, either in absolute terms or relative to aid from the US, China's most important OECD rival. In Liberia, we find that randomly assigned vignettes about donor-provided services have either null or, in the case of US aid, positive effects on tax compliance and morale as measured through surveys and behavioral games. Similarly, we find that Liberians who are most affected by US aid typically have more favorable rather than less favorable perceptions of the Liberian government. Exposure to Chinese aid is not associated with perceptions of the Liberian government one way or the other. Across other African countries as well, we use a spatial difference-in-differences estimator to show that both Chinese and US aid generally have null or even positive effects on perceptions of government after correcting for potential selection biases. These results are consistent across measurement and identification strategies, and, as we discuss below, are unlikely to be artifacts of measurement error, selection bias, or statistical power. We conclude by considering the implications of our results for the study and practice of providing aid in the future.

THEORETICAL FRAMEWORK

FOREIGN AID AND STATE LEGITIMACY

The most prominent accounts of state legitimacy posit that states gain legitimacy when citizens credit them with providing services, and with doing so equitably and effectively. Citizens who credit the state for the services it provides will view the state as legitimate, and will respond by taking actions that facilitate further service provision—most importantly, by paying taxes.[19] The result is a virtuous circle by which service provision increases state legitimacy and state legitimacy increases service provision.[20] This model, sometimes described as the "fiscal contract,"[21] is central to most studies of citizen/state relations in political science.

Two key features of the model are the quality of state-funded services and the fairness and transparency with which those services are allocated.[22] Legitimacy is thus a function of both outcomes and procedures.[23] It's not enough that citizens merely attribute services to the government (attribution). They must also perceive government service

[19] Levi 2006; Levi, Sacks, and Tyler 2009; Tyler 2006.
[20] Schmelzle and Stollenwerk 2018.
[21] Timmons 2005.
[22] Levi 2006, 5.
[23] Levi, Sacks, and Tyler 2009.

provision to be of high quality (outcome-based legitimacy) and fair and transparent (procedural legitimacy). Indeed, attribution can have de-legitimizing effects if services are delivered in a way that suggests government incompetence or inequity.[24] For example, if citizens perceive a government-built road as low quality, or the administrators at a government-run school as discriminatory and corrupt, then service provision may have the paradoxically adverse effect of decreasing state legitimacy.[25]

Scholars have long worried that foreign donors and other third parties might rupture the fiscal contract between citizens and the state, and thus break the virtuous circle of service provision and state legitimacy.[26] Underlying this concern is the intuition that if citizens see third parties supplying services, they'll conclude that government officials and institutions aren't fulfilling their end of the fiscal contract. This is a problem of attribution. By this logic, the dilemma for recipient countries isn't aid per se, but rather the risk that aid will cause citizens to attribute services to third parties. If governments could claim credit for donor-provided services—or if citizens (erroneously) gave governments credit for providing those services—then aid might bolster state legitimacy rather than erode it. This focus on attribution is common to almost all studies of aid and state legitimacy over the past decade, especially at the microlevel.[27]

Contrary to the concerns of skeptics, these studies have almost uniformly found that aid does not have de-legitimizing effects. As is appropriate for studies focused on the attribution problem, most researchers explain this null (or in some cases even positive) effect of aid on state legitimacy as a byproduct of undeserved credit claiming. By these accounts, citizens credit governments not only for government-provided services, but for donor-provided services as well. This occurs either because citizens don't have enough information to assign credit accurately, or because they interpret donor-provided services as a "signal of [politicians'] ability to extract resources from donors for the benefit of their communities."[28] Either way, the more aid the state receives, the more legitimate citizens perceive it to be.

But because existing studies focus on the attribution problem, they capture only part of the dynamic that's at work when donors provide

[24] Mcloughlin 2015.
[25] Mcloughlin 2018.
[26] Bratton 1989; Bueno de Mesquita and Smith 2009; Gubser 2002; Fowler 1991; Moyo 2010.
[27] Baldwin and Winters 2018, 4; Cruz and Schneider 2017, 399; Dietrich, Mahmud, and Winters 2018, 136; Dietrich and Winters 2015, 165, 170; Guiteras and Mobarak 2014, 6; Sacks 2012, 1.
[28] Cruz and Schneider 2017, 398.

services. We follow the influential model of state legitimacy described above and argue that aid will have (de-)legitimizing effects depending not only on (1) whether citizens credit their own government with the services provided by donors (attribution), but also on (2) the quality of the services (outcome-based legitimacy) and (3) the fairness and transparency with which they're provided (procedural legitimacy). Previous studies address this first source of (il)legitimacy (attribution), but implicitly assume away the second two (effectiveness and equity). Cesi Cruz and Christina Schneider make this assumption explicit when they write, "people in communities which obtain foreign aid projects tend to be favorably disposed to them," an assumption they describe as "straightforward and uncontroversial."[29] Our theory begins with the proposition that citizens' attitudes toward donor-provided services aren't always as straightforward and uncontroversial as they may seem.

Just as citizens may credit their own governments (perhaps undeservedly) for benefits provided by donors, so may they blame their own governments (perhaps undeservedly) for any damage caused by donors. Undeserved blame attribution is the inverse of undeserved credit claiming, and it's a phenomenon that extends well beyond the politics of aid. For example, scholars of retrospective voting have shown that voters will punish incumbents for events that are largely or entirely outside the control of elected officials, such as shark attacks,[30] floods,[31] and even the outcomes of college football games[32]—although voters' reactions to these events also depend in part on the competence of the government's response to them.[33]

In the context of aid, although governments can't control donors' actions, they can negotiate agreements with donors that maximize the benefits to citizens and minimize risks. They can also enforce regulations to ensure that aid is delivered fairly, transparently, and effectively. Citizens may punish recipient governments for allowing donors to operate in ways that are detrimental to their communities, even when recipient governments aren't directly responsible for the harm inflicted by donors. In this way, dissatisfaction with donor-provided services may result in disapproval not only of the donor itself, but also of the government officials and institutions tasked with negotiating and regulating the donor's operations. This effect may be especially stark among urban populations, in which residents tend to be better informed and

[29] Cruz and Schneider 2017, 398.
[30] Achen and Bartels 2012.
[31] Gasper and Reeves 2011.
[32] Healy, Malhotra, and Mo 2010.
[33] Healy and Malhotra 2010.

less supportive of incumbent governments,[34] and thus potentially more sensitive to the government's perceived failures in overseeing donor activities.

AID REGIMES

As noted above, we use the term *aid regimes* to refer to the different rules, principles, and operating procedures followed by different donors. Aid regimes vary along multiple dimensions, including *project type* (whether donors specialize in projects of a certain size, scope, or sector); *standards* (whether donors use benchmarks of project success); *conditionalities* (whether donors make aid conditional on political or economic reforms); *regulations* (whether donors impose restrictions on labor, corruption, or the environment); *procurement* (whether aid is tied to the procurement of goods or services from the donor country); and *transparency* (whether donors publicly disclose the terms of aid agreements). We argue that variation along these dimensions may help to explain differences in the quality, fairness, and transparency of donor-provided services. Such differences, in turn, may help to explain variation in the (de-)legitimizing effects of aid across donors.

Although we conceptualize aid regimes as a country-level phenomenon, groups of countries often follow very similar aid regimes. Consider, for example, the guidelines published by the OECD's Development Assistance Committee. The DAC explicitly encourages OECD donors to invest in political parties and other "accountability actors and institutions."[35] This guidance is intended to ensure that OECD donors coordinate around certain project types. The DAC also publishes criteria to help OECD donors set "standards of success" related to project relevance, coherence, effectiveness, efficiency, impact, and sustainability.[36] These standards are designed to ensure the quality of OECD-provided projects. The DAC also maintains a set of antibribery and anticorruption safeguards that OECD donors are expected to impose on the funds they disburse.[37] These regulations are meant to ensure fairness and transparency in the way OECD-provided projects are allocated and administered. In addition, the DAC regulates procurement for OECD-provided projects by urging OECD donors to "untie" the aid they give, so that recipient

---

[34] Harding 2010; Koter 2013.

[35] See http://www.oecd.org/dac/conflict-fragility-resilience/governance/docs/4312181e.pdf (accessed December 30, 2020).

[36] See https://www.oecd.org/dac/evaluation/daccriteriaforevaluatingdevelopmentassistance.htm (accessed December 30, 2020).

[37] See https://www.oecd.org/corruption/anti-bribery/Recommendation-Development-Cooperation-Corruption.pdf (accessed January 7, 2021).

governments can procure goods and services from the lowest bidders.[38]

Inevitably, compliance with these guidelines isn't perfect, but their existence does ensure a degree of homogeneity across the aid regimes of OECD donors. We argue that the shared features of these aid regimes may have significant implications for the effects of OECD-provided projects on the legitimacy of recipient states. Importantly, we don't argue that citizens necessarily care or even know about the general macrolevel attributes of aid regimes. Rather, we argue that these attributes shape the specific microlevel characteristics of donor-provided projects—which, in turn, shape citizens' attitudes toward those projects and, more important, toward the government officials and institutions responsible for negotiating and regulating them.

For example, aid to "accountability actors and institutions" is likely to be invisible to most citizens in recipient countries, which creates opportunities for government officials to claim credit for any subsequent improvements in government performance. The use of evaluation criteria may increase the likelihood that OECD-provided projects are high quality, while the imposition of antibribery and anticorruption safeguards may increase the likelihood that the distribution of OECD funds is fair and transparent. These attributes of the OECD aid regime make undeserved credit claiming more probable, and undeserved blame attribution less probable. Citizens may credit the government (perhaps undeservedly) for services provided by OECD donors. But even if they don't, they may nonetheless credit the government (perhaps undeservedly) for ensuring that OECD-provided services are high quality and that OECD service provision is fair and transparent.

Although existing studies have taught us much about the relationship between aid and state legitimacy, they're almost uniformly designed in ways that preclude comparisons across aid regimes—comparisons that might further illuminate the (de-)legitimizing effects of aid on recipient states. Most studies of aid and state legitimacy test the impact of aid from a single donor, such as Japan,[39] the US,[40] or the World Bank.[41] Other studies aggregate aid from multiple donors into a single summary index, obscuring potential differences between aid regimes.[42] Fewer studies compare the effects of aid from different donors, and even these studies tend to focus on OECD donors that follow similar sets

[38] See https://one.oecd.org/document/DCD/DAC(2018)33/FINAL/en/pdf (accessed January 4, 2021).
[39] Baldwin and Winters 2018.
[40] Dietrich, Mahmud, and Winters 2018.
[41] Cruz and Schneider 2017.
[42] Böhnke and Zürcher 2013; Böhnke, Koehler, and Zürcher 2010.

of rules, guidelines, and standards—for example, Canada and the US.[43] To our knowledge, no previous study has compared the effects of different aid regimes on the legitimacy of recipient states.

Differences between Chinese and OECD Aid Regimes

We advance the study of aid and state legitimacy by comparing the impact of aid from the largest and most prominent OECD donor—the US—to aid from the OECD's largest and most prominent nontraditional rival—China. The American and Chinese aid regimes are often perceived as competitors, especially in Africa. High-ranking members of the US government regularly warn African citizens and heads of state that the "empty promises" of Chinese donors breed "corruption and dependency,"[44] while the Chinese government explicitly distinguishes its aid regime from that of the US and other OECD donors.[45]

We highlight differences across five aid regime dimensions—project type, conditionalities, regulations, procurement, and transparency—that can have especially important implications for the legitimacy of recipient states. We summarize these differences and their implications in Table 1. First, whereas OECD donors typically fund a variety of relatively small-scale projects across multiple sectors, China tends to favor fewer larger-scale projects focused particularly on infrastructure (row one): transportation, communication, electricity, and so on. China explicitly promotes this focus as one of the defining characteristics of its aid regime, perhaps best embodied in its Belt and Road Initiative—the construction of transportation and energy infrastructure along corridors roughly corresponding to the historical Silk Road, accompanied by new trade deals and other agreements. Many developing countries need these higher-cost, higher-risk ventures, but more-traditional donors are often reluctant to undertake them.

Second, while many OECD donors attach political or economic conditionalities (row two) to the aid they provide, China does not. Indeed, the Chinese government views conditionalities of this sort as a threat to the sovereignty of recipient states. This position was most clearly ar-

---

[43] Dietrich and Winters 2015.
[44] Robbie Corey-Boulet, "Pompeo Closes Africa Tour with Warning about China's 'Empty Promises,'" *Yahoo News*, February 19, 2020, at https://news.yahoo.com/pompeo-closes-africa-tour-warning-chinas-empty-promises-101611608.html (accessed January 4, 2021). See also former US National Security Advisor John Bolton's comments at the Heritage Foundation in December 2018, where he accused China of engaging in "predatory practices" that "stunt economic growth in Africa" and "inhibit opportunities for US investment," at https://www.whitehouse.gov/briefings-statements/remarks-national-security-advisor-ambassador-john-r-bolton-trump-administrations-new-africa-strategy (accessed January 4, 2021).
[45] See, for example, "Full Text of Chinese President Xi Jinping's Speech at Opening Ceremony of 2018 FOCAC Beijing Summit," at http://www.xinhuanet.com/english/2018-09/03/c_129946189.htm (accessed January 7, 2021).

TABLE 1

DIFFERENCES BETWEEN CHINESE AND OECD–DAC AID REGIMES

Aid Regime Dimensions	*Attributes of OECD–DAC Aid Regime*	*Attributes of Chinese Aid Regime*	*Predicted Effects of Chinese Aid Regime on State Legitimacy*
Project types	smaller projects across many sectors	larger projects focused on infrastructure	recipient state less likely to claim credit
Conditionalities	political and economic "good governance" conditionalities	no political or economic conditionalities	citizens less likely to perceive quality, fairness, or transparency
Regulations	more antibribery and anticorruption regulations	fewer antibribery and anticorruption regulations	citizens less likely to perceive quality, fairness, or transparency
Procurement	aid not tied to procurement from donor country	aid tied to procurement from donor country	recipient state less likely to claim credit
Transparency	aid tracked through international transparency regimes	aid not tracked through international transparency regimes	citizens less likely to perceive quality, fairness, or transparency

ticulated by President Xi Jinping at the 2018 Forum on China-Africa Cooperation (FOCAC), where he laid out China's "five-no" approach to aid—which includes "no interference in African countries' pursuit of development paths that fit their national conditions," "no attachment of political strings to assistance to Africa," and "no interference in African countries' internal affairs."[46] This principle is perhaps the most defining feature of the Chinese aid regime.

Third, while OECD donors typically impose regulations (row three) to mitigate corruption, environmental degradation, and abusive or exploitative labor practices, China usually does not. For example, guided by its belief in noninterference, China generally doesn't police the taking or giving of bribes. The absence of these safeguards may explain the finding that Chinese-funded projects tend to increase local corruption, while World Bank–funded projects do not.[47] But China's lax regulations do reduce "up-front oversight in the form of planning meetings and im-

[46] The five-no approach also includes "no imposition of our will on African countries" and "no seeking of selfish political gains in investment and financing operation with Africa." See "Full Text of Chinese President Xi Jinping's Speech at Opening Ceremony of 2018 FOCAC Beijing Summit," at http://www.chinadaily.com.cn/a/201809/04/WS5b8d5c25a310add14f389592 (accessed January 7, 2021).

[47] Brazys, Elkink, and Kelly 2017; Isaksson and Kotsadam 2018a.

pact assessments," allowing it to implement projects more quickly and more cheaply than traditional donors.[48] Although efficiency is not as explicit a component of China's aid regime as noninterference, it has been a priority since at least 1964, when former premier Zhou Enlai first outlined China's "Eight Principles for Economic Aid and Technical Assistance to Other Countries"—principles that promised "quicker results" for recipient states.[49] African heads of state often cite efficiency as a difference between the Chinese aid regime and its OECD rivals.[50]

Fourth, while the DAC explicitly endorses untying aid (row four), China often ties projects to Chinese companies and products and typically relies on large numbers of expatriate Chinese workers (alongside local labor) for project implementation.[51] Given China's reluctance to impose labor regulations, Chinese contractors face few constraints on wages or hiring and management practices in recipient countries beyond the local labor laws that recipient governments may or may not enforce.

Fifth, China doesn't participate in the International Aid Transparency Initiative or in the other transparency regimes to which most OECD donors adhere (row five). As a result, there's little publicly available information about Chinese aid flows to Africa and to other parts of the developing world. This has implications both for the perceived fairness and transparency of the Chinese aid allocation process, and also for our empirical strategy—a point to which we return below.

We argue that these features of the Chinese aid regime are likely to affect the perceived legitimacy of recipient states along all three of the dimensions described above: credit claiming (attribution), quality (outcome-based legitimacy), and fairness and transparency (procedural legitimacy). On the first dimension, China's specialization in large-scale infrastructure projects may limit the opportunities for government officials to claim credit. These are precisely the sorts of projects usually financed through tax revenues, so citizens who learn about China's role in funding them may conclude that their own governments are reneging on the fiscal contract. This problem is likely to be compounded by China's procurement rules, which tie aid to the use of Chinese contractors and expatriate Chinese workers. Governments are less likely to be

[48] Hanauer and Morris 2014, 57.
[49] See "The Chinese Government's Eight Principles for Economic Aid and Technical Assistance to Other Countries," at https://digitalarchive.wilsoncenter.org/document/121560 (accessed January 7, 2021).
[50] E.g., Wade 2008.
[51] Yun Sun, "China's Aid to Africa: Monster or Messiah?" Brookings. February 7, 2014, at https://www.brookings.edu/opinions/chinas-aid-to-africa-monster-or-messiah/ (accessed January 4, 2021).

able to claim credit for projects that are so clearly a result of Chinese funds and labor.

But while recipient states may be less likely to claim credit for benefits provided by China, they may not escape blame for any harm that China causes. On the second dimension, although recipient governments may value the lack of regulations and the corresponding speed with which Chinese aid is disbursed, China is frequently accused of prioritizing speed over quality. Local and international media regularly publish reports of shoddy workmanship in Chinese-funded projects. Whether or not these reports are representative, they've clearly influenced public opinion: according to nationally representative surveys across thirty-six African countries, Africans cite the low quality of Chinese workmanship as the most important factor contributing to negative perceptions of China on the continent.[52] Citizens who share these views may conclude that Chinese aid won't actually improve their welfare, and—crucially—they may blame their own governments for setting "poor quality requirements."[53]

On the third dimension, although recipient governments may appreciate receiving aid with no strings attached, citizens may worry that a lack of conditionalities, regulations, and transparency will encourage foul play. For example, without regulations on unfair or exclusionary labor practices, Chinese contractors are often accused of paying domestic workers less than Chinese expatriates and less than the local minimum wage. Some Chinese contractors have also been accused of engaging in abusive management practices and of undermining collective bargaining.[54] This may foment grievances not only against the contractors themselves, but also against the government agencies responsible for regulating them. As a recent RAND Corporation report notes, "in some cases blame must be laid on African governments for failing to enforce their own labor laws."[55]

Of course, the schematic in Table 1 is highly stylized. Each feature of the Chinese aid regime may affect the legitimacy of recipient states through multiple mechanisms and along multiple dimensions simultaneously, making it difficult if not impossible to disentangle the dimensions empirically. For example, a lack of transparency may increase the risk of corruption during project implementation (procedural legitimacy), which may in turn diminish the quality of the project itself

[52] Lekorwe et al. 2016.
[53] Hanauer and Morris 2014, 62.
[54] Isaksson and Kotsadam 2018b.
[55] Hanauer and Morris 2014, 65.

(outcome-based legitimacy). In addition, the features of the Chinese aid regime are not mutually exclusive. For example, failures of regulation may also be perceived as failures of transparency. Our aim in Table 1 is simply to show that the Chinese aid regime differs from its OECD counterpart in multiple ways, and that these differences can have important implications for the effects of Chinese aid on the legitimacy of recipient states.

HYPOTHESES

Our theoretical framework implies that aid should have negative effects on state legitimacy when donors provide services that (1) governments can't easily claim credit for, (2) citizens perceive as low quality, or (3) citizens believe were allocated unfairly or nontransparently. In our pre-analysis plan (PAP), we hypothesized that the distinctive features of the Chinese aid regime would make it especially likely to have such negative effects, both in absolute terms and relative to aid from more traditional donors like the US. We predicted that Chinese aid would be more likely to have negative effects because citizens would be less likely to credit their own government officials for highly visible, large-scale infrastructure projects implemented by Chinese contractors and expatriate Chinese workers (attribution); because citizens would be more likely to view hastily implemented projects as low quality (outcome-based legitimacy); and because citizens would be less likely to perceive Chinese aid as fair and transparent due to China's reluctance to impose conditionalities and regulations and to participate in international aid transparency regimes (procedural legitimacy).

Our goal was not to isolate these mechanisms—a challenge even in experimental research[56]—but rather to estimate the direct effects of Chinese aid on state legitimacy while minimizing threats to inference by triangulating across multiple measurement, identification, and estimation strategies. We opted to test the impact of Chinese aid at the microlevel because existing evidence suggests that Africans' perceptions of Chinese aid are shaped less by its "political and macroeconomic implications" and more by their own "experience or exposure to Chinese economic and commercial activities."[57] But our results provide little to no evidence in support of our predictions. To be consistent with our PAP, and in the spirit of preregistration, we don't retrofit our theory to accommodate our results. Rather, we test our hypotheses as prespecified and we consider potential explanations for our findings, as well as

[56] Bullock, Green, and Ha 2010.
[57] Hanauer and Morris 2014, 60–61.

their theoretical and substantive implications, in the conclusion of this article.

RESEARCH DESIGN

To test our predictions, we combine within-country analysis of original surveys, survey experiments, and behavioral games in Liberia with cross-country analysis of data from Afrobarometer, AidData, and the Aid Information Management Systems (AIMS) of African finance and planning ministries.

WITHIN-COUNTRY RESEARCH DESIGN

Our within-country research design focuses on Liberia, a small West African nation still recovering from fourteen years of civil war and, more recently, the devastating Ebola epidemic of 2014–2015. Liberia is one of the world's least developed and most aid-dependent countries. It placed 177th out of 188 countries on the 2016 UN Human Development Index.[58] In 2015, the year we conducted our research, the OECD ranked Liberia as the second most aid-dependent country in the world, behind only the island nation of Tuvalu.[59] Tax evasion has long been pervasive in Liberia, and government resources have always been severely constrained. Underlying these problems is the nearly ubiquitous perception that government institutions are corrupt and inept—a perception that exacerbated, and was exacerbated by, the Ebola crisis.[60]

In many ways Liberia is a most likely case for detecting the adverse effects of Chinese aid on state legitimacy. Anecdotally, China's rapidly expanding presence in Liberia has stoked grievances and catalyzed civil unrest. Liberian consumers have complained of the low quality of Chinese-provided projects.[61] Liberian workers have accused Chinese contractors of "modern slavery," and have appealed to government officials to "urgently intervene before the situation turns worse."[62] A damning 2014 report by the Sustainable Development Institute observed that the Liberian government is widely believed to be "more concerned with facilitating a smooth operating environment" for Chinese firms than with

---

[58] See http://hdr.undp.org/en/content/human-development-index-hdi (accessed January 4, 2021).
[59] OECD 2015.
[60] Blair, Morse, and Tsai 2017.
[61] Rodney D. Sieh, "Liberia: Low Quality, High Volume—Dark Side of Chinese Contractors," *Front Page Africa*, June 17, 2014, at http://allafrica.com/stories/201406170710.html (accessed January 4, 2021).
[62] *New Dawn*, "Liberia: China Union Employees Claim Modern Slavery," October 8, 2013, at http://allafrica.com/stories/201310080605.html (accessed December 29, 2020).

ensuring the welfare of Liberian citizens.[63] Consistent with our theoretical framework, 75 percent of Liberians in our rural survey sample (described below) agree or strongly agree that the Liberian government deserves most of the blame for the bad things done by donors, compared to just 20 percent who agree or strongly agree that donors deserve most of the blame. The discrepancy is even starker in our urban survey sample (82 percent versus 17 percent). These descriptive statistics provide evidence of (potentially undeserved) blame attribution at work.

BEHAVIORAL GAMES

Our within-country research design begins with a modified version of the tax compliance game, a staple of behavioral economics.[64] In the conventional setup of the game, each participant receives an initial endowment—their "income"—and decides how much of that endowment to report to the tax "authorities." Reported income is taxed at a constant rate. Unreported income isn't taxed, but it's subject to the possibility of an audit. Audited participants pay a fine on unreported income. Participants decide how much to report by maximizing their expected utility over four parameters: the size of their initial endowment, the tax rate, the probability of being audited, and the size of the fine.

In August 2015, we implemented a modified version of the game with a convenience sample of 740 residents of Gbarnga, Liberia's second largest city and the capital of Bong County. Participants weren't selected randomly, but Liberian facilitators were instructed to recruit from all neighborhoods in the city and to sample men and women, youths and elders, and members of majority and minority ethnic groups. Facilitators also administered prescreening questions to test for literacy and numeracy prior to recruitment. This likely made the sample less representative—our participants were younger and more educated than the average Gbarnga resident, and more likely to be male—but non-random selection was necessary to ensure comprehension of the game. It's not obvious whether or how we should expect these features of the sample to affect our results, and in ancillary analyses we find no evidence of heterogeneous treatment effects by age, gender, education, or any of the other covariates we measure. Nonetheless, although convenience sampling is common in lab-in-the-field experiments, our results should be interpreted with this caveat in mind. The game was implemented over eleven days at the offices of Parley Liberia, a local NGO.

Upon arrival at the study site, participants were given unique iden-

---

[63] Mukpo 2014, 8.
[64] Alm, Jackson, and McKee 1992.

tification numbers and randomly assigned to one of four rooms, each corresponding to a different treatment group, with eight participants and two facilitators per room. As we show in Section A.4 of the supplementary material,[65] treatment groups are balanced along covariates. To guarantee privacy, tall wooden dividers separated participants from each other and from the facilitators. Before the game began, facilitators administered an informed-consent form and a short survey. They then described the rules of the game, following a detailed script and stopping repeatedly to answer questions. We reproduce the script in Section A.11 of the supplementary material. To illustrate the game's dynamics, facilitators guided the participants through three examples and three practice rounds, again stopping repeatedly to answer questions. We didn't conduct formal manipulation checks, but we prioritized comprehension, and the three examples and three practice rounds provided ample opportunity for facilitators to correct misunderstandings.

After the practice rounds, facilitators read the treatment group participants one of three vignettes about services provision in Liberia, focusing either on the Liberian government or on Chinese or American donors. Control group participants didn't hear a vignette. We chose to use American donors as a benchmark because of the traditional role of the US as a purveyor of aid to Africa, and because of Liberia's deep historical ties to the US. (Liberia was founded by freeborn African Americans and emancipated slaves in the nineteenth century.) The vignettes emphasized that (1) the government uses Liberian tax dollars to provide services while foreign donors do not; (2) the government (or Chinese or American donors) plays a vital role in service provision; and (3) the government (or Chinese or American donors) may nonetheless mismanage funds or mistreat Liberian workers. We reproduce the text of the vignettes in Section A.12 of the supplementary material.

The vignettes addressed both positive and negative aspects of government- and donor-provided services. This bundled treatment more accurately captures the conflicting signals that citizens receive in the real world, and allowed us to avoid generating ill will toward donors or the Liberian government as a byproduct of treatment—an important ethical consideration in and of itself. We address this point in further detail below. Given that knowledge of and exposure to both Chinese and US aid are common in Liberia,[66] we designed the vignettes to prime

[65] Section A.4 and all supplementary material can be found in Blair and Roessler 2021b.

[66] Forty-eight percent of respondents in our rural sample and 86 percent in our urban sample reported some form of exposure to Chinese aid, whether as consumers of Chinese-funded services or as labor for Chinese contractors (or both). Forty-three percent of rural respondents and 59 percent of urban respondents reported some form of exposure to US aid.

respondents to consider issues with which they were likely already familiar rather than to provide new information (although for some respondents the information may have been new).

After listening to the vignettes, participants began the live rounds of the game. For each round, participants were given an opaque manila envelope containing between zero and two hundred fake Liberian dollars (LD), redeemable for real LD at the end of the game. (Two hundred LD is roughly equivalent to US $2.75—the average daily wage for Gbarnga residents at the time.) They were also given a template for anonymously reporting their income and instructed to complete it and return it to the facilitators in each round. Reported income was taxed at a flat rate of 25 percent. After deciding how much income to report, participants took turns privately drawing colored beans from a bag. Those who drew a black bean were audited; those who drew a white bean were not. The probability of being audited was announced beforehand, and was held constant at 10 percent.[67] For ease of comprehension, audited unreported income was fined at a flat rate of one hundred LD.[68] The game was played over ten rounds, but to minimize anticipation effects, participants weren't told in advance how many rounds there would be.

After the ten rounds had been completed, facilitators conducted focus groups in which they asked participants six open-ended questions about their decision-making process during the game. (The text of the questions can be found in Section A.13 of the supplementary material.) Focus-group discussions were recorded, transcribed, and coded, allowing us to explore the internal and external validity of the game. Like all lab-in-the-field experiments, our tax compliance game was stylized. To minimize the gap between the lab and the field, facilitators informed participants at the outset that any taxes they paid would be delivered to the Liberian government, thus encouraging them to view their decisions as realistic and consequential. Results from the focus groups suggest that participants made this connection. When asked how they decided how much income to report, participants responded that they view tax compliance as necessary to "create the business-social network

[67] We set the probability of an audit relatively low in the hope of making the game more realistic. We weren't able to find data on the true probability of being audited in Liberia, but we suspect it's low (probably even less than 10 percent). It's possible that participants would have changed their behavior if we'd set the probability higher, though we see little reason to believe this would have affected our treatment effect estimates, since the probability of an audit was held constant across treatment groups.

[68] Given that the penalty associated with audits was fixed, rational participants should have reported either all or none of their income. These were indeed the two most common choices. The fixed nature of the penalty was a subtlety of our setup, however, and the degree of underreporting varied across respondents and rounds.

we are all yearning for" (focus group number 1.1.N.904), to "construct buildings and help our country build roads" (1.2.N.910), and to "be fair to the government" (10.2.U.921). Across the focus groups, participants almost invariably referred to the dynamics of taxation in the real world as motivation for their decisions. We provide additional excerpts from the focus groups in Section A.14 of the supplementary material.

SURVEYS AND SURVEY EXPERIMENTS

We supplemented the tax compliance game with surveys and survey experiments implemented in Gbarnga and in thirty-eight rural towns and villages throughout Bong, Lofa, and Nimba counties. We focused on these three counties because they have been priorities for both economic development and state consolidation in the postconflict period, and because they host a relatively large number of Chinese-funded projects, including the controversial Bong Mines iron ore project.[69] The three counties thus represent an important test case for the effects of Chinese aid on state legitimacy in Liberia. We implemented the survey in both urban and rural settings because urban populations in Africa tend to support incumbent governments at lower rates than do rural residents,[70] which we expected might make urban dwellers especially sensitive to (perceived) failures of government policy toward foreign donors.

The surveys were implemented between September and December in 2015. For the urban survey, we randomly selected 193 residents of nine randomly selected Gbarnga neighborhoods using the random walk technique, described in Section A.1 of the supplementary material. These respondents are broadly representative of Gbarnga. Although we can't say for certain how representative Gbarnga is of other Liberian cities, a comparison to a representative survey of Monrovia, the Liberian capital, conducted around the same time[71] suggests that our respondents are similar to Liberians in other urban settings. For example, the average age of respondents in our sample is thirty-nine, compared to thirty-seven in Monrovia. Ten percent of our respondents have only a primary school education, compared with 9 percent in Monrovia. Sixty-three percent of our respondents are female, compared to 57 percent in Monrovia. Ninety percent of our respondents are Christian, compared to 85 percent in Monrovia. The only marked dif-

[69] Alvin Worzi, "Bong Mines Residents Frustrated by China's Woes." *Daily Observer,* April 13, 2017, at https://www.liberianobserver.com/news/bong-mines-residents-frustrated-by-china-unions-woes/ (accessed December 29, 2020).
[70] Harding 2010; Koter 2013.
[71] Blair, Morse, and Tsai 2017.

ference is the proportion of our respondents who are farmers (15 percent, compared to 1 percent in Monrovia). This isn't surprising, given that the agricultural sector is much smaller in Monrovia than anywhere else in the country. We discuss these comparisons in more detail in Section A.3 of the supplementary material.

For our rural survey, we sampled eighteen randomly selected residents of each of the thirty-eight communities, again using the random walk technique. These respondents are broadly representative of the communities from which they were sampled, but the communities themselves aren't representative of Liberia or of the counties in which they're located. Rather, the communities were the control group for an unrelated field experiment involving the Liberian National Police.[72] Eligibility for that experiment was restricted to communities located near a usable road and with at least five hundred residents. Our sample conforms to these criteria. Although these criteria may affect the generalizability of our results, residents of these relatively accessible communities are also more likely to be affected by donor-provided projects, making their responses especially informative for our purposes.

Moreover, comparison to a nationally representative survey from 2011[73] suggests that our thirty-eight communities don't differ dramatically from the average Liberian town or village, either in the three counties covered in our survey or in the country as a whole. For example, the average age of respondents in our sample is forty, compared to thirty-seven nationwide. Eighty-six percent of our respondents are Christian, the same proportion as nationwide. Thirty percent of our sample has no education, compared to 35 percent nationwide. Sixty-one percent of our respondents work in agriculture, compared to 43 percent nationwide, but compared to 72 percent in Bong, 73 percent in Lofa, and 60 percent in Nimba. We discuss these comparisons in further detail in Section A.3 of the supplementary material. Although there are some disparities, the communities in our sample appear to be generally comparable to towns and villages across Liberia.

Both surveys included the same questions to measure respondents' prior exposure to Chinese and US aid, including whether they (1) could name or (2) had used any Chinese- or US-provided projects, and also whether (3) they or (4) any of their friends or family members had worked for a Chinese- or US-funded contractor. We aggregate responses to the first two questions into a dummy for *consumers* of donor-provided services, and the latter two into a dummy for *labor* for

[72] Blair, Karim, and Morse 2019.
[73] Vinck, Pham, and Kreutzer 2011.

donor-funded contractors. Liberian citizens don't randomly select into either of these forms of exposure to aid, but fortunately for our analysis (although not for Liberia), access to services is very limited in rural areas, and rural Liberians who need these services typically don't have multiple options from which to choose. The unemployment rate outside of subsistence agriculture is also very high—67 percent in our sample—and few rural Liberians have the luxury of declining a job offered by a potentially inequitable employer. Still, it's possible that Liberians who use donor-provided services or work for donor-funded contractors differ from those who don't along dimensions we can't control for, and that may correlate with attitudes toward government.

To overcome this problem, both surveys also included a survey experiment in which respondents were randomly assigned to hear truncated versions of the vignettes used in the tax compliance game. We reproduce the text of the vignettes in Section A.8 of the supplementary material. Respondents were then asked whether they believe Liberians have an obligation to pay taxes even if (1) donors provide most services, (2) the government "eats" some of the taxes it collects, (3) Liberian taxpayers are very poor, and (4) the government makes bad policies. We added these conditional clauses to reduce the risk of social desirability bias: by offering respondents a possible justification for tax evasion, we hoped to alleviate any pressure to provide affirmative answers. A control group was asked the same four questions without the vignette. Treatment in the survey experiments was assigned on the spot. As a result, although the probability of assignment to each condition was fixed at 0.25, the sizes of the treatment groups varied. The treatment groups are balanced on covariates, as shown in Section A.4 of the supplementary material.

CROSS-COUNTRY RESEARCH DESIGN

For our cross-country analysis we combine Afrobarometer data with information on Chinese aid from AidData[74] and information on US aid from the AIMS of African finance and planning ministries.[75] AidData tracks aid from nontraditional donors like China, Russia, and the Gulf states. Because many of these donors don't participate in international aid transparency regimes, details of the aid they provide are generally unknown. AidData fills this gap with information scraped from newspaper and academic articles, radio and television transcripts, the

---

[74] AidData Research and Evaluation Unit 2017; Bluhm et al. 2018.
[75] See the AIMS geocoded research releases at https://www.aiddata.org/datasets (accessed January 4, 2021).

websites of recipient governments, and NGO reports. AIMS captures aid from various traditional donors, including the US. We merge these data sets with rounds two through six of the Afrobarometer survey, which AidData geocoded to the community (village) level.[76] Our analysis covers the six countries for which Afrobarometer, AidData, and AIMS data are all available: Burundi, Malawi, Nigeria, Senegal, Sierra Leone, and Uganda. (In Section B.9 of the supplementary material, we drop US aid and expand our analysis to cover all thirty-eight Afrobarometer countries. Our conclusions about the effects of Chinese aid on state legitimacy do not change.)

To operationalize exposure to aid, we code whether each Afrobarometer respondent lives within thirty kilometers of at least one Chinese or US project, the narrowest bandwidth for which precise geolocation data are available. Since aid isn't randomly distributed, either within or across countries, we use a spatial difference-in-differences estimator to compare respondents living near planned (future) projects to those living near projects that were completed by the time of the survey, as in analyses by Ryan Briggs[77] and by Ann-Sofie Isaksson and Andreas Kotsadam.[78] Our identifying assumption is that the same selection process underlies the location of planned and completed projects, such that respondents living near planned projects serve as valid counterfactuals for those living near completed ones. If this assumption holds, we can estimate the net effect of the services actually provided by subtracting the coefficient on planned projects from the coefficient on completed ones.[79]

One potential threat to this assumption is that recipient governments may prioritize projects that are sited in politically important areas. If these projects are more likely to be completed than those sited elsewhere, then different selection processes underlie project planning and completion, potentially biasing our results. We view this risk as relatively minor, as we see no reason to expect recipient governments to prioritize in this way during the implementation phase but not during the planning phase.

---

[76] BenYishay et al. 2017.

[77] Briggs 2019.

[78] Isaksson and Kotsadam 2018a; Isaksson and Kotsadam 2018b.

[79] As we discuss in Section B.2 of the supplementary material, some respondents live near more than one completed or planned project. Following most spatial difference-in-differences analyses of this sort (Briggs 2019; Isaksson and Kotsadam 2018a, Isaksson and Kotsadam 2018b), we focus on estimating the effect of Chinese and US projects on the extensive rather than the intensive margin. In Section B.7, we show that our conclusions are unchanged when we instead use the number of projects within a thirty-kilometer radius of each respondent. In addition, some respondents live near both completed and planned projects. In Section B.8, we show that our conclusions are unchanged when we drop these respondents.

In other words, if governments favor politically important areas, then intuitively we should expect this favoritism to manifest during project planning. Still, to mitigate the risk of bias, we control for a dummy indicating whether each project was sited in the region of the president's birth, which has been shown to be an important source of variation in the subnational distribution of aid.[80]

A second potential threat is that donors or recipient governments may complete the highest-priority projects first, such that planned projects are of lower priority than completed ones at any given moment in time. We explore this possibility in Section B.4 of the supplementary material, in which we compare the distribution of sectors (agriculture, communications, education, and so on) for all planned and completed projects in our analysis. If donors complete the highest-priority projects first, then intuitively we should expect high-priority sectors (like transportation and communications infrastructure) to constitute a larger proportion of completed projects than planned ones. As we show in Section B.4, this doesn't appear to be the case.

A third potential threat is that citizens may protest projects during the planning or implementation phase, causing either the donor or the recipient government to cancel the projects outright. This is especially relevant for Chinese aid, which has generally been more controversial than aid from other donors. Again, we view this risk as relatively minor, as there seem to be few (if any) examples of Chinese projects being canceled due to protests in Africa during the period under study.[81] But as an imperfect proxy, we use the Armed Conflict Location and Event Dataset (ACLED) to control for the number of protests within thirty kilometers of each respondent before the first Chinese or US project was planned within that same radius.[82] We also include additional controls to mitigate any remaining biases, including gender, age, religion, distance to the capital city, a dummy for cities, and, as discussed above, a dummy indicating whether the respondent lives in the president's home region.

Given the Chinese government's lack of transparency about the aid it provides, an additional challenge is to identify the universe of Chi-

---

[80] Dreher et al. 2016.

[81] There have been cases of cancellation due to protests in Asia—notably in Myanmar, Pakistan, and Nepal. The closest African analogue we're aware of was the cancellation of a 2007 inauguration ceremony at Zambia's Chambishi copper mine (Alden 2007), but that project was completed despite the protests.

[82] Ideally, we could control for protests that occurred after the first project was planned but before it was completed, but this would induce posttreatment bias. Ours is a second-best approach. We interpret past protests as a proxy for potential future protests.

nese-provided services. AidData has advanced our ability to do this, but because it relies on information scraped from secondary sources, the sample of projects in the data set, including their location and timing, is inevitably incomplete. In Section B.3 of the supplementary material, we discuss in detail our approach to dealing with missing data and show that the projects in our sample don't appear to differ systematically from those in the rest of the AidData data set (or AIMS), which suggests that our coding rules aren't likely to induce or exacerbate bias.

Another challenge is the often-blurry line between Chinese aid (Overseas Development Assistance, or ODA) and the less concessionary loans and credits that China offers (Other Official Flows, or OOF). Unfortunately, information on concessionality and intent is not available for many Chinese projects, and these distinctions can only be drawn suggestively.[83] We combine project types in our analysis and interpret our results as the impact of Chinese "aid," broadly defined. But in Section B.10 of the supplementary material, we do distinguish infrastructure from other Chinese projects to assess the possibility that highly visible, large-scale infrastructure projects have de-legitimizing effects that projects in other sectors do not. We find that in either case, the risk of de-legitimizing effects is low.

In addition, it's possible that planned projects themselves affect citizens' attitudes through the anticipation of future services or employment opportunities. This wouldn't bias our results, but it would change their interpretation. If planned projects affect citizens' attitudes, then subtracting the coefficient on planned projects from the coefficient on completed ones would yield the impact of project completion net of both selection and anticipation effects. If planned projects don't affect citizens' attitudes, then subtracting coefficients would yield the effect of project completion net of selection effects alone. Either way the difference-in-differences would return the quantity of interest, which is the net effect of services that are actually provided. Nonetheless, to avoid ambiguities in interpretation, we define planned projects as those for which a formal agreement hasn't yet been reached. These projects are in such a nascent stage of development that anticipation effects are highly unlikely, as we discuss further in Section B.5 of the supplementary material.

---

[83] For example, Dreher et al. 2018, 11, adopt a "second-best" approach by distinguishing "ODA-like" projects from "OOF-like" ones, recognizing that this distinction only partially aligns with more-standard definitions of ODA and OOF.

ADVANTAGES OF OUR APPROACH

None of the components of our research design is flawless, but each helps to compensate for the limitations of the others. Although surveys capture real-world, individual-level exposure to donor-provided services, they also rely on self-reports and are susceptible to selection bias and to nonrandom recall. AidData and AIMS capture real-world exposure to aid without relying on self-reports, but only indirectly, since not all individuals who live near donor-provided services use or even know about those services. AidData and AIMS also suffer from missing data. Survey experiments solve potential selection problems via randomization, but they're more stylized and they still rely on self-reports to measure tax compliance and other outcomes. The tax compliance game solves the selection problem while capturing observed rather than self-reported tax compliance, but it's even more stylized. Our within-country data are rich and highly detailed, but may not generalize to other countries; our cross-country data are coarser, but may have more external validity. Our ability to triangulate across the various components of our research design is an important advantage of our approach. The more consistent our results, the less likely they are to be artifacts of a particular method or sample.

WITHIN-COUNTRY RESULTS

BEHAVIORAL GAME RESULTS

Figure 1 reports average treatment effects for our modified tax compliance game in Liberia.[84] We present our results graphically for ease of interpretation. Our proxy for tax compliance is the ratio of unreported income to total income over the ten rounds of the game; the lower the ratio, the higher the degree of compliance. The income received by participants in each round was randomized, but the total received (1,000 LD) was not, and the distribution of possible incomes was held constant across participants to eliminate differential income effects.[85] To increase precision and to mitigate incidental imbalance on individual covariates, we include controls for gender, age, employment, religion (a dummy for Muslims), and education (dummies for primary, junior high, and high school education).

---

[84] In our PAP, we prespecified five hypotheses about heterogeneous treatment effects as well. These results are irrelevant to our analysis here.

[85] In each round, each participant sampled without replacement from the same distribution of incomes. Over ten rounds, each participant received each of ten possible incomes exactly once—LD 0, 70, 75, 80, 85, 95, 105, 130, 160, or 200.
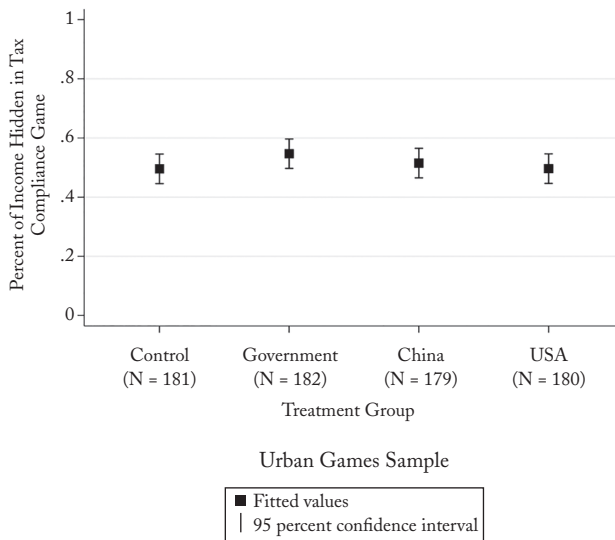
FIGURE 1
AVERAGE TREATMENT EFFECTS FOR TAX COMPLIANCE GAME IN LIBERIA[a]

[a]Average treatment effects on the percent of income hidden across ten rounds in the tax compliance game. We control for age, gender, education, employment, and religion at the individual level.

The squares in Figure 1 denote fitted values from an ordinary least squares regression given by $y_i = \alpha + \sum_{j=1}^{3} \beta_j T_{ji} + \sum_{k=1}^{L} \delta_k X_{ki} + e_i$ where y_i denotes the ratio of unreported to total income, T_{ji} indexes the three treatment groups (with the control as the base category), and X_{ki} indexes our L controls. The lines denote 95 percent confidence intervals. Our results in Figure 1 provide no evidence to suggest that information about aid decreases tax compliance. Although rates of tax evasion are generally high in the game—at or above 50 percent across the board[86]—

[86] Interestingly, observed rates of tax evasion in the game are much higher than self-reported rates in the survey, which, as we show in Section A.2 of the supplementary material, range from 8 percent to 15 percent, depending on the sample. One possible explanation for this disparity is that participants in the game didn't take their decisions seriously, perhaps because the government wasn't physically present to punish them, and so they evaded at a higher rate than they would outside the lab. We view this explanation as unlikely. As we discuss above and in Section A.14 of the supplementary material, when we conducted focus groups after the game, we found that most participants directly connected the experiment to real-world taxation and took their decisions seriously. Moreover, facilitators administered audits during the game itself, so the government's physical absence shouldn't have affected participants' beliefs about the probability of detection and punishment.

Another potential explanation is that survey respondents overreported compliance, perhaps due to social desirability bias. Indeed, avoiding this problem is the reason the tax compliance game was created in the first place. As Alm, Jackson, and McKee 1992, 107, explain, "by its very nature, people have an incentive to hide information on their [tax] evasion behavior, and this concealment makes empirical work quite difficult." Tax evasion is both illegal and socially undesirable, which together may incentivize overreporting of compliance. Although this strikes us as more likely, we can't adjudicate between these two explanations empirically. Our ability to combine survey- and game-based data in a single study is one of the advantages of our research design.

they're substantively similar and statistically indistinguishable across treatment groups, ranging from a low of 50 percent in the control and US treatment groups to a high of 54 percent in the government treatment group. Our confidence intervals are narrow, suggesting that these nulls are not likely to be artifacts of statistical imprecision.

Survey Experiment Results

Figure 2 reports average treatment effects for our urban (panel [a]) and rural (panel [b]) survey experiment. The dependent variable is an additive index of dummies indicating whether respondents agree or strongly agree with each of the four statements about tax morale described above. (Our results are unchanged if we add the Likert scales without dichotomizing.) To increase precision and to mitigate any incidental imbalance, all specifications include individual-level controls for gender, age, employment, religion, education, and household wealth. For the rural sample, we also include district fixed effects and community-level controls for population, average household wealth, average literacy, average employment, and average education, gleaned from the 2008 census. Figure 2 reports fitted values and 95 percent confidence intervals with standard errors clustered by neighborhood (for the urban sample) or community (for the rural sample).

Again, we find little evidence to suggest that information about aid diminishes tax morale. Our point estimates are nearly identical across treatment groups in the rural sample, with almost fully overlapping confidence intervals. These nulls are also consistent across demographic and partisan subgroups, as we show in Section A.9 of the supplementary material. In the urban sample, respondents in the government and US treatment groups are more likely to believe they have an obligation to pay taxes relative to respondents in both the control and China treatment groups. On average, respondents in the government treatment group score a 3.38 out of 4 on our index of tax morale, compared to 2.84 in the control group and 2.79 in the China treatment group—an increase of 19 percent ($p = 0.07$) and 21 percent ($p = 0.05$), respectively. Respondents in the US treatment group score 3.28 on average, which is statistically larger than the average score in the control group (an increase of 16 percent, $p = 0.08$) and in the China treatment group (17 percent, $p = 0.04$).

This partially confirms our hypothesis that the American and Chinese aid regimes have differential effects on the legitimacy of recipient states. But contrary to our expectations, the difference is driven not by the adverse effects of Chinese aid, but rather by the beneficial effects of
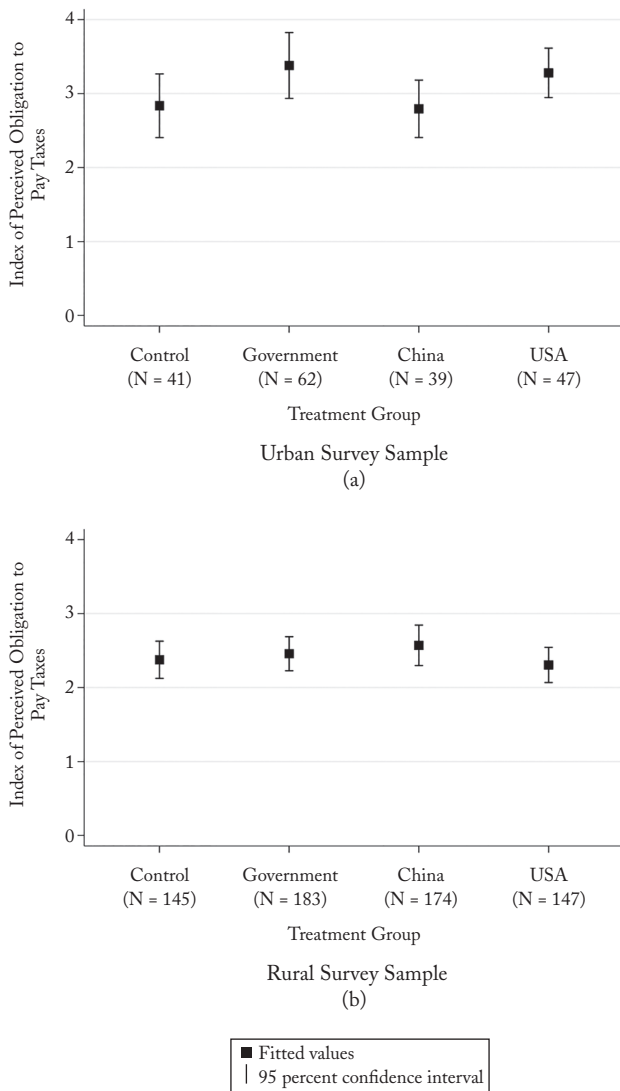
FIGURE 2

AVERAGE TREATMENT EFFECTS FOR SURVEY EXPERIMENT IN LIBERIA[a]

[a]Average treatment effects on the perceived obligation to pay taxes in the survey experiment in urban and rural Liberia. The dependent variable is an index scaled from zero to four. We control for age, gender, education, employment, religion, and wealth at the individual level in both samples. We also control for population, education, literacy, employment, and wealth at the community level in the rural sample, with district fixed effects. Standard errors are clustered by neighborhood (urban sample) or community (rural sample).

US aid. Priming urban respondents to reflect on the role taxation plays in service provision appears to reinforce their belief that tax compliance is an obligation shared by all citizens. But priming them to reflect on US aid has the same effect—which suggests that, if anything, exposure to US aid strengthens rather than weakens tax morale. Average scores between the control and the China treatment groups on our tax morale index are substantively similar and statistically indistinguishable. This suggests that information about China's role in service provision neither increases nor decreases tax morale, relative to no information at all.

SURVEY RESULTS

Table 2 reports correlations between exposure to aid and attitudes toward the Liberian government in our urban and rural survey samples. We distinguish between two types of exposure: consumers (those who know about or have used donor-provided projects) and labor (those who have worked for a donor-funded contractor or know someone who has). These different forms of exposure should affect citizens' attitudes through different mechanisms. For consumers, it is primarily through the quality of donor-provided services. For labor, it is primarily through the wages and hiring and management practices of donor-funded contractors. These correspond to outcome-based and procedural legitimacy, respectively—although of course, consumers may also be sensitive to procedural legitimacy and labor may also be sensitive to outcome-based legitimacy.

To avoid posttreatment bias, we only control for individual-level characteristics that either are fixed over time or are extremely unlikely to be affected by exposure to aid: age, gender, and religion. We rely on the 2008 census as a source of additional community-level controls. At the community level we control for population, wealth, literacy, proportion of residents with no education, and proportion of residents who are unemployed. We report coefficients from OLS regressions with standard errors clustered by neighborhood (for the urban sample) or community (for the rural sample).

Aid may affect tax compliance in part through its effect on citizens' perceptions of government as fair and transparent (procedural legitimacy). To operationalize these perceptions, the dependent variable in column 1 of Table 2 is an additive index of three dummies indicating whether respondents believe the Liberian government (1) treats all Liberians equally, (2) makes decisions in an open and transparent manner, and (3) is free of corruption. As an additional measure of perceived government responsiveness, the dependent variable in column 2 is an in-

TABLE 2

FOREIGN AID AND STATE LEGITIMACY IN LIBERIA USING SURVEY DATA[a]

	Perceptions of Government (Index)	Believes Democracy Is High Quality	Has Ever Refused to Pay Taxes	Believes Government Has the Right to Tax	Believes It's Easy to Avoid Paying Taxes
Urban Sample					
User of Chinese projects	0.02	−0.12	−0.03	0.12	−0.14
	(0.13)	(0.10)	(0.10)	(0.09)	(0.11)
User of US projects	0.09	0.24	−0.10	0.03	−0.05
	(0.17)	(0.06)***	(0.05)*	(0.06)	(0.08)
Worker for Chinese contractors	0.03	0.10	0.01	−0.00	−0.08
	(0.17)	(0.10)	(0.05)	(0.04)	(0.05)
Worker for US contractors	0.16	−0.14	0.13	−0.03	0.10
	(0.16)	(0.10)	(0.06)**	(0.07)	(0.06)
Observations	196	196	186	186	186
Individual-level controls	y	y	y	y	y
Community-level controls	n	n	n	n	n
District fixed effects	n	n	n	n	n
Rural Sample					
User of Chinese projects	0.15	−0.05	−0.01	0.03	0.01
	(0.09)*	(0.04)	(0.03)	(0.04)	(0.04)
User of US projects	0.02	−0.05	0.00	0.12	0.02
	(0.10)	(0.04)	(0.03)	(0.04)***	(0.06)
Worker for Chinese contractors	0.11	−0.03	0.01	−0.01	0.01
	(0.11)	(0.06)	(0.04)	(0.04)	(0.05)
Worker for US contractors	−0.04	0.08	0.03	−0.01	−0.04
	(0.13)	(0.05)	(0.04)	(0.04)	(0.07)
Observations	685	652	652	671	652
Individual-level controls	y	y	y	y	y
Community-level controls	y	y	y	y	y
District fixed effects	y	y	y	y	y

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.01$.

[a] Correlation between exposure to foreign aid and perceptions of government in urban and rural Liberia. The dependent variable in column 1 is an index scaled from zero to three; the dependent variables in columns 2 through 5 are dummies. We control for age, gender, and religion at the individual level in both samples. We also control for population, education, literacy, employment, and wealth at the community level in the rural sample, with district fixed effects. Standard errors, clustered by neighborhood (urban sample) or community (rural sample), are in parentheses.

dicator for whether respondents express satisfaction with the quality of democracy in Liberia. In columns 3, 4, and 5, our dependent variables are indicators for whether respondents have ever evaded taxes (a measure of tax compliance), believe the government has the right to collect taxes (tax morale), and believe taxes are easy to avoid, respectively. By measuring not only whether citizens believe government has a right to collect taxes, but also whether they believe taxes are easy to avoid, we're better able to disentangle whether citizens pay taxes because they believe they have an obligation to do so or because they fear punishment if they don't—or both.[87] The wording of the questions in columns 2 through 5 was taken from the Afrobarometer survey to facilitate comparison with the other countries in our study.

In the rural sample, the correlations between exposure to aid and attitudes toward government are almost uniformly null, regardless of which indicator we use. The only notable exception is that respondents who know about or have used US-provided projects are twelve percentage points more likely to affirm the government's right to tax than respondents who haven't used those projects (rural sample, column 4). This is consistent with the results of our (urban) survey experiment, as described above. Rural respondents who know about or have used Chinese-provided projects also tend to have more favorable perceptions of the Liberian government (rural sample, column 1), although this result is only weakly statistically significant at conventional levels. In Section A.7 of the supplementary material, we replicate our analysis from the rural sample using AidData to operationalize exposure to Chinese aid. Our results are similar.

Likewise, in the urban sample, respondents who know about or who have used a US-provided service are twenty-four percentage points more likely to describe Liberian democracy as high quality (column 2) and ten percentage points less likely to report ever refusing to pay taxes (column 3). This, too, is consistent with our urban survey experiment. Surprisingly, those who have worked for a US-funded contractor or know someone who has are thirteen percentage points more likely to report refusing to pay taxes (column 3) and ten percentage points more likely to believe taxes are easy to avoid (column 5)—although this correlation falls short of statistical significance at conventional levels ($p = 0.12$). But these workers are no more or less likely to endorse the government's right to tax (column 4), which suggests that they may engage

[87] Blair 2018.

in tax evasion not because they believe it's acceptable, but because they believe (probably correctly) that they're unlikely to get caught. Taken together, these results suggest that exposure to aid generally has benign effects on state legitimacy in the eyes of Liberian citizens.

CROSS-COUNTRY RESULTS

Table 3 tests the effects of Chinese- and US-provided projects on state legitimacy across the six African countries for which AidData, AIMS, and Afrobarometer data are all available. The dependent variable in column 1 is an additive index of dummies indicating citizens' trust in six state institutions: police, military, local council, parliament, president, and courts. All other dependent variables are identical to those in Table 2. Two of our taxation measures (columns 3 and 5) are available only for rounds five and six of the Afrobarometer survey. Our remaining dependent variables are available for rounds two through four, as well. All specifications include country and Afrobarometer-round fixed effects, and controls for gender, age, religion, distance to the capital city, number of previous protests within a thirty-kilometer radius, a dummy for cities, and a dummy indicating whether the respondent lives in the president's home region. Standard errors are clustered by community.

We find no evidence to suggest that Chinese aid is eroding state legitimacy among these six countries. If anything, the opposite appears to be true. Completed Chinese projects are positively and statistically significantly correlated with our index of trust in government (column 1), and the difference between the coefficients on completed and planned projects is also positive and statistically significant. Nor do we find evidence to suggest that US aid diminishes state legitimacy. US projects tend to be sited in locations where trust in government is relatively low (column 1) and where perceptions of democracy are relatively negative (column 2). Interestingly, US projects also tend to be sited in locations where belief in the government's right to tax is relatively high (column 4). Completed US projects are positively correlated with belief in the government's right to tax (column 4) and negatively correlated with the perceived ease of tax evasion (column 5), but in both cases the coefficient on completed projects isn't statistically different from the coefficient on planned ones. Nevertheless, US projects do appear to improve citizens' perceptions of democracy after differencing away the selection effect. To the extent that US aid affects state legitimacy in the eyes of African citizens, the effects appear to be beneficial.

TABLE 3

FOREIGN AID AND STATE LEGITIMACY ACROSS SIX AFRICAN COUNTRIES[a]

	Trust in Government (Index)	Believes Democracy Is High Quality	Has Ever Refused to Pay Taxes	Believes Government Has the Right to Tax	Believes It's Easy to Avoid Paying Taxes
Near completed Chinese project	0.23	0.01	−0.01	−0.02	0.03
	(0.06)***	(0.02)	(0.01)	(0.03)	(0.02)
Near planned Chinese project	0.05	0.01	0.04	0.04	−0.001
	(0.06)	(0.02)	(0.04)	(0.01)***	(0.05)
Near completed US project	0.01	0.01	−0.02	0.04	−0.03
	(0.07)	(0.01)	(0.01)	(0.01)***	(0.01)**
Near planned US project	−0.10	−0.05	0.07	0.06	0.25
	(0.06)*	(0.02)***	(0.09)	(0.02)***	(0.19)
Completed vs. planned Chinese project p-value	0.043	0.922	0.337	0.137	0.570
Completed vs. planned US project p-value	0.180	0.003	0.333	0.427	0.147
Observations	21857	23759	11426	24910	9914
Individual-level controls	y	y	y	y	y
Community-level controls	y	y	y	y	y
Country fixed effects	y	y	y	y	y
Afrobarometer round fixed effects	y	y	y	y	y
Cutoff for planned projects	2010	2010	2010	2010	2010
Buffer	30 km	30 km	30 km	30 km	30 km

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.01$.
[a] Correlation between exposure to foreign aid and perceptions of government across six African countries. The dependent variable in column 1 is an index scaled from zero to six; the dependent variables in columns 2 through 5 are dummies. Exposure is operationalized as a dummy for any completed or planned Chinese or US projects within a thirty-kilometer radius. We control for gender, age, religion, distance to the capital city, number of previous protests within a thirty-kilometer radius, a dummy for cities, and a dummy indicating whether the respondent lives in the president's home region. Standard errors, clustered by community, are in parentheses.

DISCUSSION AND CONCLUSION

In this study, we develop and test a theory to explain how different aid regimes might have disparate effects on the perceived legitimacy of recipient states. In our PAP, we hypothesized that exposure to aid would diminish state legitimacy in Africa and that this effect would be especially pronounced for aid delivered by China. Our findings do not conform to these expectations. Why? What explains these null (and in

some cases positive) results? Might they be an artifact of flaws in our research design, such as measurement error or lack of statistical power? Although we can't definitively eliminate these concerns, there are several reasons to believe that they don't explain our findings.

First and most important, our results are generally consistent across multiple approaches to measurement, identification, and estimation. Second, although scholars debate the extent to which tax compliance is a reliable proxy for state legitimacy,[88] it remains one of the most widely used proxies in the literature[89]—and in any event, our results are robust to other potential proxies as well (like trust in government). This suggests that problems with the conceptualization and measurement of our dependent variables are unlikely to explain our results. Third, as we show in Section A.5 of the supplementary material, in Liberia at least, our survey-based and AidData-based proxies for exposure to Chinese aid are highly positively correlated with each another. Both proxies are measured with some error, but the correlation between them suggests that they're capturing exposure to aid in an empirically meaningful way. This suggests that problems with the conceptualization and measurement of our independent variables are unlikely to explain our results either. (Unfortunately, we can't replicate this exercise for US aid because we don't have AIMS data for Liberia.)

Fourth, most of our nulls are precisely estimated, with narrow confidence intervals and point estimates that are close to zero or, in the case of the behavioral game and survey experiment, similar to one another. If statistical power were a problem, we'd expect to see wider confidence intervals around more idiosyncratic point estimates. But we do not. A more formal ex ante power analysis confirms these intuitions. In both our behavioral game and our rural survey experiment, we have 80 percent power to detect an effect equal to roughly 0.29 standard deviations of the control group mean. There are no clear rules of thumb to apply here, but this would generally be considered a small effect. In our urban survey experiment, we have 80 percent power to detect an effect of roughly 0.51 standard deviations. This would be considered a moderate effect. While ex post power calculations should be interpreted with caution, they similarly suggest that statistical power is unlikely to explain our results, as we discuss in Section A.6 of the supplementary material.

Is it possible that the hypotheses we tested were implausible to begin with, or that our research design somehow guaranteed a null result, perhaps due to the bundled nature of our experimental vignettes?

[88] Blair and Winters 2020; Dolan 2020.
[89] Levi, Sacks, and Tyler 2009.

Again, we view these explanations as unlikely. As discussed above, there are many reasons to expect aid to diminish the legitimacy of recipient states and many reasons to believe Chinese aid is especially likely to have pernicious effects. Moreover, while our vignettes emphasized both the positive and negative aspects of Chinese (and US) aid, the fact that we find statistically significant beneficial treatment effects in the urban survey experiment suggests that nulls were not a foregone conclusion.

These bundled treatments also more accurately capture the mixed signals about aid that people actually receive in the real world and were primarily intended as a prime, not as a source of new information. We could have manufactured adverse treatment effects by heavily priming the (purported) problems with the Chinese aid regime, but this would have taught us little about the real world and would have raised serious ethical concerns. That neither our experimental nor our observational measures of exposure to Chinese aid are negatively correlated with perceptions of government lends additional credence to this interpretation of our results.

A potential concern with the design of the tax compliance game is that the treatment was weak—that given the extended length of the vignettes and given the novelty and complexity of the game, participants simply ignored or forgot about the vignettes as they decided how much of their income to report. This is indeed a concern. But the survey experiment is potentially susceptible to the opposite problem—that given the directness and simplicity of the vignettes and given the outcomes were measured immediately after treatment was administered, participants fixated on the vignettes when answering our survey questions. In other words, the tax compliance game may have risked underestimating treatment effects while the survey experiments may have risked overestimating them. The consistency of our conclusions across these two components of our research design suggests that they're unlikely to be artifacts of overly weak or overly strong treatments. This, again, is an advantage of triangulation.

Is it possible that Chinese aid just happens to have had benign or beneficial effects on state legitimacy in Liberia and in the six countries in our cross-country analysis? Might we have found adverse effects in other settings? To explore this possibility, in Section B.9 of the supplementary material we expand our sample to the thirty-eight African countries for which AidData and Afrobarometer data are both available, but AIMS data are not. Our results are substantively similar to those reported here. After we difference away the selection effect, completed Chinese projects appear to improve citizens' perceptions of the quality

of democracy and to have no effect on trust in government, tax compliance, or tax morale. It's possible that US aid might have more pernicious effects in an expanded sample, although we see no a priori reason to expect that.

If selection bias, measurement error, statistical power, and sample selection don't explain our results, then what does? We can't say for certain, and because we expected to find adverse effects—an expectation that we preregistered—any explanation is inevitably post hoc. Nonetheless, there are possibilities that future scholars might consider. First and most obvious, although critics view certain features of the Chinese aid regime as problematic, citizens of recipient countries may not share that view. For example, if citizens reject good governance conditionalities as paternalistic, they may prefer the unconditionality of Chinese aid. If citizens prioritize low cost over high quality, they may prefer the speed of China's approach to implementation. And if citizens value large-scale infrastructure projects over smaller-scale alternatives, then the recipient government's inability to claim credit for these projects may be irrelevant, especially if government officials can at least claim credit for attracting Chinese funds.[90]

Second, aid may have positive effects on some citizens' perceptions and negative effects on the perceptions of others, resulting in a net null. In Section A.9 of the supplementary material, we find only limited evidence of heterogeneous treatment effects among our sample of rural Liberians, but heterogeneity along dimensions we didn't measure is possible. More intriguingly, if citizens value both outcome-based and procedural legitimacy, then aid may simultaneously have positive and negative effects on the perceptions of the same citizens. For example, citizens may perceive donor-provided services as high quality (outcome-based legitimacy), but may also perceive the implementation of those services as unfair and nontransparent (procedural legitimacy). If these positive and negative perceptions counterbalance one another, the result will be, at worst, a net null.

Our survey provides some evidence of these conflicting views at work. Respondents were asked to evaluate Chinese and American donors along seven dimensions, including the quality of the services they provide (outcome-based legitimacy) and the fairness of their hiring and management practices (procedural legitimacy). We use respondents'

---

[90] There is some suggestive evidence to support this explanation. In a recent Afrobarometer report, respondents who view China's influence as positive tend to cite investment in infrastructure as the most important reason for their support (Lekorwe et al. 2016), but we can't infer the effects of Chinese aid from these numbers alone.

answers to these questions to construct additive indices ranging from zero to seven; the higher the score on the index, the more favorable the respondent's perceptions of the corresponding donor. In our rural sample, we find that nearly a quarter of all respondents (23 percent) score 3 or 4 on the index of perceptions of American donors—indicating an almost even mix of positive and negative perceptions—and nearly half (49 percent) score between 2 and 5. The average score is 3.3, which is almost exactly in the center of the distribution of possible scores. Perceptions of China are generally more negative—the average score on the index of perceptions of Chinese donors is 2.1—but more than one quarter of all respondents (28 percent) nonetheless score 3 or 4. Perceptions are similarly mixed in the urban sample. If respondents are ambivalent about aid, then exposure to it may have a net null effect on their perceptions of their own government.

Focus groups conducted after the tax compliance game provide further evidence of these conflicting views. Some participants reported that they "feel nice to see those people [donors] coming in to help us…. So I feel nice to our government for bringing them in" (focus group code 7.2.N.917). Other participants said they "feel happy when the government comes in with those investors" (10.1.U.919), or that they "feel good because it is good for foreigners to come to the country to help development" (2.2.G.911). But some were more circumspect, arguing that Liberians should not "depend on the outside world to keep developing our own country" (1.11.U.921), and that development "is something the government needs to do for our own country people" (1.1.C.921). As one participant pointedly stated, "It's like my government is not able to handle her own country" (1.1.N.904). In some cases, the same participants expressed conflicting opinions in response to different questions, which also suggests ambivalence about the role of aid in Liberia.

Again, we predicted that Chinese aid would have negative effects on state legitimacy, and these explanations for our results are post hoc. But they are amenable to empirical evaluation. One way to test the first explanation might be to elicit citizens' preferences before their exposure to aid, and then to assess whether the effects of aid vary with the intensity of those preferences. The second explanation could be tested by using a conjoint experiment to randomize information about different characteristics of the Chinese aid regime, and then to assess whether particular combinations of characteristics generate greater or lesser support.

Whatever the explanation, our results suggest that the relationship between aid and state legitimacy may need to be rethought. We find

that Chinese aid is not damaging state/society relations, despite the warnings of policymakers in the US and elsewhere. More generally, our findings, combined with those of multiple recent studies across multiple contexts,[91] suggest that aid is not eroding the legitimacy of recipient states in the eyes of citizens, as critics have long feared it might.

SUPPLEMENTARY MATERIAL

Supplementary material for this article can be found at https://doi.org/10.1017 /S004388712000026X.

DATA

Replication files for this article can be found at https://doi.org/10.7910/DVN/5J H0WJ.

DATA AVAILABILITY STATEMENT

Preanalysis Plan was registered at Evidence in Governance and Politics (EGAP) and can be found at https://osf.io/qek4v.

REFERENCES

Achen, Christopher H., and Larry M. Bartels. 2004. "Blind Retrospection: Electoral Responses to Drought, Flu, and Shark Attacks." Working Paper. At https://international.ucla.edu/media/files/PERG.Achen.pdf, accessed December 29, 2020.

Achen, Christopher H., and Larry M. Bartels. 2012. "Blind Retrospection: Why Shark Attacks Are Bad for Democracy." Working Paper 5-2013. Nashville, Tenn.: Vanderbilt University. Center for the Study of Democratic Institutions. At https://my.vanderbilt.edu/larrybartels/files/2011/12/CSDI_WP_05-2013 .pdf, accessed December 29, 2020.

AidData Research and Evaluation Unit. 2017. "Geocoding Methodology, Version 2.0.2." Williamsburg, Va.: William & Mary AidData. At https://www.aiddata .org/publications/geocoding-methodology-version-2-0, accessed December 28, 2020.

Alden, Chris. 2007. *China in Africa.* London, UK: Zed Books.

Alm, James, Betty R. Jackson, and Michael McKee. 1992. "Estimating the Determinants of Taxpayer Compliance with Experimental Data." *National Tax Journal* 45, no. 1: 107–14. At https://www.jstor.org/stable/41788949.

Bababekov, Yanik J., Sahael M. Stapleton, Jessica L. Mueller, Zhi Ven Fong, and David C. Chang. 2018. "A Proposal to Mitigate the Consequences of Type 2 Error in Surgical Science." *Annals of Surgery* 267, no. 4: 621–22. At https://doi .org/10.1097/sla.0000000000002547.

Baldwin, Kate, and Matthew S. Winters. 2018. "Bypass Aid and Perceptions of Local Government Performance and Legitimacy." Working Paper no. 56. Williamsburg, Va.: William & Mary AidData. At https://www.aiddata.org /publications/bypass-aid-and-perceptions-of-local-government-performance -and-legitimacy, accessed January 4, 2021.

[91] Cruz and Schneider 2017; Dietrich and Winters 2015; Dietrich, Mahmud, and Winters 2018.

BenYishay, Ariel, Renee Rotberg, Jessica Wells, Zhonghui Lv, Seth Goodman, Lidia Kovacevic, and Dan Runfola. 2017. "Geocoding Afrobarometer Rounds 1–6: Methodology & Data Quality." Williamsburg, Va.: William & Mary AidData. At https://www.aiddata.org/publications/geocoding-afrobarometer-rounds-1-6-methodology-data-quality, accessed December 28, 2020.

Blair, Robert A. 2018. "Legitimacy after Violence: Evidence from Two Lab-in-the-Field Experiments in Liberia." SSRN Working Paper 2326671. At http://dx.doi.org/10.2139/ssrn.2326671.

Blair, Robert A., Sabrina M. Karim, and Benjamin S. Morse. 2019. "Establishing the Rule of Law in Weak and War-Torn States: Evidence from a Field Experiment with the Liberian National Police." *American Political Science Review* 113, no. 3: 641–57. At https://doi.org/10.1017/S0003055419000121.

Blair, Robert A., Benjamin S. Morse, and Lily L. Tsai. 2017. "Public Health and Public Trust: Survey Evidence from the Ebola Virus Disease Epidemic in Liberia." *Social Science & Medicine* 172, January: 89–97. At https://doi.org/10.1016/j.socscimed.2016.11.016.

Blair, Robert A., and Philip Roessler. 2021a. "Replication files for: Foreign Aid and State Legitimacy: Evidence on Chinese and US Aid to Africa from Surveys, Survey Experiments, and Behavioral Games." Harvard Dataverse, V1. At https://doi.org/10.7910/DVN/5JH0WJ.

Blair, Robert A., and Philip Roessler. 2021b. "Supplementary material for: Foreign Aid and State Legitimacy: Evidence on Chinese and US Aid to Africa from Surveys, Survey Experiments, and Behavioral Games." At https:/doi.org/10.1017/S004388712000026X.

Blair, Robert A., and Matthew S. Winters. 2020. "Foreign Aid and State-Society Relations: Theory, Evidence, and New Directions for Research." *Studies in Comparative International Development* 55, no. 2: 123–42. At https://doi.org/10.1007/s12116-020-09301-w.

Bluhm, Richard, Axel Dreher, Andreas Fuchs, Bradley Parks, Austin Strange, and Michael Tierney. 2018. "Connective Financing: Chinese Infrastructure Projects and the Diffusion of Economic Activity in Developing Countries." Working Paper no. 64. Williamsburg, Va.: William & Mary AidData. At https://www.aiddata.org/publications/connective-finance-chinese-infrastructure-projects, accessed December 29, 2020.

Böhnke, Jan, Jan Koehler, and Christoph Zürcher. 2010. "Assessing the Impact of Development Cooperation in North East Afghanistan 2005–2009: Final Report." Evaluation Reports 049. Bonn, Germany: Bundesministerium für wirtschaftliche Zusammenarbeit und Entwicklung. At https://www.oecd.org/countries/afghanistan/46785983.pdf, accessed December 29, 2020.

Böhnke, Jan, and Christoph Zürcher. 2013. "Aid, Minds and Hearts: The Impact of Aid in Conflict Zones." *Conflict Management and Peace Science* 30, no. 5: 411–32. At https://doi.org/10.1177/0738894213499486.

Bratton, Michael. 1989. "The Politics of Government-NGO Relations in Africa." *World Development* 17, no. 4: 569–87. At https://doi.org/10.1016/0305-750x(89)90263-5.

Bräutigam, Deborah. 2009. *The Dragon's Gift: The Real Story of China in Africa*. New York, N.Y.: Oxford University Press.

Bräutigam, Deborah, Odd-Helge Fjeldstad, and Mick Moore, eds. 2008. *Taxa-*

tion and State-Building in Developing Countries: Capacity and Consent. New York, N.Y.: Cambridge University Press.

Bräutigam, Deborah A., and Stephen Knack. 2004. "Foreign Aid, Institutions, and Governance in Sub-Saharan Africa." *Economic Development and Cultural Change* 52, no. 2: 255–85. At https:/doi.org/10.1086/380592.

Brazys, Samuel, Johan A. Elkink, and Gina Kelly. 2017. "Bad Neighbors? How Co-located Chinese and World Bank Development Projects Impact Local Corruption in Tanzania." *Review of International Organizations* 12, no. 2: 227–53. At https:/doi.org/10.1007/s11558-017-9273-4.

Briggs, Ryan C. 2019. "Receiving Foreign Aid Can Reduce Support for Incumbent Presidents." *Political Research Quarterly* 72, no. 3: 610–22. At https:/doi.org/10.1177/1065912918798489.

Bueno de Mesquita, Bruce, and Alastair Smith. 2009. "A Political Economy of Aid." *International Organization* 63, no. 2: 309–40. At https://doi.org/10.1017/s0020818309090109.

Bullock, John G., Donald P. Green, and Shang E. Ha. 2010. "Yes, But What's the Mechanism? (Don't Expect an Easy Answer)." *Journal of Personality and Social Psychology* 98, no. 4: 550–58. At https://doi.org/10.1037/a0018933.

Cruz, Cesi, and Christina J. Schneider. 2017. "Foreign Aid and Undeserved Credit Claiming." *American Journal of Political Science* 61, no. 2: 396–408. At https://doi.org/10.1111/ajps.12285.

Dietrich, Simone, Minhaj Mahmud, and Matthew S. Winters. 2018. "Foreign Aid, Foreign Policy, and Domestic Government Legitimacy: Experimental Evidence from Bangladesh." *Journal of Politics* 80, no. 1: 133–48. At https://doi.org/10.1086/694235.

Dietrich, Simone, and Matthew S. Winters. 2015. "Foreign Aid and Government Legitimacy." *Journal of Experimental Political Science* 2, no. 2: 164–71. At https://doi.org/10.1017/xps.2014.31.

Dolan, Lindsay R. 2020. "Rethinking Foreign Aid and Legitimacy: Views from Aid Recipients in Kenya." *Studies in Comparative International Development* 55, no. 2: 143–59. At https://doi.org/10.1007/s12116-020-09302-9.

Dreher, Axel, and Andreas Fuchs. 2015. "Rogue Aid? An Empirical Analysis of China's Aid Allocation." *Canadian Journal of Economics/Revue Canadienne d'Économique* 48, no. 3: 988–1023. At https://doi.org/10.1111/caje.12166.

Dreher, Axel, Andreas Fuchs, Roland Hodler, Bradley Parks, Paul Raschky, and Michael J. Tierney. 2016. "Aid on Demand: African Leaders and the Geography of China's Foreign Assistance." Working Paper no. 3, revised. Williamsburg, Va.: William & Mary AidData. At https://doi.org/10.2139/ssrn.2531966.

Dreher, Axel, Andreas Fuchs, Brad Parks, Austin M. Strange, and Michael J. Tierney. 2018. "Apples and Dragon Fruits: The Determinants of Aid and Other Forms of State Financing from China to Africa." *International Studies Quarterly* 62, no. 1: 182–94. At https://doi.org/10.1093/isq/sqx052.

Fain, Haskell. 1972. "The Idea of the State." *Noûs* 6, no. 1: 15–26. At https:/doi.org/10.2307/2214510.

Fowler, Alan. 1991. "The Role of NGOs in Changing State-Society Relations: Perspectives from Eastern and Southern Africa." *Development Policy Review* 9, no. 1: 53–84. At https://doi.org/10.1111/j.1467-7679.1991.tb00175.x.

Gasper, John T., and Andrew Reeves. 2011. "Make It Rain? Retrospection and the Attentive Electorate in the Context of Natural Disasters." *American Jour-*

nal of Political Science 55, no. 2: 340–55. At https://doi.org/10.1111/j.1540 -5907.2010.00503.x.

Gelman, Andrew, and John Carlin. 2014. "Beyond Power Calculations: Assessing Type S (Sign) and Type M (Magnitude) Errors." *Perspectives on Psychological Science* 9, no. 6: 641–51. At https:/doi.org/10.1177/1745691614551642.

Gilley, Bruce. 2009. *The Right to Rule: How States Win and Lose Legitimacy*. New York, N.Y.: Columbia University Press.

Gubser, Peter. 2002. "The Impact of NGOs on State and Non-State Relations in the Middle East." *Middle East Policy* 9, no. 1: 139–48. At https://doi.org/10.1111/14 75-4967.00047.

Guiteras, Raymond P., and Ahmed Mushfiq Mobarak. 2015. "Does Development Aid Undermine Political Accountability? Leader and Constituent Responses to a Large-Scale Intervention." NBER Working Paper no. 21434. Cambridge, Mass.: National Bureau of Economic Research. At https://doi.org/10.3386/w2 1434.

Hanauer, Larry, and Lyle J. Morris. 2014. *Chinese Engagement in Africa: Drivers, Reactions, and Implications for U.S. Policy*. Santa Monica, Calif.: RAND Corporation.

Harding, Robin. 2010. "Urban-Rural Differences in Support for Incumbents across Africa." Working Paper no. 120. Afrobarometer. At https://afrobarom eter.org/publications/wp120-urban-rural-differences-support-incumbents -across-africa, accessed December 29, 2020.

Healy, Andrew, and Neil Malhotra. 2010. "Random Events, Economic Losses, and Retrospective Voting: Implications for Democratic Competence." *Quarterly Journal of Political Science* 5, no. 2: 193–208. At https://doi.org/10.1561/100 .00009057.

Healy, Andrew J., Neil Malhotra, and Cecilia Hyunjung Mo. 2010. "Irrelevant Events Affect Voters' Evaluation of Government Performance." *Proceedings of the National Academy of Sciences* 107, no. 29: 12804–09. At https://doi.org/10.10 73/pnas.1007420107.

Isaksson, Ann-Sofie, and Andreas Kotsadam. 2018a. "Chinese Aid and Local Corruption." *Journal of Public Economics* 159, March: 146–59. At https://doi.org/10 .1016/j.jpubeco.2018.01.002.

Isaksson, Ann-Sofie, and Andreas Kotsadam. 2018b. "Racing to the Bottom? Chinese Development Projects and Trade Union Involvement in Africa." *World Development* 106, June: 284–98. At https://doi.org/10.1016/j.worlddev.2018 .02.003.

Kelly, Gina, Samuel Brazys, and Johan A. Elkink. 2016. "The Dragon's Curse? China, the World Bank, and Perceptions of Corruption in Tanzania." Working Paper no. 26. Williamsburg, Va.: William & Mary AidData. At https://www .aiddata.org/publications/the-dragons-curse-china-the-world-bank-and-pe rceptions-of-corruption-in-tanzania, accessed January 4, 2021.

Knack, Stephen. 2004. "Does Foreign Aid Promote Democracy?" *International Studies Quarterly* 48, no. 1: 251–66. At https://doi.org/10.1111/j.0020-8833 .2004.00299.x.

Koter, Dominika. 2013. "Urban and Rural Voting Patterns in Senegal: The Spatial Aspects of Incumbency, c. 1978–2012." *Journal of Modern African Studies* 51, no. 4: 653–79. At https://doi.org/10.1017/S0022278X13000621.

Lekorwe, Mogopodi, Anyway Chingwete, Mina Okuru, and Romaric Samson.

2016. "China's Growing Presence in Africa Wins Largely Positive Popular Reviews." *Afrobarometer Dispatch* no. 122. At https://www.afrobarometer.org/publications/ad122-chinas-growing-presence-africa-wins-largely-positive-popular-reviews, accessed January 4, 2021.

Levi, Margaret. 2006. "Why We Need a New Theory of Government." *Perspectives on Politics* 4, no. 1: 5–19. At https://doi.org/10.1017/s1537592706060038.

Levi, Margaret, Audrey Sacks, and Tom Tyler. 2009. "Conceptualizing Legitimacy, Measuring Legitimating Beliefs." *American Behavioral Scientist* 53, no. 3: 354–75. At https://doi.org/10.1177/0002764209338797.

Luttmer, Erzo F. P., and Monica Singhal. 2014. "Tax Morale." *Journal of Economic Perspectives* 28 no. 4: 149–68. At https://doi.org/10.1257/jep.28.4.149.

Mcloughlin, Claire. 2015. "When Does Service Delivery Improve the Legitimacy of a Fragile or Conflict-Affected State?" *Governance* 28, no. 3: 341–56. At https://doi.org/10.1111/gove.12091.

Mcloughlin, Claire. 2018. "When the Virtuous Circle Unravels: Unfair Service Provision and State De-Legitimation in Divided Societies." *Journal of Intervention and Statebuilding* 12, no. 4: 527–44. At https://doi.org/10.1080/17502977.2018.1482126.

Moyo, Dambisa. 2010. *Dead Aid: Why Aid Makes Things Worse and How There Is Another Way for Africa*. London, UK: Penguin.

Mukpo, Ashoka. 2014. *Community Relations in the China Union Concession*. Briefing no. 1. Monrovia, Liberia: Sustainable Development Institute. At https://sdiliberia.org/sites/default/files/publications/Community%20Relations%20in%20China%20Union%27s%20Concession.pdf, accessed December 29, 2020.

Naím, Moisés. 2007. "Missing Links: Rogue Aid." *Foreign Policy* 159: 96. At https://www.jstor.org/stable/25462153.

OECD. 2015. *States of Fragility 2015: Meeting Post–2015 Ambitions*. Paris, France: Organization for Economic Cooperation and Development. At https://doi.org/10.1787/9789264227699-en.

Sacks, Audrey. 2012. "Can Donors and Non-State Actors Undermine Citizens' Legitimating Beliefs?" Working Paper no. 140. Afrobarometer. At https://afrobarometer.org/publications/can-donors-and-non-state-actors-undermine-citizens%E2%80%99-legitimating-beliefs, accessed December 29, 2020.

Schmelzle, Cord, and Eric Stollenwerk. 2018. "Virtuous or Vicious Circle? Governance Effectiveness and Legitimacy in Areas of Limited Statehood." *Journal of Intervention and Statebuilding* 12, no. 4: 449–67. At https://doi.org/10.1080/17502977.2018.1531649.

Svensson, Jakob. 2000. "Foreign Aid and Rent-Seeking." *Journal of International Economics* 51, no. 2: 437–61. At https://doi.org/10.1016/s0022-1996(99)00014-8.

Timmons, Jeffrey F. 2005. "The Fiscal Contract: States, Taxes, and Public Services." *World Politics* 57, no. 4 (July): 530–67. At https://doi.org/10.1353/wp.2006.0015.

Tyler, Tom R. 2006. "Psychological Perspectives on Legitimacy and Legitimation." *Annual Review of Psychology* 57: 375–400. At https://doi.org/10.1146/annurev.psych.57.102904.190038.

Vinck, Patrick, Phuong Pham, and Tino Kreutzer. 2011. *Talking Peace: A Popula-*

tion-Based Survey on Attitudes about Security, Dispute Resolution, and Post-Conflict Reconstruction in Liberia. Berkeley, Calif.: University of California, Berkeley, Human Rights Center. At https://escholarship.org/uc/item/59p3f3mn, accessed December 29, 2020.

Wade, Abdoulaye. 2008. "Time for the West to Practise What It Preaches." *Financial Times.* January 23. At https://www.ft.com/content/5d347f88-c897-11dc -94a6-0000779fd2ac, accessed January 6, 2021.

Authors

Robert A. Blair is an assistant professor of political science and international and public affairs at Brown University. His research focuses on international intervention and the consolidation of state authority in conflict and postconflict settings. Blair is the author of *Peacekeeping, Policing, and the Rule of Law after Civil War* (2020). He can be reached at robert_blair@brown.edu.

Philip Roessler is an associate professor of government at The College of William & Mary. His research focuses on conflict, political economy of development, and African politics. He is the author of *Ethnic Politics and State Power in Africa: The Logic of the Coup-Civil War Trap* (2016) and coauthor, with Harry Verhoeven, of *Why Comrades Go to War: Liberation Politics and the Outbreak of Africa's Deadliest Conflict* (2016). He can be reached at proessler@wm.edu.

Acknowledgments

Funding

Key Words

Africa, behavioral games, fiscal contract, foreign aid, legitimacy, Liberia, state/society relations, survey experiments, taxation, tax compliance, tax morale

Review Article

THE STATUS OF STATUS IN WORLD POLITICS

By PAUL K. MACDONALD and JOSEPH M. PARENT

ABSTRACT

What is status? How does it work? What effects does it tend to have? A new wave of scholarship on status in international relations has converged on a central definition of status, several causal pathways, and the claim that the pursuit of status tends to produce conflict. The authors take stock of the status literature and argue that this convergence is not only a sign of progress, but also an obstacle to it. They find that the consensus definition conceals critical contradictions between standing and membership, that its causal pathways are promising but often in tension with each other, and that the literature may be overlooking the ways in which status can help states avoid conflict and promote cooperation under certain conditions.

Deborah Welch Larson and Alexei Shevchenko. 2019. *Quest for Status: Chinese and Russian Foreign Policy.* New Haven, Conn.: Yale University Press, 352 pp.

Michelle Murray. 2019. *The Struggle for Recognition in International Relations: Status, Revisionism, and Rising Powers.* New York, N.Y.: Oxford University Press, 280 pp.

Jonathan Renshon. 2017. *Fighting for Status: Hierarchy and Conflict in World Politics.* Princeton, N.J.: Princeton University Press, 328 pp.

Steven Ward. 2017. *Status and the Challenge of Rising Powers.* New York, N.Y.: Cambridge University Press, 282 pp.

I N a 2018 address to the United Nations, American President Donald J. Trump asserted, "My administration has accomplished more than almost any administration in the history of our country." Humiliatingly, the audience laughed at him. "The United States is stronger, safer, and a richer country than it was when I assumed office," he continued undaunted, "we are standing up for America and for the American people."[1] It is tempting to dismiss these boasts as those of an unusual leader, but while Trump may be unusual,[2] he is not alone. From Brazil and Hungary to the Philippines, leaders the world over are bluntly asserting their international status and bristling at encroachments upon it. In a 2014 address justifying the annexation of Crimea, Russian Presi-

[1] Trump 2018.
[2] Drezner 2020.

dent Vladimir Putin accused the West of "constantly trying to sweep us into a corner because we have an independent position, because we maintain it, and because we call things like they are and do not engage in hypocrisy. But there is a limit to everything."[3] In a 2019 speech, Chinese President Xi Jinping boasted, "The Chinese nation has realized a tremendous transformation: it has stood up, grown rich and is becoming stronger; it has come to embrace the brilliant prospects of national renewal. This phenomenal transformation brings infinite pride to every son and daughter of the Chinese nation."[4]

Why do leaders invoke national status in their public statements? What consequences does this have on world politics? In recent years, a growing literature has developed to answer these questions. Status scholarship, William Wohlforth argues, "has become mainstream. It has gone global."[5] By our count, there have been at least eighteen scholarly monographs in this past decade alone that focus on status, prestige, recognition, and related topics.[6] Scholars have pointed to status as the primary cause of arms races, territorial expansion, and diplomatic crises, as well as of the outbreak and intensity of interstate wars.[7] On this view, Trump, Putin, and Xi are not outliers. Their obsession with their countries' standing reflects impulses that have driven states' foreign policies throughout history.[8]

The four books under consideration here represent some of the best recent attempts to place status at the center of the study of world politics. Their arrival could not be more timely. As China's relative power increases, there are worries that this could shake up the membership of the great powers, exacerbate concerns over standing, and become a source of conflict.[9] So, too, the ascent of populist demagogues has generated anxiety about how questions of status could inflame domestic divisions and fuel international rivalries.[10] If world politics is entering a period in which boasts and brags are supplanting discretion and di-

---

[3] Putin 2014.

[4] Xi 2019.

[5] Wohlforth 2019.

[6] Lebow 2010; Volgy et al. 2011; Miller 2014; Coggins 2014; Larson, Paul, and Wohlforth 2014; de Carvalho and Neumann 2014; Cooley and Snyder 2015; Paul 2016; Pouliot 2016; Lamont et al. 2016; Renshon 2017; Ward 2017b; Gilady 2018; Pu 2019; Murray 2019; Larson and Shevchenko 2019; Charoenvattananukul 2020; Barnhart 2020.

[7] Volgy and Mayhall 1995, 67; Lebow 2010, 15; Wolf 2011, 105; Wohlforth 2014, 139; Sambanis, Skaperdas, and Wohlforth 2015, 280; Barnhart 2017; Ward 2017b, 38; Ward 2017a, 822; Renshon 2017, 154–57; Hall 2017, 12–13; Greve and Levy 2018, 175–76; Murray 2019, 5.

[8] Wohlforth 1993, 28; Deng 2008, 5; Fordham 2011, 593; Onea 2014, 127; Barnhart 2016, 386; Larson and Shevchenko 2019, 174–75.

[9] Onea 2014; Wolf 2014.

[10] Destradi and Plagemann 2019.

plomacy, then instability will likely increase. It is essential that scholars help decipher these disquieting trends.

The aim of this review is to assess the progress of the status literature. We ask three questions: What is status? How does it work? And what effects does it tend to have? On the one hand, the literature has impressively convergent answers to these questions. There is near unanimity that status consists of collective beliefs about a state's standing and membership, based on valued attributes, and is recognized by voluntary deference. It also agrees that status hierarchies are common in world politics, that states crave high perches within these hierarchies, and that a combination of psychological and domestic political factors push states to engage in status-seeking behavior. It concurs that although status does not predestine violence, the pursuit of status does tend to destabilize interstate relations. Therefore, states should recognize each other's status claims and find ways to accommodate them.

On the other hand, this convergence is not only a sign of progress, but also an obstacle to it. By agreeing on what status is, why states want it, and how they tend to compete for it, the status literature has made great strides in addressing fundamental dynamics in international politics. But this apparent consensus conceals crosscutting logics that ought to be openly juxtaposed. For instance, by defining status as both standing and membership, the literature has unreconciled contradictions in its core concept, and makes it tougher to measure persuasively. By developing psychological and domestic pathways, the literature has fleshed out the causal processes through which status can influence foreign policy without fully recognizing tensions within and between these pathways. By focusing on status as a destabilizer, the literature misses the ways that status can help states avoid conflict and promote cooperation under certain conditions. For continued progress, we should acknowledge these tensions and refine our theories and evidence to help adjudicate them.

A single review article cannot cover the entirety of any body of literature, so in what follows, we focus on the most popular topic: studies connecting status to foreign policy and interstate conflict, which is the focus of all four books under review. This means that we must exclude studies that look at how status affects internal conflict, global governance, and other issues, although some of our points may also apply to them.[11] We begin in Section I with a brief overview of the four books. In Section II, we consider some of the challenges scholars have encoun-

---

[11] See, for example, Petersen 2002, 2; Johnston 2008, 76–84.

tered when seeking to define and measure status. Section III explores the varied and contradictory ways scholars have thought about how status influences foreign policy. Section IV examines the nexus between status and competition, and sketches a wider range of possible ways status may shape the prospects for peace. In Section V, we conclude with suggestions about how the status literature can improve its theoretical foundations and expand the scope of its empirical applications.

I. Summary of the Books Reviewed

The four books under review have the same essential purpose: to highlight how status matters in world politics. In *Status and the Challenge of Rising Powers,* Steven Ward argues that status provides the most compelling explanation for why rising powers pursue revisionist foreign policies designed to overturn existing international orders (pp. 3–4). In *The Struggle for Recognition in International Relations: Status, Revisionism, and Rising Powers,* Michelle Murray contends that the failure to recognize rising powers' status claims is the primary cause of spirals of competition and conflict during power transitions (pp. 14–17). In *Fighting for Status: Hierarchy and Conflict in World Politics,* Jonathan Renshon demonstrates that states frequently fight with one another to improve their standing within particular status communities (pp. 21–25). In *Quest for Status: Chinese and Russian Foreign Policy,* Deborah Welch Larson and Alexei Shevchenko maintain that it is the pursuit of status, more than wealth or power, which drives the foreign policy choices of great powers, such as Russia and China (pp. 14–16).

Previous reviews of the status literature focus on the challenge of defining status and distinguishing it from related concepts like "honor," "prestige," or "reputation."[12] These reviews also push the status literature to more clearly delineate the effects of status from those driven by power or interests.[13] These four books make considerable progress in addressing these concerns. First, all rely upon the same definition of status.[14] In its most general form, status refers to "an actor's position within a social hierarchy. It may mean either membership in a highly regarded group . . . or rank within a group" (Ward, p. 35). Collective assessments of status depend on "others' perceptions of a state's ranking on a set of valued characteristics" (Murray, p. 45). In international

---

[12] Dafoe, Renshon, and Huth 2014.
[13] Thompson 2014; Lake 2014. Mercer 2017.
[14] For similar definitions, see Dafoe, Renshon, and Huth 2014, 374–75; Larson, Paul, and Wohlforth 2014, 7–8.

politics, these "prized attributes" can include "military power, economic development, cultural achievements, diplomatic skill, and technological innovation" (Larson and Shevchenko, p. 3). But status is more than the mere possession of valued attributes. Status is also social. It "clarifies what rights, obligations, and patterns of *deference* from others the actor should expect as well as how the actor is expected to behave with respect to others in dominant and subordinate positions" (Renshon, p. 33). High-status states enjoy certain rights and responsibilities, which low-status states accept.

Second, all four provide detailed accounts of why status is not a marginal concern in international politics, but a core and continuous one. Renshon observes that leaders are "plainly obsessed with investing in, seizing, and defending" status, making it "one of the most sought-after qualities in world politics" (pp. 1, 3). Similarly, Ward argues that status is "not the *only* resource that motivates states, but it is a prominent and underappreciated one" (p. 38). Leaders may covet status for a variety of reasons. Some desire status because it is a "valuable resource" that "confers benefits on its holders" in interactions with rival states (Renshon, pp. 52–53). Others seek status because "having higher status increases collective self-esteem and pride" (Larson and Shevchenko, p. 3). Still others value status because of "its significance for domestic political legitimacy" (Ward, p. 37). It is precisely because status is so valuable, whether for instrumental or intrinsic reasons, that the decision to deny another state's status claims can be so consequential (Murray, pp. 12–13). Because status can appear zero-sum and particular status positions, such as that of a great power, are relatively scarce, this raises the stakes for whether a state is included or excluded.

Third, all four detail the ways that the pursuit of status can have destabilizing consequences for world politics, which neither a focus on power nor interests would predict. Ward argues that "anxiety about status" can advantage "hardliners over moderates in domestic contests over the direction of foreign policy" (p. 204). He attributes the poisonous revisionism of Wilhelmine Germany, Imperial Japan, and Interwar Germany to these "obstructed or thwarted status ambitions" (p. 208). Murray contends that rising powers that are denied "world power status" lash out by aggressively pursuing high-status markers, such as large navies (p. 80). She contrasts Britain's successful recognition of American aspirations with its rejection of German ambitions, which resulted in an arms race, then war (p. 191). Broadening the analysis beyond rising powers, Renshon finds states that are dissatisfied with their status in general are more likely to go to war, favoring weak, high-status tar-

gets they can embarrass and defeat (p. 258). Using a combination of experimental, statistical, and historical evidence, he demonstrates that "once triggered, heightened status concerns set in motion a set of consequences at the individual and state level" that more often than not end in conflict (p. 256). Larson and Shevchenko acknowledge that states may "pursue varying strategies for attaining status, depending on the openness of elite clubs and the similarity of their values with those of the established powers" (p. 6). Yet after surveying more than five hundred years of diplomatic history, they conclude that "both China and Russia are hypersensitive to perceived slights and have used military power to assert superiority" (p. 244). Even attempts to bolster one's status through peaceful means can "shade into social competition, insofar as a state is stressing new criteria for status" (p. 245). Conflicts over status are not inevitable, but other things equal, status makes cooperation harder.

II. What Is Status? The Concept and Its Complications

Unlike many social science concepts, status is not essentially contested. The building blocks of status are well understood: there exists a set of collectively valued attributes in world politics, states occupy different positions on these valued attributes, and high-status states have different rights and responsibilities than low-status ones. Despite surface agreement, however, the works considered here reveal significant differences. This is most striking in two related areas: whether status should refer primarily to standing or membership and whether status is best captured using quantitative or qualitative methods. We consider each in turn.

Defining Status: Standing versus Membership

A central ambiguity in consensus definitions of status is whether it is best thought of as a continuous measure of relative standing, a dichotomous attribute of group membership, or some combination of both. Most definitions of status are agnostic on this question. They acknowledge that status can refer to either "membership in a defined club" or "relative standing within such a club."[15] If states with high-status attributes tend to enter high-status clubs, this ambiguity would be unproblematic. Yet theoretically, standing and membership are different concepts, and empirically, they are often mismatched. Postwar Japan ranked high

---

[15] Larson, Paul, and Wohlforth 2014, 7.

on measures of economic influence, but "[lacked] the institutional privileges accorded 'legitimate great powers'" in the UN system.[16] By contrast, "material factors would surely have predicted France's relegation to the international system's periphery" after the Second World War, yet it was nevertheless rewarded with a permanent seat on the UN Security Council.[17]

In practice, different authors end up emphasizing different aspects of status. Some focus on standing, which equates with esteem, and see status as a kind of metric states can use to establish baselines, to draw comparisons, and to assess worth. Renshon comes closest to this position. He notes that "status as rank is not about 'having' versus 'not having'; it concerns how *much* we have relative to others" (p. 35). The benefit of standing is that it allows one to make fine-grained assessments of where states are positioned. Austria-Hungary and Germany were both members of the great power club prior to the First World War, but the former was falling from the ranks while the latter was ascending them, a difference that mattered in their foreign policy orientations.[18] The drawback is that standing is underspecified. Do valued attributes refer primarily to the impressive means states possess or to the virtuous ends they pursue? Who decides which attributes are prized and how? Should we treat standing as a universal metric or disaggregate it into one's standing in a particular issue area, institutional context, or geographic region?[19] Consider Brazil and India, two states with high status aspirations. Brasília's claim has rested on Brazil's economic strength and the establishment of a kind of "consensual hegemony" over states in South America, while New Delhi has focused on projecting India's military strength and cultivating bilateral ties with the hegemon, the United States.[20] One, the other, both, or neither could be considered high status, depending on what criteria and whose judgments we rely on.

Other authors choose to focus on membership, which treats status as closely related to recognition. The attributes that states possess are less important than whether high-status states choose to grant others membership in high-status clubs. Murray is a leading proponent of this approach. She emphasizes that status "refers to a *recognized identity*, not the acknowledgement or acceptance of a state's characteristics or capabilities" (p. 46). A state cannot "simply assert its social status . . . only when recognized does it assume the authority it needs to secure the

[16] Suzuki 2008, 52.
[17] Heimann 2015, 186.
[18] Volgy et al. 2014, 58.
[19] Thompson 2014, 211.
[20] Carranza 2014.

identity it seeks" (p. 46). The advantage of using membership is that which states are members of elite clubs and which are not is often uncontroversial. We can simply assume that states that sit on the UN Security Council, the IMF executive board, or the WTO core negotiation group have higher status than those that are excluded. Yet membership exhibits many of the same problems as standing. Why are some states admitted over others? Why are some clubs more prestigious than others, and how do we know? For example, outsiders attribute outsized influence to the World Economic Forum held annually in Davos, Switzerland, while insiders dismiss it as a "big cocktail party."[21] These issues are compounded in the case of certain clubs, such as the "regional powers," for which membership is more implicit, informal, and sometimes contested.[22] The reasons for, and consequences of, club membership may also have little to do with status. States grant or withhold recognition to one another for a variety of reasons, including to exercise strategic leverage, to avoid third-party punishments, and to achieve geopolitical goals.[23] States join clubs for a variety of reasons, too, only some of which may be tied to status aspirations. And if membership tends to be influenced by nonstatus considerations, then membership loses much of its meaning.

Many authors adopt a flexible approach, lumping standing and membership together and shifting between them. An example of this is Ward's discussion of Wilhelmine Germany. There are passages in which Ward suggests that Kaiser Wilhelm II and his advisors simply desired more respect among the great powers (pp. 73, 78). Yet there are others in which he argues that what they really wanted was recognition as a "world power" (p. 81), which included a right to "naval equality" and the "splitting [of] the entente" (pp. 83, 86). The difference matters a great deal. If we focus on standing, Germany was a rising great power, an economic and military model for much of Europe, whose revisionism appeared premature and self-defeating. Yet if we emphasize membership, Germany rated below Britain, and its revisionism seemed to be the only way it could force its way into the world-power club. Germany embraced revisionism not because it cared about status in general, but because it became obsessed with one kind of status and not the other.

In short, what is missing is a clear set of theoretical expectations about why states may care more about membership than rank or vice versa, and how one goes about aggregating valued attributes in each

[21] Graz 2003, 321.
[22] Nolte 2010.
[23] Coggins 2014.

case. Hybrid solutions sidestep these essential questions and allow scholars to conceptualize status however is convenient for their claims. Hybrid solutions also downplay essential differences between these two kinds of status. For example, when states focus on standing, their assessments tend to be zero-sum, yet when they emphasize membership, they need not be.[24] Just as light may be a particle and a wave, status may be standing and membership, but one should not equate the two, and which kind of status states are preoccupied with matters a great deal.

Measuring Status: Quantitative versus Qualitative Approaches

Assuming a clear conception of status, there remains the matter of measuring it. Some scholars use quantitative data to generate cross-national measures of rank. Renshon represents the cutting-edge of this method. He uses diplomatic exchange data to identify states that reside in central positions within diplomatic networks (pp. 124–32), and to identify the boundaries of distinct status communities (pp. 140–48). An alternative approach uses qualitative methods to identify status motives in the language of policymakers. Ward typifies this approach. "When actors speak in terms of the rights that the state is owed on the basis of its position," he observes, "they are articulating a claim to status" (p. 63). When leaders invoke their national honor or bristle at perceived humiliations, this is taken as evidence that they are driven by status, especially when uttered in private (Murray, p. 85).

The choice of method is tied in part to how one conceives of status. One of the main advantages of quantitative measures is that they allow scholars to capture where states rank in the international hierarchy. This can help to illuminate the extent to which status departs from material capabilities (see Renshon, pp. 135–40) and to provide insights as to why some states "over-" or "under-perform" on status given their material endowments.[25] An obvious challenge of quantitative approaches, of course, is deciding how best to operationalize rank. Some scholars, including Renshon, use measures of diplomatic exchange, based on the assumption that establishing an embassy is an act of social recognition; others look at membership in international organizations, based on the idea that belonging to multiple clubs conveys social prominence.[26] Some rank states based on their aggregate attributes, such as the number of diplomats they host; others, including Renshon, rank states based

---

[24] Lake 2014, 268.
[25] Duque 2018; Røren and Beaumont 2019.
[26] Compare Renshon 2017, 120–23, and Hafner-Burton and Montgomery 2006, 11–12.

on their centrality within broader networks of diplomatic exchange.[27] Some rank states relative to the entire international community; others, including Renshon, rank states within their home regions or among relevant "status communities."[28] Each of these empirical choices is defensible in the abstract, and Renshon does an admirable job validating his particular method, which uses Google's PageRank algorithm with other more-direct measures of status, such as official state visits (pp.132–35). But the more baroque the technique for generating status rankings, the less plausible it is that politicians will practice similar methods of accounting when making their own status assessments.

More important, none of these quantitative measures capture voluntary deference, and as a result, they are at best rough proxies for status. Centrality within diplomatic networks, for example, is partly a product of money: wealthy states can afford to send more diplomats abroad, while their large economies entice diplomats in return.[29] Diplomatic recognition can also be an outgrowth of coercion. The fact that North Korea and Taiwan "receive fewer embassies than their capabilities would warrant"[30] could be evidence of their pariah status or a reflection of the arm-twisting that the United States and China apply to others. And even if diplomatic actions are independent of bribery or threats, they could reflect shared interests or ideological affinities and be a product of choice, not deference. Unless wealth, power, and interests can be disentangled from diplomatic recognition, it is unclear that status is driving recognition.

Qualitative measures approach status differently. Instead of trying to measure status *ex ante*, they look for evidence of status *ex post*, in the statements and actions of policymakers. The advantage of this approach is that it has the potential to capture the subjective and perceptual aspects of status: if leaders perceive that their states occupy particular status positions, and if they describe their actions as efforts to increase their status, then this supports status accounts. The primary problem is that leaders rarely talk about status as scholars define it. Instead, they use words or phrases that seem to evoke status: *national honor*, *national dignity*, *national greatness*, and so on. Yet these phrases may have little or nothing to do with status defined as high rank or membership in elite clubs. After all, there is honor among thieves, too. Leaders may also favor this language for psychological reasons, as a way to convey emo-

[27] Compare Volgy and Mayhall 1995 and Renshon 2017, 124–29.
[28] Compare Cline et al. 2011 and Renshon 2017, 40–48.
[29] Mercer 2017, 138; Ward 2020, 161–62.
[30] Duque 2018, 583.

tions, such as pride, that may be unrelated to where their states stand in social hierarchies.[31] Alternatively, leaders may use emotional language strategically to signal their interests to foreign audiences or to mobilize domestic support.[32] In these cases, it is the rhetorical necessities rather than the status hierarchies that are doing the causal work.

A related problem is the existence of mixed motives and the challenge of disentangling security or economic interests from status concerns. Consider Germany's naval buildup prior to the First World War, a case mentioned by all four authors (Larson and Shevchenko, pp. 8–9; Murray, pp. 94–95; Renshon, pp. 187–88; Ward, pp. 79–80). While status may have been a motive for some German policymakers, others saw a large navy as a military instrument to deter British intervention in a continental war, a diplomatic tool to buttress claims to territorial concessions in China, and a domestic political maneuver to discredit Social Democrats.[33] The fact that naval officers, in particular Admiral Alfred von Tirpitz, were the strongest proponents of becoming a world power suggests that organizational culture and parochial interests can drive status language. Even if we could be clear about what kind of rhetoric reflects status concerns, it can be difficult to determine whose words matter most.

One final issue with qualitative assessments is their tendency to focus almost exclusively on examples of the pursuit of status, not deference to it. Larson and Shevchenko, Murray, and Ward provide ample evidence that states like Germany, Japan, Russia, and China desired respect, but little evidence that status translates into deference. Of course, states might still seek esteem by asserting high status regardless of whether others offer deference. If that's the case, then the pursuit of status would verge on solipsism. If status hierarchies are genuinely social and shape world politics in more systematic ways, we need to see evidence that other states at least recognize and likely defer to those at the top of the totem pole. The problem here is that states align their policies with the preferences of the powerful for a variety of reasons that have nothing to do with status. In his study of the 1899–1902 South African War, Jonathan Mercer finds that Britain's adversaries did not acknowledge its status and refused to defer, while its allies deferred to it out of "a desire for security" rather than "mutual admiration."[34] The weak submit to the strong all the time in world politics, but deference alone is not proof that status hierarchies exist, or that status motives are driving behavior.

[31] Mercer 2017, 139–40.
[32] Götz 2019.
[33] Rüger 2011, 602–05.
[34] Mercer 2017, 157–59.

In sum, both quantitative and qualitative approaches have their advantages. The former allows us to construct useful cross-national metrics of relative rank. The latter provides compelling evidence that status is driving the assessments and actions of policymakers. Yet neither approach fully captures the primary collective manifestation of status as voluntary deference. And if other states do not defer to high status, then the pursuit of status itself may be a chimera.

III. How Status Works: Conditions and Mechanisms

Let us assume that we can settle on a clear conception of status and measure it in practice. How does status work? Status scholars agree that international politics is full of status hierarchies and that states are driven to improve their positions within them. Yet they differ on what conditions activate status concerns and, once activated, by what mechanisms status concerns influence foreign policy. We delve into conditions and mechanisms, respectively.

Conditions: Prompted by Position or a Response to Events?

While scholars maintain that states desire esteem, they acknowledge that status-driven behavior can vary in frequency and intensity. A number of scholars argue that there are certain positions that compel states to be concerned with status. Larson and Shevchenko place particular emphasis on powerful states and their desire to attain great power status. "Great power status carries with it the expectation that the state will be consulted on important issues," they observe, and as a result, "governments have spent enormous sums on efforts to achieve or maintain great power standing, at the expense of their state's power and wealth" (pp. 233–34). Ward also argues that powerful states have a particular interest in status, especially if their material capabilities are rising. "Increasing wealth and military power make a rising state more like established high-status powers," he observes, "which prompts people who identify with the riser to expect—and demand—convergence in terms of standing, influence, and rights" (p. 39). Renshon provides perhaps the broadest positional argument. He argues that states will be most interested in status when there is a "divergence between the status accorded an actor and what they believe themselves to deserve" (p. 53). When there is an inconsistency between a state's objective capabilities and its subjective rank within a particular community, "status dissatisfaction" can prompt greater interest in achieving elevated status (pp. 63–64).

It is worth noting that this list of positions is not exhaustive. While Larson and Shevchenko emphasize powerful states' interests in status, others contend that middle powers should be particularly status-obsessed because they have the most to gain if they are admitted to exclusive clubs.[35] Still others observe that small states should be the most invested in status because their material weakness means that elevated status is one of the few remaining ways they can gain influence.[36] While Ward contends that rising powers should be preoccupied with status, others maintain that declining powers should be most sensitive to their relative rank because status can help offset drops in relative capabilities.[37] Ward admits that the literature has "not developed a comprehensive account of variation in the salience of status concerns" (p. 39).

The downside of rooting status in particular structural positions is that it becomes difficult to disentangle the influence of status from relative power. Great powers may be more assertive in defending their rank and rights, as Larson and Shevchenko suggest, but they may also be more assertive in general, regardless of the issue. Rising powers may have reasons to be more status conscious, as Ward claims, yet they also tend to have expanding interests and growing capabilities, both of which can correlate with aggressive behavior. "If power determines prestige," Mercer observes, "then distinguishing the concepts is pointless."[38] Status is then the language with which states discuss power positions. Renshon provides a way around this issue by underscoring the disjuncture between status and power. But Renshon's approach introduces a second issue: How do we identify inconsistencies or deficits in status? We have already noted the challenges in measuring status. Identifying status deficits requires an additional step: we must generate baseline expectations of rank based on material capabilities, which we then compare to those derived from status. Yet there are a range of contested ways to measure national power, doubling the difficulty of demonstrating where a state stands and making it hard to know whether a state is receiving too little, too much, or the proper amount of respect.

An alternative approach is to see status as triggered less by position than by humiliating events. Murray provides the best example of this approach in her discussion of what she calls "misrecognition" (p. 47).[39] When established powers "treat a rising power as an inferior actor" (p. 71), denying it the status that it believes it deserves, this is "experienced

[35] Karim 2018.
[36] de Carvalho and Neumann 2014; Wohlforth et al. 2018.
[37] Onea 2014, 135; Greve and Levy 2018, 156.
[38] Mercer 2017, 136.
[39] Ringmar 2014.

by the rising power as disrespect" (p. 73). Rising powers respond to humiliation by engaging in "forceful contestation with the established powers" designed to "compel these states to recognize its aspirant status" (p. 74). Ward offers a similar, though slightly different, account in his discussion of "status immobility" (pp. 3–4). When established powers engage in repeated "acts of status denial," it creates the perception that rising powers face a "glass ceiling" (p. 47). But rather than responding by trying to compete with the established powers on their own terms, Ward argues that humiliated states will embrace rejectionist policies that aim to "protest, delegitimize, or overthrow" the established international order (p. 51).

While Murray and Ward focus on cases of disrespect by established powers, other authors highlight how humiliating events can spark an obsession with status. Ayşe Zarakol argues that defeat in major wars can saddle states with shameful stigmas that they strive to correct.[40] Joslyn Barnhart finds that states that experience an involuntary territorial loss seek to restore their status through their own acts of territorial aggression.[41] Both Paul Saurette and Ahsan Butt contend the shock and humiliation of the 9/11 terrorist attacks prompted the United States to respond in an aggressive manner.[42] Although not a central part of their account, Larson and Shevchenko list "humiliating military defeats, exclusion from elite clubs, disregard for their interests, or economic difficulties" as the kinds of events that can prompt a state to seek to restore its tarnished status (p. 240). As with the structural positions described earlier, the list of events that could spark status seeking is long to the point of indeterminacy. If everything from defeat in major wars to economic fluctuations to unkind diplomatic exchanges is sufficient to activate status concerns, then humiliations are wildly over-predictive. There will always be some injury that scholars can point to after the fact to explain a state's behavior.

There are additional challenges in connecting status to cycles of action and reaction. One is that it can be difficult to separate insults and humiliations from regular hard bargaining. Ward cites Britain's refusal to accept "naval equality" with Germany prior to the First World War as evidence that London rejected Berlin's status aspirations (p. 83). Yet Britain had sensible strategic reasons to maintain naval supremacy. As Foreign Secretary Sir Edward Grey observed, "If the German Navy ever became superior to ours, the German Army can conquer this coun-

[40] Zarakol 2011, 11.
[41] Barnhart 2017.
[42] Saurette 2006, 512–21; Butt 2019, 268.

try. There is no corresponding risk of this kind to Germany."[43] Britain was hardly dismissive of German concerns. London engaged Berlin in arms control negotiations on multiple occasions—it simply refused to accept the demand that it remain neutral in a continental war in exchange for reductions in German shipbuilding. All of this is consistent with a traditional defense of British interests rather than with an attempt to denigrate Germany's status.

A second is that diplomats have strategic incentives to use the language of humiliation. They may be seeking a more favorable settlement or be signaling a willingness to use force. As Renshon acknowledges, German Foreign Minister Alfred von Kiderlen-Waechter confessed to "fanning the flames of nationalist fervor" during the Second Moroccan crisis in 1911 "in order to signal resolve" (p. 204). Appeals to status in these cases are a consequence of incompatible interests, not a cause of them. A final challenge is that states often respond to dramatic events not only because they are humiliating, but also because it is in their interest to do so. Larson and Shevchenko argue that Russia and China embraced military reforms after "humiliating military defeats" because these events were "damaging to morale" and "[made their] inferiority visible to others" (p. 24). Yet an equally compelling reason to embrace reforms is pragmatism. They can help to remedy institutional defects and to prevent future defeats.[44] Just because events are humbling does not mean that the actions that follow them are necessarily rooted in status.

Of course, it may well be that status concerns are catalyzed both by a state's position and by events. But there are tensions within and between positional and events-based explanations. There is no consensus logic for which structural positions accentuate status concerns, nor a convincing technique for separating status from power. Event-based explanations often boil down to humiliations, but what constitutes a humiliation is capacious and hard to extricate from international and domestic bargaining. And if states are doomed to seek status by position, then events are largely epiphenomenal. Yet if humiliating events drive status, then it becomes harder to predict in advance which kinds of states will be most obsessed with status.

MECHANISMS: LEADER PSYCHOLOGY AND DOMESTIC PATHOLOGIES

Once status concerns are activated, they must influence foreign policy decisions in some clear and consistent way. In general, scholars have focused on two distinct, although potentially connected, mechanisms.

[43] Grey 1908.
[44] MacDonald and Parent 2018, 29–32.

The first centers on the psychology, perceptions, and emotions of leaders. Drawing on Social Identity Theory (SIT), Larson and Shevchenko argue that "people derive part of their identity from membership in social groups" and that there exists "an innate human desire for one's group to be superior" (p. 3). Leaders respond to perceived inferiority by pursuing "identity management strategies" designed to "improve [their state's] standing" (p. 5). Renshon likewise argues that "leaders are typically assumed to identify with the status concerns of the states they represent" (p. 10). Once status concerns are triggered, leaders will attach an "increased value for status" (p. 60) and be willing to run greater risks to acquire it. He argues that this is particularly true for leaders with high Social Dominance Orientation (SDO), defined as "one's preference for (or comfort with) dominance and hierarchy" (p. 64).

There is no shortage of examples of prideful leaders around the world, so it makes sense that scholars would place them at the center of how status shapes foreign policy. But as Ward and others have pointed out, we must be careful about how we translate social psychology to world politics.[45] Many of these theories, such as SIT, were originally designed to understand individual attitudes toward group membership, not to provide a complete account of intergroup relations.[46] Thus, although studies have established that individuals exhibit in-group favoritism, there is little evidence that assignment to a group increases out-group antipathy or that the intensity of in-group bias is correlated with aggression.[47] Moreover, studies suggest that individuals in lower-status groups tend to display less in-group favoritism and to perceive more in-group variability "so as to mitigate the consequences of being tarred with the same brush."[48] How one translates these findings to world politics is unclear, but one possibility is that individuals in lower-status states would actually be less nationalistic and less invested in their state's status, a prediction that is at odds with Larson and Shevchenko and others. Although the alternatives might be circumscribed, individuals in lower-status states may identify more with subnational, regional, or ethnic identities or alternatively, with pan-national, religious, or even cosmopolitan identities, than with their nations.

A related issue concerns the levels-of-analysis problem. Theories that emphasize leadership psychology assume, as Renshon does, that leaders attach as much value to their state's reputation as they do to

[45] Ward 2017a; Hymans 2002.
[46] Ellemers and Haslam 2012, 386.
[47] Hymans 2002, 7–9.
[48] Brown 2000, 748, 751.

their own. But many leaders are more concerned with their personal prestige; they want to remain in office and reap the rewards of power. So if the pursuit of status abroad comes at the expense of prestige at home, most will likely choose the latter. There may even be some cases, as Rebecca Adler-Nissen emphasizes, in which international stigmas can be a source of individual pride, a sign that a leader is willing to defy the international community to defend the unique moral virtues of their state.[49]

More significantly, scholars have not provided clear guidance about what kinds of leaders will be most emotionally invested in status. Some allude to a leader's personal history. Murray references the Kaiser's penchant for delivering "excited" and "emphatic" speeches (pp. 106, 118), which some have attributed to his sense of insecurity after being born with a withered arm.[50] Others emphasize a leader's ideological commitments. Ward speculates that leaders who are "strong nationalists" will be particular invested in their state's rank (p. 55). Renshon emphasizes dispositional features of leaders' personalities, singling out "high-SDO subjects" as "particularly susceptible to the deleterious effects of status concerns" (p. 64). Whichever factor one focuses on, the challenge is how to separate a leader's dispositions and attitudes from his or her context and behavior. Recent studies, for example, have shown that SDO is "not a relatively stable, fixed individual difference variable" but a product of the "specific forms of group-based inequality" relevant to the respondent.[51] It is exceedingly difficult, in other words, to divorce individual dispositions from their social context. Leaders who are strong nationalists may be particularly sensitive to humiliations and thus more likely to clash with rivals, or they may be more likely to lash out at rivals simply because they are strong nationalists.

An alternative mechanism through which status can shape foreign policy is domestic politics. Ward provides the best example of this argument. He contends that hardliners can exploit perceived foreign insults to "undermine moderate leaders" and force a state to adopt "an aggressive, rejectionist foreign policy couched in the language of status" (p. 57). These dynamics are most likely to occur when leaders are "less secure from replacement by rivals" (p. 60) and when nationalists "represent a significant part of a leader's governing coalition" (p. 58). Although not the main focus of their theory, Larson and Shevchenko also hint at the importance of domestic motives for status seeking. They document how "the drive for political equality with the United States" was a "key ingredient . . . of domestic legitimacy for both Soviet and

[49] Adler-Nissen 2014, 153.
[50] Röhl 2015, 20–21.
[51] Turner and Reynolds 2003, 200, 202.

post-Soviet rulers" (p. 183). States pursue status not only because of co-alitional pressures, but also to cultivate broad public support.

The focus on domestic politics is compelling because it provides a plausible political account of how status concerns can shape the policy process. At the same time, elements of the domestic political story are underspecified. First, it has not been demonstrated that domestic constituencies place much stock in status. Most publics in most places tend to hold favorable views of their countries and suspicious views of other countries, while assessments of "national pride" tend to be driven more by domestic factors, such as levels of economic inequality, than international ones.[52] There are good reasons to believe that most people either do not know or do not care what other countries think about them.[53] Nor do public sentiments appear to be strongly tied to international triumphs or defeats. In an extensive statistical analysis, Andreas Wimmer finds that "countries that fought many wars with other states since 1816 are neither more nor less proud than more peaceful countries . . . nor are countries that lost those wars less proud."[54] Of course, there may be concentrated interest groups that worry about international rankings. Yet nationalist lobbies do not necessarily agree on which policies will improve their state's status, nor do they always possess sufficient influence to impose their preferences on policymakers. The literature posits large, status-sensitive constituencies, but this is more assumed than demonstrated.

Second, nationalist groups articulate collective grievances almost constantly, so it is not clear that status can explain why publics shift their support from moderates to hardliners. Interwar Germany provides a useful example. Ward argues that the Nazi party was able to exploit public anger about reparations, especially the 1929 Young Plan (pp. 150–51). However, the evidence for domestic outrage at this international insult is thin. While extremists rejected the plan, the Reichstag voted to accept it by a 318 to 82 margin, and a subsequent popular vote on the issue had low turnout (15 percent) but high approval (95 percent). Adolf Hitler did not take power until four years later, and foreign policy played a limited role. Indeed, as Jack Snyder points out, Hitler "soft-pedaled his *Lebensraum* theme in the crucial years when the Nazis were winning huge electoral successes."[55] Hardliners may try to exploit foreign policy setbacks, but their political success often depends more on their capacity to seize resources, attract recruits, build parties, and of-

[52] See, for example, Evans and Kelley 2002; Solt 2011.
[53] Mercer 2017, 168.
[54] Wimmer 2018, 223.
[55] Snyder 1991, 106.

fer compelling domestic programs. Status claims may follow nationalist groups' electoral or political successes without necessarily causing them.

Third, scholars have not specified which regimes produce the strongest pressures to pursue status. One might assume that democracies, in which leaders are responsive to domestic publics, may be most inclined to status seeking. Yet democratic institutions are also governed by norms of equality, which would seem to diminish the salience of arguments based on hierarchy and rank. Conversely, autocratic leaders tend to have the kinds of personalities, such as a high SDO, which would make them receptive to status arguments. At the same time, they are less beholden to their publics and are more invested in maintaining domestic control, which may be only loosely connected to their state's global standing. Ward's claim that insecure leaders who are beholden to nationalist parties are particularly prone to status seeking is reasonable, but borders on tautology.

In sum, each of the mechanisms linking status to foreign policy is plausible, but the literature lacks a developed theory of the interaction between leader psychology and domestic politics. Do leaders manipulate status to bolster their domestic authority? Do domestic hardliners force reluctant leaders to pursue status? Or do elites and masses sing from the same nationalist hymnal? Whether leaders are hypocrites, hostages, or true believers has profound implications, and the literature has muddied matters by portraying them as all three.

IV. The Effects of Status Seeking: Competition and Its Competitors

If states are compelled to seek status for wide-ranging reasons, it raises the question of what impact so doing has on world politics. For the most part, scholars worry that the pursuit of status will be destabilizing. While careful to note that the pursuit of status does not always generate competition or conflict, the consensus is that status impedes cooperation. The association of status seeking with conflict and instability, however, raises a number of issues. We elaborate two: the extent to which the pursuit of status demands competition and the extent to which status competition results in conflict.

How Pervasive Is Status Competition?

One issue the literature is divided on concerns the extent to which threats and the use of force are required to adjudicate status rankings. Some authors see a relatively tight connection between status and competition.

Robert Gilpin famously argued that prestige is "ultimately imponderable and incalculable," and is only known when "tested . . . on the field of battle."[56] Renshon provides compelling evidence to support this claim. His statistical results indicate that "even just the initiation of conflict—independent of outcome—boosts a state's status ranking by three ranks over the course of ten years" (p. 263). While less deterministic, Murray finds that rising powers that have been denied great power status often respond by investing in "exemplary military power" to "compel the recognition" they desire (p. 80). Such "struggles for recognition" can quickly devolve into arms races, territorial scrambles, containment, and war (p. 84).

But it is unclear how far one can push this logic. Most scholars acknowledge that states are unlikely to compete over status unless there is already some underlying conflict of interest.[57] Most accept that power political concerns will moderate status competition. States are unlikely to engage in competition if the prospects of victory are dim (Renshon, p. 168).[58] As a consequence, there is a tendency to focus on great power rivalries over status. But this creates a problem: great powers are more likely to compete in general, so it can be difficult to know whether competitive behaviors are a response to status or a byproduct of power. It is also unclear why being an aggressor should raise a state's esteem and why a reputation for achieving one's interests through skilled diplomacy would not be equally, if not more, valuable. As Renshon acknowledges (pp. 159–61), the meaning of victory and defeat varies depending on circumstances and who the relevant audiences are—and even when outcomes are clear, they can be paradoxical: overpowering a social inferior can diminish a state's status, while a valiant defeat at the hands of a social superior can augment it.[59] The meaning of conflict can similarly change with the broader normative context. If war comes to be seen as less legitimate, this diminishes, if not reverses, the status gains states can reap from engaging in it. As noted above, Wimmer finds no relationship between war and national pride, while Jennifer Miller and her coauthors find that contemporary states ascribe positive status to peers that respect human rights and engage in peaceful dispute resolution.[60] These are perfectly plausible findings, but ones that undermine the core logic of status competition.

[56] Gilpin 1983, 32–33.
[57] Wohlforth 2009, 39–40; Wohlforth 2014, 139.
[58] Greve and Levy 2018, 158; Barnhart 2016, 383.
[59] Johnson and Tierney 2006, 32–36.
[60] Miller et al. 2015.

Acknowledging these concerns, other authors allow for a wider range of responses to status anxieties. Larson and Shevchenko draw on SIT to argue that lower-status states can embrace multiple strategies to remedy their situation, including social mobility, where "aspiring states adopt the political, economic, and social norms of the dominant powers to be admitted to more prestigious institutions or clubs" (p. 6), and social creativity, where they "seek prestige in a different area . . . such as promoting international norms or a particular model for economic development" (p. 11). They argue that states will be drawn to more competitive strategies when elite clubs are impermeable and the status hierarchy is insecure, meaning it is perceived "to be *illegitimate* (unfair or unjust) and/or *unstable* (susceptible to change)" (p. 7). Ward concurs that there are multiple "logics of identity management" that lower-status states can adopt, some of which work within the normative constraints of the status hierarchy and others that reject it altogether (p. 49).

Although these approaches move us beyond status competition, they still do not provide a compelling account for why states choose some strategies over others. Part of the challenge here is translating findings from SIT, which studies individuals, to the realm of world politics, which concerns groups. Recent psychological studies suggest that the potential for social mobility "need not be very extensive" for actors to prefer individual adaptation to collective action in response to status deficits.[61] The most common psychological response to low status, in other words, is to distance oneself from a group rather than to fight for it, an option that may not be available to every individual, and a possibility at odds with the thrust of the status literature. More broadly, how do we know when elite clubs are permeable? How can we determine whether a status hierarchy is stable or unstable? For Larson and Shevchenko, the answer often comes down to "prevailing power relations" (p. 13). They point to the 2008 global financial crisis, for example, as a critical juncture that "undermined the stability and legitimacy of the status hierarchy" (p. 204). Yet this event also shook up the distribution of economic and military power, providing revisionist opportunities that were independent of status considerations. As Zarakol points out, how states respond to international stigmas can be shaped by many factors outside of the stability and legitimacy of status hierarchies, including the historical context, prevailing domestic conditions, and the rhetorical strategies adopted by political entrepreneurs.[62]

[61] Huddy 2001, 140.
[62] Zarakol 2011, 105–07.

How Destabilizing Is Status Competition?

Once states have chosen to compete over status, the literature assumes that the risks of conflict outbreak increase significantly. This pessimistic view is driven by two interrelated assumptions. The first is that status anxieties, once triggered, are difficult to alleviate. Leaders will cling to anger and resentment, hardliners will dominate domestic debates, and publics will rally around the flag, even as costs mount (Ward, pp. 59–61). The second is that international audiences will be unwilling or unable to address the status concerns of aggrieved states. Although most scholars agree that policies of "status accommodation" would ameliorate conflict, those at the top of status hierarchies are often reluctant to sacrifice their privileged positions (Larson and Shevchenko, p. 250; Murray, p. 202).[63] The perceived zero-sum character of status traps states in "status dilemmas," in which each side issues inflexible demands for recognition that results in a spiral of escalating tensions (Murray, p. 208).[64]

Though tenable, there are reasons to doubt whether these arguments hold as often as the literature asserts. First, states are often willing to abandon their status claims, especially when competition turns out to be costly. One of the curious features of the Anglo-German naval rivalry, which only Renshon stresses (pp. 214–15), is that it ended on relatively amicable terms. As Jan Rüger argues, the "naval race [was] effectively decided between 1909 and 1912, resulting in a more positive image of Germany in Britain." Indeed, in the two years prior to the war, Britain and Germany worked together to defuse a series of crises related to Portuguese colonies, the Baghdad railway, military advisors in Constantinople, and a host of other issues based on "a mutual feeling of responsibility."[65] Although Renshon may be overstating matters when he claims this proves that status competition paid dividends for Germany, it does suggest that while status anxieties may be acutely felt, they can be quickly forgotten.

Second, high-status states are often willing to accommodate status demands, especially if they can exchange recognition for political, economic, or institutional support. Although their focus is primarily on status competition, Larson and Shevchenko acknowledge that other states have accommodated Russian and Chinese status concerns on numerous occasions. Indeed, by our count Larson and Shevchenko cite at least nine examples in the post–Cold War period in which the international community recognized Russian and Chinese status concerns

[63] Paul 2016, 16.
[64] Wohlforth 2014, 114.
[65] Rüger 2011, 68.

(pp. 185, 191, 196–97, 206, 208–209, 216–17). As Phillip Lipscy and others have argued, there may be certain issues areas in which states have to sacrifice their privileged positions if they want to maintain institutionalized cooperation.[66] Clubs can be exclusive status markers, but they can also be useful mechanisms for peacefully distributing scarce resources. At the same time, recent Russian and Chinese revisionism raises questions about just how effective these acts of recognition have been, and Larson and Shevchenko concede that accommodation may fail if it does not meet stringent conditions, including that it is "made from a position of relative strength" (pp. 206, 219, 250).

All of this suggests a third point: although the notion of status dilemmas depends on the spiral model of conflict, it is equally plausible that status may operate based on the deterrence model. Refusing to accommodate status demands may dissuade states from seeking to overturn the status quo while accommodating status demands might prompt calls for more extreme forms of recognition. Consider British appeasement. As Stacie Goddard argues, Hitler legitimated his expansionist policies through status appeals, in particular to notions of "equality" and "self-determination," and British policymakers accommodated him because they perceived themselves to be "honest brokers" who were upholding the norms of the Versailles system.[67] Ward claims that it was the allies' refusal to acknowledge Germany's status that paved the way for Hitler's rise (pp. 156–57). But Goddard's evidence suggests the opposite: it was Britain's willingness to indulge German status demands that blinded them to the danger of Hitler's revisionism. It is worth noting that aggressors often use the language of status grievances. The Kaiser Wilhelms, Hideki Tojos, and Adolf Hitlers of the world have justified aggression by claiming that they had been disrespected, that they were the true victims, that they had no choice to but to lash out at their tormentors. We are under no obligation to take them at their word, nor to assume that accommodation would have satiated them.

POSSIBLE FUNCTIONS OF STATUS IN WORLD POLITICS

Status scholars have provided one possible pathway connecting status to conflict, but as the discussion above suggests, much of the basic theoretical and empirical work remains incomplete. We still do not know how often states issue status demands, how often other states accommodate or deny these demands, and how often these varied responses produce either competition or quiescence. All this raises the possibil-

[66] Lipscy 2017; Kruck and Zangl 2019.
[67] Goddard 2018, 135–37.

ity that status hierarchies may promote stability and cooperation under certain conditions. There are some grounds for this claim. Hegemonic stability theorists have long argued that the global economy functions best when there is a recognized leader who accepts the responsibility of providing public goods.[68] Recent studies suggest that states that sit in subordinate positions in global security and economic hierarchies tend to spend less on defense and are less likely to participate in militarized disputes.[69]

To explore the varied functions status can perform in world politics, consider two potential dimensions. The first is the *frequency of status inconsistencies*. In some cases, states will tend to be granted roughly the amount of status that their wealth or power would suggest, while in other cases they will be routinely deprived of the rank and recognition they believe they are entitled to. We anticipate states will make more frequent and intense status demands in the latter case. The second is the *character of responses to status*. Sometimes, status tends to produce peaceful responses: states either defer to high status or attempt to acquire or creatively redefine the attributes that deliver status. Other times, status tends to produce destabilizing reactions: states either embrace strategies of status competition or engage in violent rejection of status hierarchies.

Pulling these dimensions together, we can imagine four different ways status may shape world politics, which we present together in Table 1. When status inconsistencies are commonplace and status induces destabilizing responses (lower right-hand quadrant), status tends to have a particularly harmful impact on world politics. States will find themselves locked in deep and enduring rivalries over irreconcilable differences in status, along the lines described by Ward or Murray. When status is destabilizing yet status inconsistencies are relatively rare (upper right-hand quadrant), status can still be dangerous, but not necessarily debilitating. States may sometimes pick fights to bolster their status, as Renshon suggests. They may sometimes seize unimportant territories or participate in arms races to prove their place in the great power ranks. But most of the time, states will be afforded the status they believe they deserve, and competition among them will be shaped by factors unrelated to status.

Status hierarchies play a much different role if we move to the other side of the table. When status inconsistencies are rare and status prompts nonviolent responses (upper left-hand quadrant), status will tend to reflect and reinforce the distribution of power. States that pos-

---

[68] Lake 1993.
[69] Lake 2009; McDonald 2015.

TABLE 1
FUNCTIONS OF STATUS IN WORLD POLITICS

| | | Character of Status Responses | |
		Stabilizing	Destabilizing
Frequency of Status Inconsistencies	Infrequent	status reinforces the distribution of power (Gilpin)	status sometimes generates conflict (Renshon)
	Frequent	status challenges the distribution of power (Larson and Shevchenko?)	status frequently generates conflict (Ward, Murray)

sess wealth and power will tend to receive status, will assume the rights and responsibilities associated with being great powers, and subordinate states will tend to defer to them, as Gilpin suggests. Status in this world becomes a way in which systems of hegemony or great power cliques get normalized and legitimated. When status elicits peaceful responses yet status inconsistencies are common (lower left-hand quadrant), status operates as a kind of autonomous, normative standard that states use to evaluate and to challenge one another. Higher-status states claim special rights based on perceived moral, economic, or social superiority, while lower-status states seek to emulate or appropriate characteristics of their higher-status peers. Contestation in such a world will be driven less by geopolitical competition than by debates about how to construct status hierarchies—about what attributes should be valued and who should be included or excluded. Passages in Larson and Shevchenko suggest this kind of contestation, such as their discussion of the "new development models" offered by Deng Xiaoping and Mikhail Gorbachev (p. 135). Yet they describe these policies as reactions to prior failures of competitive policies (p. 138) rather than as choices shaped by the character of status itself. The bottom line is that status need not destabilize world politics, though more work needs to be done to clarify when and why this may be the case.

V. CONCLUSIONS

Over the last generation, there has been an explosion of scholarship on status. As demonstrated by the four books reviewed in this article, this work is theoretically ambitious, methodologically diverse, and rich

with insights. Nonetheless, there remain significant gaps in our under-standing of what status is, how it works, and what its effects are. While scholars have coalesced around a single definition of status, they have not resolved how membership and standing combine to create status, nor provided direct and compelling measures of status itself. They are also unclear about which states will be most interested in status and the pathways through which status shapes foreign policy. And schol-ars agree that status exerts a pernicious pull on world politics, although they disagree about the extent to which the pursuit of status requires competitive behavior and they neglect the ways in which status can de-fuse conflict and promote cooperation.

Paradoxically, then, the status of status is established, but not settled. It is established in the sense that it has arrived: the status literature is massive, is produced by scholars from around the world, and is widely recognized by the field. Yet it is unsettled in the sense that its coun-tervailing concepts and claims jostle for preeminence, and its stand-ing relative to other literatures, which wrestle with similar problems, remains undetermined. To its critics, status is an illusion, just another way of speaking about power. To its defenders, status supplements or supplants traditional theories, providing a unique perspective on what drives leaders and states. We have suggested that the empirical pos-sibilities are more varied, depending on the assumptions one makes about how frequently states make status claims and the extent to which competition is required to adjudicate them. The convergence of claims among the four books considered here highlights just how far the sta-tus literature has come, yet we have made the case that the tensions and inconsistencies in the literature deserve equal attention. The next steps are for status scholars to more clearly articulate their divergent theo-retical positions and to devise new empirical strategies to resolve their differences.

More specifically, we venture a few recommendations. First, scholars should tighten their definitions and explore other measures of status. Neither the qualitative nor quantitative measures developed in these books directly capture what status is or how it is expressed. We have little data about which attributes policymakers or publics tend to value, where they think their own state and other states rank on these attri-butes, and whether these assessments are consistent or vary from one country to another. Anecdotal evidence suggests that these kinds of as-sessments vary widely. In 2019, the Pew Research Center asked people in thirty-four countries who they considered the globe's top economic power; majorities in twenty-one countries named the United States,

while those in thirteen countries named China.[70] Along the same lines, we do not have systematic data about how often policymakers invoke status in their public or private statements. The data we do have suggests that status concerns are one among many. In his data set of the almost three thousand inferences British policymakers drew about the future behavior of European great powers in private diplomatic documents between 1855 and 1914, Robert Trager finds that only about 2 percent related to "prestige or to reputations for resolve."[71] Most important, we still do not have a clear sense of how often states defer voluntarily to those with higher status. Much turns on how we define voluntary deference and how we distinguish between decisions driven by coercion from those rooted in consent. Going forward, scholars should focus on identifying cross-national measures that more directly capture how policymakers and publics talk about status, and whether this actually translates into voluntary deference.

Second, scholars should do more to develop the specific causal mechanisms that connect status to foreign policy outcomes, paying particular attention to the conditions under which different mechanisms operate and how different mechanisms may interact with one another. One approach would be to focus on narrower events. When do wartime humiliations generate feelings of shame versus romanticization of a lost cause? When do diplomatic insults prompt resentment rather than indifference? Jennifer Lind, for example, examines the varied reactions to apologies, describing the conditions that can help them facilitate reconciliation.[72] An alternative approach would be to explore mechanisms related to particular kinds of status hierarchies. When is membership in the ranks of the great powers contested and when is it stable? When do states make explicit appeals to racial hierarchies and when do they gesture toward racial equality? Adom Getachew provides a compelling account of how principles of self-determination were reinvented to challenge racial hierarchies and spur anticolonial movements.[73] A third approach would be to use different methods to explore specific links in the causal chain. In a creative experiment, Renshon finds that the fear of losing status increases the tendency of actors to take risks in international crises, but only for "low-power" individuals (pp. 112–13). This suggests that psychological mechanisms work for some leaders in certain situations, but not others. Scholars could employ similar survey ex-

[70] Pew Research Center 2019.
[71] Trager 2017, 36.
[72] Lind 2010.
[73] Getachew 2019, 92–100.

periments of public opinion to identify scope conditions for domestic pressure and coalitional outbidding mechanisms.

Last, scholars need to better engage counterarguments. All the books under review take care to disentangle the effects of status from power and interests, but they could do more to explore the occasions when status does not work as we expect. When do states minimize or ignore status? When does the pursuit of status not exacerbate tensions, but instead facilitate cooperation? Because much of the literature focuses on the ways status precedes conflict, we know less about how status operates in other domains. Preliminary evidence suggests that it works much differently. In his study of socialization in international institutions, for example, Iain Johnston finds that China's desire to "maximize the normatively accepted markers of a high-status actor" prompted it to accept various multilateral arms control agreements.[74] Beth Simmons and Zachary Elkins present evidence that states are more likely to adopt liberal economic policies when high-income states and their cultural peers do so.[75] Wohlforth and his coauthors find that small states, such as Norway, often engage in "do-gooder status seeking," providing foreign aid or humanitarian assistance to bolster their moral authority.[76] Writing in a more speculative mode, Robert Keohane wonders whether creating an "economy of esteem" around greenhouse-gas mitigation could lead to more concerted action on climate change.[77] Taken together, these studies suggest that the initial emphasis on great powers and geopolitics may lose sight of the big picture. Status may appear zero-sum and competition for it may seem destabilizing. But this may be because power politics is a domain in which relative gains, fears of cheating, and tragic outcomes already predominate. Ultimately, status may buttress world order.

REFERENCES

Adler-Nissen, Rebecca. 2014. "Stigma Management in International Relations: Transgressive Identities, Norms, and Order in International Society." *International Organization* 68, no. 1: 143–76. At https://doi.org/10.1017/S0020818313000337.

Barnhart, Joslyn. 2016. "Status Competition and Territorial Aggression: Evidence from the Scramble for Africa." *Security Studies* 25, no. 3: 385–419. At https://doi.org/10.1080/09636412.2016.1195620.

[74] Johnston 2008, 197.
[75] Simmons and Elkins 2004.
[76] Wohlforth et al. 2018, 543.
[77] Keohane 2010, 21.

Barnhart, Joslyn. 2017. "Humiliation and Third-Party Aggression." *World Politics* 69, no. 3 (July): 532–68. At https://doi.org/10.1017/S0043887117000028.

Barnhart, Joslyn. 2020. *The Consequences of Humiliation: Anger and Status in World Politics*. Ithaca, N.Y.: Cornell University Press.

Brown, Rupert. 2000. "Social Identity Theory: Past Achievements, Current Problems, and Future Challenges." *European Journal of Social Psychology* 30, no. 6: 745–78. At https://psycnet.apa.org/doi/10.1002/1099-0992(200011/12)30:6%3C745::AID-EJSP24%3E3.0.CO;2-O.

Butt, Ahsan I. 2019. "Why Did the United States Invade Iraq in 2003?" *Security Studies* 28, no. 2: 250–85. At https://doi.org/10.1080/09636412.2019.1551567.

Carranza, Mario E. 2014. "Rising Regional Powers and International Relations Theories: Comparing Brazil and India's Foreign Security Policies and Their Search for Great-Power Status." *Foreign Policy Analysis* 13, no. 2: 255–77. At https://doi.org/10.1111/fpa.12065.

Charoenvattananukul, Peera. 2020. *Ontological Security and Status-Seeking: Thailand's Proactive Behaviours during the Second World War*. New York, N.Y.: Routledge.

Cline, Kirssa, Patrick Rhamey, Alexis Henshaw, Alesia Sedziaka, Aakriti Tandon, and Thomas J. Volgy. 2011. "Identifying Regional Powers and Their Status." In Thomas J. Volgy, Renato Corbetta, Keith A. Grant, and Ryan G. Baird, eds., *Major Powers and the Quest for Status in International Politics*. New York, N.Y.: Palgrave Macmillan.

Coggins, Bridget. 2014. *Power Politics and State Formation in the Twentieth Century: The Dynamics of Recognition*. New York, N.Y.: Cambridge University Press.

Cooley, Alexander, and Jack Snyder, eds. 2015. *Ranking the World: Grading States as a Tool of Global Governance*. New York, N.Y.: Cambridge University Press.

Dafoe, Allan, Jonathan Renshon, and Paul Huth. 2014. "Reputation and Status as Motives for War." *Annual Review of Political Science* 17, no. 1: 371–93. At https://doi.org/10.1146/annurev-polisci-071112-213421.

de Carvalho, Benjamin, and Iver B. Neumann, eds. 2014. *Small States and Status Seeking: Norway's Quest for International Standing*. New York, N.Y.: Routledge.

Deng, Yong. 2008. *China's Struggle for Status: The Realignment of International Relations*. New York, N.Y.: Cambridge University Press.

Destradi, Sandra, and Johannes Plagemann. 2019. "Populism and International Relations: (Un)predictability, Personalisation, and the Reinforcement of Existing Trends in World Politics." *Review of International Studies* 45, no. 5: 711–30. At https://doi.org/10.1017/S0260210519000184.

Drezner, Daniel W. 2020. *The Toddler in Chief: What Donald Trump Teaches Us about the Modern Presidency*. Chicago, Ill.: University of Chicago Press.

Duque, Marina G. 2018. "Recognizing International Status: A Relational Approach." *International Studies Quarterly* 62, no. 3: 577–92. At https://doi.org/10.1093/isq/sqy001.

Ellemers, Naomi, and S. Alexander Haslam. 2012. "Social Identity Theory." In Paul A. M. Van Lange, Arie W. Kruglanski, and E. Tory Higgins, eds, *Handbook of Theories of Social Psychology*. London, UK: SAGE Publications: 379–98.

Evans, M. D. R., and Jonathan Kelley. 2002. "National Pride in the Developed World: Survey Data from 24 Nations." *International Journal of Public Opinion Research* 14, no. 3: 303–38. At https://doi.org/10.1093/ijpor/14.3.303.

Fordham, Benjamin O. 2011. "Who Wants to Be a Major Power? Explaining the Expansion of Foreign Policy Ambition." *Journal of Peace Research* 48, no. 5: 587–603. At https://doi.org/10.1177%2F0022343311411959.

Getachew, Adom. 2019. *Worldmaking after Empire: The Rise and Fall of Self-Determination*. Princeton, N.J.: Princeton University Press.

Gilady, Lilach. 2018. *The Price of Prestige: Conspicuous Consumption in International Relations*. Chicago, Ill.: University of Chicago Press.

Gilpin, Robert. 1983. *War and Change in World Politics*. New York, N.Y.: Cambridge University Press.

Goddard, Stacie E. 2018. *When Right Makes Might: Rising Powers and World Order*. Ithaca, N.Y.: Cornell University Press.

Götz, Elias. 2019. "Strategic Imperatives, Status Aspirations, or Domestic Interests? Explaining Russia's Nuclear Weapons Policy." *International Politics* 56, no. 6: 810–27. At https://doi.org/10.1057/s41311-018-0168-7.

Graz, Jean-Christophe. 2003. "How Powerful Are Transnational Elite Clubs? The Social Myth of the World Economic Forum." *New Political Economy* 8, no. 3: 321–40. At https://doi.org/10.1080/1356346032000138041.

Greve, Andrew Q., and Jack S. Levy. 2018. "Power Transitions, Status Dissatisfaction, and War: The Sino-Japanese War of 1894–1895." *Security Studies* 27, no. 1: 148–78. At https://doi.org/10.1080/09636412.2017.1360078.

Grey, Edward. 1908. "Memorandum (31 July)". In G. P. Gooch and H. W. V. Temperley, eds., *British Documents on the Origins of the War 1898–1914*, vol. 6, appendix 3. London, UK: His Majesty's Stationary Office.

Hafner-Burton, Emilie M., and Alexander H. Montgomery. 2006. "Power Positions: International Organizations, Social Networks, and Conflict." *Journal of Conflict Resolution* 50, no. 1: 3–27. At https://doi.org/10.1177%2F0022002705281669.

Hall, Todd H. 2017. "On Provocation: Outrage, International Relations, and the Franco–Prussian War." *Security Studies* 26, no. 1: 1–29. At https://doi.org/10.1080/09636412.2017.1243897.

Heimann, Gadi. 2015. "What Does It Take to Be a Great Power? The Story of France Joining the Big Five." *Review of International Studies* 41, no. 1: 185–206. At https://doi.org/10.1017/S0260210514000126.

Huddy, Leonie. 2001. "From Social to Political Identity: A Critical Examination of Social Identity Theory." *Political Psychology* 22, no. 1: 127–56. At https://doi.org/10.1111/0162-895X.00230.

Hymans, Jacques E. C. 2002. "Applying Social Identity Theory to the Study of International Politics: A Caution and an Agenda." Paper presented at the annual meeting of the International Studies Association, New Orleans, March 24–27.

Johnson, Dominic D. P., and Dominic Tierney. 2006. *Failing to Win: Perceptions of Victory and Defeat in International Politics*. Cambridge, Mass.: Harvard University Press.

Johnston, Alastair Iain. 2008. *Social States: China in International Institutions, 1980–2000*. Princeton, N.J.: Princeton University Press.

Karim, Moch Faisal. 2018. "Middle Power, Status-Seeking and Role Conceptions: The Cases of Indonesia and South Korea." *Australian Journal of International Affairs* 72, no. 4: 343–63. At https://doi.org/10.1080/10357718.2018.1443428.

Keohane, Robert O. 2010. "The Economy of Esteem and Climate Change." *St. Anthony's International Review* 5, no. 2: 16–28.

Kruck, Andreas, and Bernhard Zangl. 2019. "Trading Privileges for Support: The Strategic Co-optation of Emerging Powers into International Institutions." *International Theory* 11, no. 3: 318–43. At https://doi.org/10.1017/S1752971 919000101.

Lake, David A. 1993. "Leadership, Hegemony, and the International Economy: Naked Emperor or Tattered Monarch with Potential?" *International Studies Quarterly* 37, no. 4: 459–89. At https://doi.org/10.2307/2600841.

Lake, David A. 2009. *Hierarchy in International Relations*. Ithaca, N.Y.: Cornell University Press.

Lake, David A. 2014. "Status, Authority, and the End of the American Century." In T. V. Paul, Deborah Welch Larson, and William C. Wohlforth, eds., *Status in World Politics*, 246–70. New York, N.Y.: Cambridge University Press.

Lamont, Michèle, Joshua A. Guetzkow, Hanna Herzog, Graziella Moraes Dias da Silva, Nissim Mizrachi, Elisa Pereira Reis, and Jessica S. Welburn, eds. 2016. *Getting Respect: Responding to Stigma and Discrimination in the United States, Brazil, and Israel*. Princeton, N.J.: Princeton University Press.

Larson, Deborah Welch, T. V. Paul, and William C. Wohlforth. 2014. "Status and World Order." In T. V. Paul, Deborah Welch Larson, and William C. Wohlforth, eds., *Status in World Politics*. New York, N.Y.: Cambridge University Press.

Larson, Deborah Welch, and Alexei Shevchenko. 2019. *Quest for Status: Chinese and Russian Foreign Policy*. New Haven, Conn.: Yale University Press.

Lebow, Richard Ned. 2010. *Why Nations Fight: Past and Future Motives for War*. New York, N.Y.: Cambridge University Press.

Lind, Jennifer M. 2010. *Sorry States: Apologies in International Politics*. Ithaca, N.Y.: Cornell University Press.

Lipscy, Phillip Y. 2017. *Renegotiating the World Order: Institutional Change in International Relations*. New York, N.Y.: Cambridge University Press.

MacDonald, Paul K., and Joseph M. Parent. 2018. *Twilight of the Titans: Great Power Decline and Retrenchment*. Ithaca, N.Y.: Cornell University Press.

McDonald, Patrick J. 2015. "Great Powers, Hierarchy, and Endogenous Regimes: Rethinking the Domestic Causes of Peace." *International Organization* 69, no. 3: 557–88. At https://doi.org/10.1017/S0020818315000120.

Mercer, Jonathan. 2017. "The Illusion of International Prestige." *International Security* 41, no. 4: 133–68. At https://doi.org/10.1162/ISEC_a_00276.

Miller, Jennifer L., Jacob Cramer, Thomas J. Volgy, Paul Bezerra, Megan Hauser, and Christina Sciabarra. 2015. "Norms, Behavioral Compliance, and Status Attribution in International Politics." *International Interactions* 41, no. 5: 779–804. At https://doi.org/10.1080/03050629.2015.1037709.

Miller, Manjari Chatterjee. 2014. *Wronged by Empire: Post-Imperial Ideology and Foreign Policy in India and China*. Stanford, Calif.: Stanford University Press.

Murray, Michelle K. 2019. *The Struggle for Recognition in International Relations: Status, Revisionism, and Rising Powers*. New York, N.Y.: Oxford University Press.

Nolte, Detlef. 2010. "How to Compare Regional Powers: Analytical Concepts and Research Topics." *Review of International Studies* 36, no. 4: 881–901. At https://doi.org/10.1017/S026021051000135X.

Onea, Tudor A. 2014. "Between Dominance and Decline: Status Anxiety and Great Power Rivalry." *Review of International Studies* 40, no. 1: 125–52. At https://doi.org/10.1017/S0260210512000563.

Paul, T. V. 2016. "The Accommodation of Rising Powers in World Politics." In T. V. Paul, ed., *Accommodating Rising Powers*. New York, N.Y.: Cambridge University Press.

Petersen, Roger Dale. 2002. *Understanding Ethnic Violence: Fear, Hatred, and Resentment in Twentieth-Century Eastern Europe*. New York, N.Y.: Cambridge University Press.

Pew Research Center. 2019. "Views of the Balance of Power between U.S. and China." December 5. At https://www.pewresearch.org/global/2019/12/05/views-of-the-balance-of-power-between-u-s-and-china-2019/.

Pouliot, Vincent. 2016. *International Pecking Orders: The Politics and Practice of Multilateral Diplomacy*. New York, N.Y.: Cambridge University Press.

Pu, Xiaoyu. 2019. *Rebranding China: Contested Status Signaling in the Changing Global Order*. Stanford, Calif.: Stanford University Press.

Putin, Vladimir. 2014. "Address by President of the Russian Federation." March 18. At http://en.kremlin.ru/events/president/news/20603.

Renshon, Jonathan. 2017. *Fighting for Status: Hierarchy and Conflict in World Politics*. Princeton, N.J.: Princeton University Press.

Ringmar, Erik. 2014. "International Politics of Recognition." In Erik Ringmar and Thomas Lindemann, eds., *The International Politics of Recognition*. New York, N.Y.: Paradigm.

Röhl, John C. G. 2015. *Young Wilhelm: The Kaiser's Early Life, 1859–1888*. New York, N.Y.: Cambridge University Press.

Røren, Pål, and Paul Beaumont. 2019. "Grading Greatness: Evaluating the Status Performance of the BRICS." *Third World Quarterly* 40, no. 3: 429–50. At https://doi.org/10.1080/01436597.2018.1535892.

Rüger, Jan. 2011. "Revisiting the Anglo-German Antagonism." *Journal of Modern History* 83, no. 3: 579–617. At https://doi.org/10.1086/660841.

Sambanis, Nicholas, Stergios Skaperdas, and William C. Wohlforth. 2015. "Nation-Building through War." *American Political Science Review* 109, no. 2: 279–96. At https://doi.org/10.1017/S0003055415000088.

Saurette, Paul. 2006. "You Dissin me? Humiliation and Post 9/11 Global Politics." *Review of International Studies* 32, no. 3: 495–522. At https://doi.org/10.1017/S0260210506007133.

Simmons, Beth A., and Zachary Elkins. 2004. "The Globalization of Liberalization: Policy Diffusion in the International Political Economy." *American Political Science Review* 98, no. 1: 171–89. At https://doi.org/10.1017/S0003055404001078.

Snyder, Jack L. 1991. *Myths of Empire: Domestic Politics and International Ambition*. Ithaca, N.Y.: Cornell University Press.

Solt, Frederick. 2011. "Diversionary Nationalism: Economic Inequality and the Formation of National Pride." *Journal of Politics* 73, no. 3: 821–30. At https://doi.org/10.1017/S002238161100048X.

Suzuki, Shogo. 2008. "Seeking 'Legitimate' Great Power Status in Post-Cold War International Society: China's and Japan's Participation in UNPKO." *International Relations* 22, no. 1: 45–63. At https://doi.org/10.1177%2F0047117807087242.

Wolf, Reinhard. 2011. "Respect and Disrespect in International Politics: The Significance of Status Recognition." *International Theory* 3, no. 1: 105–42. At https://doi.org/10.1017/S1752971910000308.

Wolf, Reinhard. 2014. "Rising Powers, Status Ambitions, and the Need to Reassure: What China Could Learn from Imperial Germany's Failures." *Chinese Journal of International Politics* 7, no. 2: 185–219. At https://doi.org/10.1093/cjip/pou029.

Xi, Jinping. 2019. "Speech at the Reception in Celebration of the 70th Anniversary of the Founding of the People's Republic of China." September 30. At https://www.fmprc.gov.cn/mfa_eng/wjdt_665385/zyjh_665391/t1704400.shtml.

Zarakol, Ayşe. 2011. *After Defeat: How the East Learned to Live with the West.* New York, N.Y.: Cambridge University Press.

AUTHORS

PAUL K. MACDONALD is an associate professor of political science at Wellesley College. His books include *Networks of Domination: The Social Foundations of Peripheral Conquest in International Politics* (2014) and *Twilight of the Titans: Great Power Decline and Retrenchment* (2018), coauthored with Joseph M. Parent. He can be reached at pmacdona@wellesley.edu.

JOSEPH M. PARENT is a professor of political science at the University of Notre Dame. His books include *Uniting States: Voluntary Union in World Politics* (2011) and *Twilight of the Titans: Great Power Decline and Retrenchment* (2018), coauthored with Paul K. MacDonald. He can be reached at jparent@nd.edu.

ACKNOWLEDGEMENTS

KEY WORDS

Cambridge Core

For further information about this journal,
go to the publisher's journal website at
cambridge.org/wpo

CAMBRIDGE
UNIVERSITY PRESS