

The Supply of Conspiracy Theories in State-Controlled Media*

Gabriel Koehler-Derrick[†]

Richard A. Nielsen[‡]

David Romney[§]

May 14, 2021

Abstract

When, and why, do conspiracy theories appear in state-controlled media? We argue that regime type and political threats combine to explain the supply of conspiracy theories by the state. Regimes under threat can stave off dissent by promoting conspiracy theories but this can backfire and damage regime legitimacy. These costs make governments that do not feel threatened unlikely to circulate conspiracy theories, while governments that feel threatened will. Given the same level of threat, democratic governments supply fewer conspiracy theories than autocrats because they are more concerned with legitimacy. We test this argument by comprehensively examining conspiracy theories in Egypt's primary state-controlled newspaper between 1998 and 2018. The state-controlled newspaper responds to threatening events by publishing more conspiracy theories than its independent counterpart. However, this relationship is moderated by changes in regime: the state-controlled media promoted fewer conspiracy theories during periods of low threat and democracy.

Abstract word count: 146

Document word count (text, footnotes, references): 11,861

*Thanks to Ala' Alrababa'h, Lisa Blaydes, Melani Cammett, Daniel Carpenter, Fotini Christia, Moustafa Hefny, Amaney Jamal, Brendan McElroy, and audiences at the AALIMS conference, the University of Cairo, and the University of Colorado Boulder for feedback and suggestions. Amani Selmi, Mehdy Elloumi, and Elisha Cohen provided excellent research assistance.

[†]Postdoctoral Fellow, Brown University, gabriel_koehler-derrick@brown.edu

[‡]Associate Professor, Massachusetts Institute of Technology, rnielsen@mit.edu. Corresponding author.

[§]Postdoctoral Fellow, Harvard University, dromney@fas.harvard.edu

1 Introduction

On August 14th, 2013, Egyptian security forces killed approximately 1,000 demonstrators protesting the ouster of Egypt's first democratically-elected President, Mohamed Morsi in Cairo's Raba'a square. Thirteen days later, the state-run newspaper, *Al-Ahram*, ran a front-page headline alleging a "New Conspiracy to Undermine Stability: Politicians, Journalists, and Businessmen Involved."¹ Citing unnamed official sources, the article outlined a plot organized by unlikely conspirators — United States' Ambassador to Egypt, Anne Peterson, and Muslim Brotherhood leader Khayrat al-Shater — to smuggle "300 armed men from Gaza to Egypt through tunnels, to spread chaos throughout Cairo and storm numerous prisons." This presents a puzzle: Why would the state-controlled newspaper promote a preposterous conspiracy theory certain to increase diplomatic tension between the United States and Egypt during an already tense time?

In this paper, we offer theory and evidence to explain the supply of conspiracy theories in state-controlled news media. A surge of recent scholarship examines the politics of conspiracy theories, mostly using surveys and experiments, to understand citizens' demand for conspiracy theories. This demand-side literature is approaching consensus that conspiracy theories serve as heuristics to cope with uncertainty (Oliver and Wood, 2014) and that psychological traits such as mistrust, belief in unseen forces, and propensity to see patterns in randomness largely explain individuals' predisposition to conspiricism (Radnitz and Underwood, 2017; van Prooijen, Douglas and De Inocencio, 2018). These psychological traits are common in individuals worldwide, so it seems unlikely that variation in individual demand alone can fully explain the wide variation in the prevalence of conspiracy theories across political contexts. Instead, we argue that variation in the supply of conspiracy theories by various actors, including the state, is essential to understanding why conspiracy theories are more politically relevant in some places than others.

The supply of political misinformation has also received scholarly attention (Peisakhin and Rozenas, 2018; Rozenas and Stukal, 2019), but there is less consensus about why states promote conspiracy theories. There is evidence that some regimes supply more conspiracy theories than others, but scholars disagree about whether this is best explained by the psychology of political leaders (Gray, 2010*b*), political culture (Hofstadter, 1965), autocratic political institutions (Gray, 2010*a*), or diversionary incentives (Alrababa'h and Blaydes, 2020). Many of these studies rely on anecdotal examination of elaborate, pervasive, or implau-

¹<https://gate.ahram.org.eg/daily/News/228698.aspx>, accessed 3/22/2021.

sible conspiracy theories. While this approach may have advantages, it does not allow systematic analysis of changes in state supply over time. A few recent studies show the promise of systematically analyzing trends in comprehensive corpora from state-controlled media (Rozenas and Stukal, 2019; Alrababa'h and Blaydes, 2020), and we follow their lead.

The supply-side literature does not often compare state-controlled and independent media, but we make this comparison central to our approach. The presence of conspiracy theories in state-controlled media may not always be evidence of regime manipulation; state-controlled media outlets may be responding to factors that lead all media outlets to produce conspiratorial content. For this reason, analyzing state-controlled media in isolation makes it challenging to infer when conspiracy theories appear because of regime manipulation or for other reasons. Ideally, to learn about regime strategy, we would focus only on the conspiracy theories that appear in official media at the behest of the regime, but this is impossible without intimate knowledge of private government communication to editors and journalists.

To resolve this inferential challenge, we compare the supply of conspiracy theories in state-controlled media and independent media. By controlling for conspiracy theory supply in independent media, we estimate government *oversupply*: when official media produces conspiracy theorizing in excess of the independent media under the same political conditions. Our concept of oversupply suggests a mirror-image possibility that the literature has largely ignored: that states can also *undersupply* conspiracy theories relative to independent media for political reasons. Of course, under authoritarianism no paper is truly independent from official influence. However, because some media enjoy *relatively* more independence than their official counterparts, the difference between the two that we estimate with our approach is indicative of oversupply or undersupply.

We develop theory and evidence about the oversupply and undersupply of conspiracy theories in state-controlled media. Like many of the scholars who precede us, we argue that the state supply of conspiracy theories depends on political conditions. We build on the idea that autocratic governments face incentives to circulate conspiracy theories to divert from poor performance (Rozenas and Stukal, 2019) and threat (Alrababa'h and Blaydes, 2020). However, we incorporate recent findings showing that state propaganda and information manipulation can backfire in the short-term (Hobbs and Roberts, 2018) and degrade public perceptions and trust in the long-term (Huang, 2018). We argue that these side-effects raise the cost of

promoting conspiracy theories for states. In the absence of threats, the costs of conspiracy theory promotion are likely to outweigh the benefits, leading official media to undersupply conspiracy theories relative to independent media. However, when events threaten incumbent elites, they are likely to oversupply conspiracy theories for the short-term benefits, regardless of the long-term costs. Governments that feel a high baseline level of threat—from past experience or from observing the experience of prior regimes—are more likely to turn to conspiracism when threatened. Governments that feel a relatively low baseline level of threat should be slower to switch from undersupply to oversupply.

We test these arguments by examining state supply of conspiracy theories in Egypt between 1998 and 2018. We collect all of the articles in Egypt's main state-controlled newspaper, *Al-Ahram*, and compare the prevalence, timing, and content of conspiracy theories to Egypt's most prominent independent newspaper, *Al-Masry Al-Youm*. Identifying conspiracy theories in over one million newspaper articles is a formidable challenge that we overcome by combining qualitative coding with automated text analysis methods (Lucas et al., 2015).

Dramatic, unexpected developments in Egyptian politics over the last two decades allow us to plausibly estimate the effects of regime type and political threat on the state supply of conspiracy theories in Egypt. In the twenty years we observe, Egypt was a stable autocracy under president Hosni Mubarak, an unstable autocracy under military rule after the 2011 revolution forced Mubarak out, an unstable democracy under the Muslim Brotherhood, and an unstable and repressive autocracy following a military coup in 2013. The sudden changes in Egyptian politics over this time surprised virtually all observers, including regime elites. We cannot experimentally manipulate regime type or threats to the regime, but we believe that the unexpected changes in regime type and threat that occurred in Egypt allow us to estimate credible causal effects from our observational data. Readers who are skeptical about giving a causal interpretation to our results should nevertheless be interested in our descriptive findings, which provide a complete record of the conspiracy theories that the Egyptian government has promoted through its official newspaper over the last twenty years.

We find that the state-controlled newspaper undersupplies conspiracy theories relative to the independent newspaper when regime stability is not threatened. When the number of stability-threatening events increases, the regime begins to oversupply conspiracy theories. The degree of oversupply varies

by administration—most of the state over-supply of conspiracy theories in our data has occurred under Egypt’s current president, Abdel Fattah Al-Sisi. By contrast, the autocratic government of Hosni Mubarak undersupplied conspiracy theories most of the time, supporting our argument that even autocrats are likely to undersupply conspiracy theories when they face no appreciable threats because of the high reputational costs of conspiracism.

After examining the quantity of conspiracy theories that the state supplies, we turn to examining them qualitatively. We use nested analysis (Lieberman, 2005), sampling thirty days worth of oversupplied conspiracy theories for close analysis based on the predictions of our regression model. We find that conspiracies in the official newspaper rarely identify specific perpetrators, victims, or incidents, and when they do, they are often only tenuously connected to recent events. Rather, these conspiracies seem designed to increase a sense of threat and foreboding, portraying the state as the only defense against chaos and division sown by nefarious enemies, mostly unnamed, outside and inside Egypt.

Beyond Egypt, our theory offers a possible explanation for the supply of conspiracy theories by other states, both autocratic and democratic. Our approach opens new possibilities for measuring and predicting the supply of conspiracy theories by various actors in other countries and contexts.

2 Conspiracy Theories: Supply and Demand

Longstanding scholarly interest in conspiracy theorizing (Popper, 1945; Hofstadter, 1965), has been reinvigorated by new interest in political misinformation (Berinsky, 2017) and “fake news” (Lazer et al., 2018). One of the most foundational contributions of this recent literature is a definition of conspiracy theories that we adopt directly from Oliver and Wood (2014). We consider conspiracy theories to be a form of political discourse with three key features: (1) a propensity to “locate the source of unusual social and political phenomena in unseen, intentional, and malevolent forces,” (2) a propensity to interpret events “in terms of a Manichean struggle between good and evil,” and (3) the implication that “mainstream accounts of political events are a ruse” (953).

Scholars in a range of disciplines, including political science and political psychology (Goertzel, 1994; Wood, 2015), have studied the causes and effects of individual belief in conspiracy theories using surveys and experiments. Some focus on individual characteristics that make some people more or less receptive

to conspiracy theories (Wood, Douglas and Sutton, 2012), while others examine the impact of conspiracy theories on beliefs and attitudes (Dixon and Jones, 2015) or behavior (Einstein and Glick, 2014).

Together, these studies suggest that some individuals are attracted to conspiracy theories for psychological reasons. There is still debate about the individual traits that correlate with a greater receptiveness to conspiracy theories, but psychological traits (Radnitz and Underwood, 2017; van Prooijen, Douglas and De Inocencio, 2018), partisanship (Enders, Smallpage and Lupton, 2020; Ryan and Aziz, 2020), and gender (Bruder et al., 2013; Cassese, Farhart and Miller, 2020) feature prominently in the literature. Despite disagreements, scholars generally accept that conspiracy theories provide a mechanism for establishing a sense of order and purpose in the face of complex events that are difficult to understand (Oliver and Wood, 2014).

The Demand for Conspiracy Theories

Prior scholarship establishes that belief in conspiracy theories is widespread. According to survey results reported by Oliver and Wood (2014), “half of the American public consistently endorses at least one conspiracy theory.” Egypt, the country we study, is no exception. Nyhan and Zeitzoff (2018) find that “adherence to conspiracy theories is widespread” in their sample of 2,015 respondents in Egypt and Saudi Arabia, with more than 80 percent of respondents believing at least two conspiracy theories. Despite empirical evidence that conspiracy theories are common across cultures, scholars of the Middle East and North Africa sometimes claim that conspiracy theories are particularly prevalent in the region for cultural reasons. These claims are based on conjecture and anecdotes, rather than systematic comparison across countries and regions; the biased tenor of some of these claims may even have stifled more systematic research on the question.² The implicit companion claim is that the supply of conspiracy theories in Middle East news media is simply a function of demand. If Arabs, or Muslims, or Middle Easterners are particularly prone to conspiracy theorizing, the story goes, then it is no surprise to find conspiracy theories widely circulating in the press. Without more rigorous analysis, anyone interested in the subject risks coming away with the impression that culture is the primary explanation conspiracy theories in the MENA region.

Perhaps Middle Eastern culture is particularly prone to conspiricism, but we see scant evidence for

²According to Brynjar Lia, “The famous and notorious Daniel Pipes, a very prominent right-wing pundit and activist. . . [is] author of one of the most cited books on conspiracy theories in the region. And obviously, the theme has become a bit toxic. . . because if you write about it, in a sense, you have to cite Daniel Pipes, and who’d want to have Daniel Pipes in his footnotes?” (Nordic Society for Middle East Studies Conference, 2019, <https://www.youtube.com/watch?v=xaj7fQ2S424>, minute 23:18-23:40, accessed 3/22/2021).

this. Belief in conspiracy theories is high in some countries in the region (Bruder et al., 2013; Nyhan and Zeitzoff, 2018), but it is also high elsewhere (Oliver and Wood, 2014). We expect that the supply of conspiracy theories in the Middle East is not just meeting cultural demand, but is instead shaped by politics.

The Supply of Conspiracy Theories

The surge in scholarship on conspiracy theories has focused primarily on demand, but important studies also examine their supply, especially by states promoting misinformation to their own citizens. The literature broadly agrees that regime type matters; autocratic governments are more likely to manipulate the media environment generally (Black, 2008; Peterson, 2011), and to use state-controlled media to promote propaganda (Peisakhin and Rozenas, 2018; Rozenas and Stukal, 2019). Following Gray (2010*a*, Chapter 4), we organize our discussion around three main reasons for state supply of conspiracy theories: genuine belief, national narratives, and diversion.

Some leaders of states genuinely believe they were being conspired against (Gray, 2010*b*, 32). In the Middle East, the long list of leaders who have lost power at the hands of foreign and domestic conspirators suggests that this belief is not irrational. Autocrats may communicate the conspiracies they fear to explain policies aimed at stopping these conspiracies, especially if these policies are unpopular, repressive, or violent.

Alternatively, regimes might promote conspiracy theories as a way of crafting national narratives that enhance the state's symbolic power (Gray, 2010*a*, 133-134). For example, (Yablokov, 2015) argues that Russia's media conglomerate, Russia Today, is a tool of public diplomacy used to convince audiences of anti-Western and pro-Russian ideas. State-promoted conspiracy theories do not need to be convincing in order to sustain symbolic state power. Wedeen (2019) argues that the Syrian government's conspiracy theories are not convincing to regular Syrians but are instead symbols "specifying the form and content of civic obedience" (Wedeen, 2019, viii). Conspiracy theories can also communicate political priorities to elites, as well as the masses, and may assist coordination among political elites (Radnitz, 2016).

Finally, conspiracy theories may be a diversionary tactic: a cynical but rational response of autocrats when "calls for reform or democratization need to be silenced" (Gray, 2010*a*, 120, 130). Facing threatening circumstances at home, the state may promote conspiracy theories to explain its failings and direct rage away from the government and toward an alternative source, creating a "channel for popular disquiet or mistrust"

(121). Examples abound. Rozenas and Stukal (2019) conclude from daily news reports on Russia's largest state-owned television network that "bad news is not censored, but it is systematically blamed on external factors, whereas good news is systematically attributed to domestic politicians." Alrababa'h and Blaydes (2020) examine the Syrian state-controlled media from 1987-2018 and show that diversionary claims of external threats are common: "the Assad regime long sought to focus public attention on forces external to the regime, consistent with a logic of diversionary threat" (2).

There is reason to believe that the diversionary rhetoric might work in the short term. Experimental evidence shows that support for Egyptian government repression increases support when it accuses its opponents of conspiring to commit violence (Williamson and Malik, 2020). Even when diversionary conspiracy theories are not believed, they prop up authoritarian governments by sowing fear. Huang (2015) shows that Chinese students exposed to state propaganda are still dissatisfied with the government but, nevertheless, become unwilling to engage in dissent because they believe the regime is strong. Wedeen (2019, Chapter 4) similarly argues that the circulation of conspiracy theories during the 2011 Syrian uprising and civil war helped the regime by making citizens ambivalent. Conspiracy theories provided "evidence" to support contradictory narratives of the war, leaving open the possibility that the regime deserved support, or at least passivity, even as repression mounted.

Autocratic diversion is the most prevalent explanation for conspiracy theory supply, but it faces theoretical challenges. Research suggests that even autocrats pay a price when they indulge in too much conspiracy theorizing (Huang, 2018). Conspiracy theories may render citizens passive, but they also spread narratives that the regime is beset on all sides by powerful conspirators, which may undermine regime goals of projecting strength. And blaming foreign conspirators can backfire diplomatically; Alrababa'h and Blaydes (2020) find that anti-Israel propaganda appears less often on the front page of Syria's state-controlled newspaper during periods where the Syrian government is in political negotiations with Israel.

The diversion argument is also difficult to test because comprehensively identifying conspiracy theories in state-controlled media is challenging. Alrababa'h and Blaydes (2020), for example, work with an impressive corpus of Syrian state-controlled media, but are forced to hand-code a single theme on the front page for the early years because the text is not digitized. In later years, they are able to obtain the full text of articles, but they turn to topic modeling to identify conspiracy theories, even though unsupervised methods

are not ideal for this task. In the remainder of this paper, we offer new theoretical arguments about the state supply of conspiracy theories and then test these arguments with better automated detection of conspiracy theories in newspaper text.

3 Threats and Regime Type Explain State Supply of Conspiracy Theories

Humans turn to conspiracy theories to make sense of the world, especially complex events. News organizations report on complex events as they unfold, so we expect that these events are likely to increase the supply of conspiracy theories in the news media even without government influence. Secretive or violent political events are especially likely to prompt conspiracy theories for at least two reasons. First, humans may have evolved to be interested in violent events because understanding them enhances individual and group survival in competitive environments. Second, contentious politics heighten anxiety and uncertainty. Journalists and readers might be more likely to entertain conspiracy theories during periods of violent unrest than they would during periods of relative calm. Thus, in the absence of state influence, we expect the media to supply more conspiracy theories when violent political events are more frequent for several reasons, including reporters' genuine beliefs and perceived reader demand.

Journalists at state-controlled media outlets may also respond to complex and threatening events by increasing the supply of conspiracy without any pressure from the state. Journalists working for state-controlled media are subject to many of the same forces as their counterparts at independent outlets. They may have latitude about what appears in print because state-controlled media tend to be run bureaucratically, with layers of principals and agents making decisions that affect the ultimate composition of the official media. Any conspiracy theories promoted by these journalists are technically being promoted by the state, by virtue of appearing in state-controlled media, but it is not safe to assume that all of them were intentionally placed there by the regime. While “naturally occurring” conspiracy theories that appear without state influence are important, they are not the main subject of our inquiry.

We seek to explain the conspiracy theories that are supplied by state-controlled media at the behest of the regime. We pose a counterfactual question that isolates the role of the state: how would the quantity and quality of conspiracy theories in the state-controlled media be different if it were not state-controlled? This question leads us to introduce two concepts that are central to our argument: the *oversupply* and *undersupply*

of conspiracy theories by state-controlled media. When the state-controlled media circulates more conspiracy theories than it would have if it were not state-controlled, we call this condition oversupply. Conversely, when the state-controlled media circulate fewer conspiracy theories than if it were not state-controlled, this is undersupply.

By introducing these two concepts, we reframe the problem of conspiracy theory supply. It is not sufficient to note that conspiracy theories become more prevalent in the state-controlled media during moments of crisis and conclude that this is all government strategy. Rather, the literature should recognize that complex events that precipitate crisis naturally lead to more conspiracy theories in the media. The key question is when and why the state puts its thumb on the scale in one direction or the other.

Previous scholarship tends to suggest that governments, especially autocratic ones, typically oversupply conspiracy theories because of their effectiveness at neutralizing dissent. We argue that the costs of promoting conspiracy theories are substantial and under-appreciated by previous theories. A more complete account of state supply of conspiracy theories needs to take these costs into account. If promoting conspiracy theories has no downside, then we would expect the state to promote conspiracy theories all the time, regardless of political developments. After all, if the primary effect of spreading conspiracy theories is to increase support for government policies among some citizens and neutralize dissent among others, then why wouldn't governments facing opposition just go back to the conspiracy theory well over and over again?

The upfront costs of promoting conspiracy theories are admittedly low. Conspiracy theories are easy to produce and easy to disseminate. They need not be logically coherent to be effective (Wood, Douglas and Sutton, 2012). The costs of promoting conspiracy theories come after they circulate. Conspiracy theories pollute the media commons and decrease human capital by cluttering the minds of media consumers (Wedeen, 2019). Although sowing confusion and paranoia among citizens can be a short-term ploy to retain power, it is not a recipe for long-term national success. Conspiracy theories can also have unintended consequences—conspiracies intended to distract citizens may instead create a Streisand effect that draws more attention to events the government would prefer to conceal (Hobbs and Roberts, 2018). An overreliance on conspiracy theories might make the government less able to communicate credibly about real threats and conspiracies. And outside of a state's borders, promoting conspiracy theories invites international criticism and ridicule because it violates norms and complicates diplomatic relationships, especially

with alleged foreign perpetrators of conspiracies (Alrababa'h and Blaydes, 2020).

We argue that the decision by a state to promote conspiracy theories weighs these benefits and costs. Threatening events are likely to raise the benefits and decrease the costs for many governments, both autocratic and democratic. Threatening events might lead the government to favor short-term gains over the long-term costs of promoting conspiracy theories. Additionally, the reputational costs of official conspiracism vary. As uncertainty grows in the wake of destabilizing events, official allegations that might have been implausible previously may become more credible, or at least less ridiculous, to both domestic and international audiences.

Considering the benefits and costs of promoting conspiracy theories also provides an explanation for why democratic governments are less likely to promote conspiracy theories than autocracies. It is not that the benefits of conspiracism are necessarily lower for elected officials. Rather, institutions and norms increase the costs of official promotion of conspiracy theories in a democracy. This is because political elites in democracies face incentives to satisfy relatively large portions of their populations to stay in office, which encourages them to provide good policies and public goods (Bueno de Mesquita et al., 2003). Ruling elites in autocracies are more likely to be able to stay in office by satisfying only a small portion of their population which encourages strategies of buying off supporters and repressing the opposition. Because democracies have incentives to provide public goods, they are less likely to destroy the media commons and decrease human capital. Political parties may have longer time horizons than individual politicians, so they don't want to ruin long-term reputations for short term gains (though individual politicians in democracies may angle for the short-term gains of promoting conspiracy theories, especially if they are trying to win over a narrow base of supporters). Elected officials and bureaucrats in democracies also face greater costs for destroying the media commons and promoting misinformation because freedom of the press is seen as a democratic norm internationally.

Even though these costs are higher for democracies, there may be times when a democratically-elected government promotes conspiracy theories because it expects the benefits to outweigh the costs. In the same spirit, we argue that not all autocratic governments are equally likely to promote conspiracy theories. Some autocrats base more of their rule on perceptions of good governance and public goods provision; these will face higher costs for promoting conspiracy theories. And some autocratic governments face more threats to

their rule than others. These threats, or at least the perception of threat, raise the value of conspiracy theories, especially during a time of crisis, when the government has a strong incentive to neutralize challenges to their continued political power.

Thus, regime type explains broad tendencies in conspiracy theory supply, but is not a complete explanation by itself. We argue that in addition to regime type, the degree of threat perceived by the government will impact the supply of conspiracy theories by the government.

Government perception of threat from political events operate at two timescales. In the short term, contentious political events threaten the government's ability to rule. In general, the types of political events that increase the threat a regime faces are oppositional or violent. Recent mass protests, terror attacks, strikes, etc, are likely to increase the level of threat perceived by ruling elites. This means that there is significant overlap in the types of events that are likely to inspire conspiracy theories by independent journalists and those that threaten a regime enough that it responds by increasing the supply of conspiracy theories in the official media.

Events also shape the threat perceptions of political leaders in the long term. When elites face relatively few political events that challenge their rule over long periods of time, they may conclude that threats to their continued rule are low. When elites face many threatening events over a long time period, or observe threats to prior leaders, they are more likely to assess their baseline level of threat to be high.

We can now summarize our argument for why political events and regime type interact to shape state supply of conspiracy theories and draw out observable implications of our theory. We expect that as threatening political events occur, conspiracy theories will increase in the news media generally. The presence of conspiracy theories in government media is not itself evidence that the government is promoting conspiracy theories. Conspiracy theories can appear for many reasons. Without inside information about the editorial process, journalists' beliefs, and the inner operations of a government, it is difficult to determine which of these explanations most plausibly accounts for the publication of any specific conspiracy theory. Instead, we seek to explain patterns of oversupply and undersupply of conspiracy theories in the state-controlled media. Even when we cannot detect government influence in any single instance, we are confident that when a state-controlled newspaper is printing more conspiracy theories on average than comparable independent newspapers, this is evidence of government influence.

We expect that democracies will generally undersupply conspiracy theories relative to autocracies facing the same levels of long- and short-term threat. This is not to say that democracies will not respond to events with conspiracy theories, but rather that they do so less because their costs are higher. For a regime attempting to retain office with good policies, promoting conspiracy theories is of limited use as a general strategy and has substantial risks. On the other hand, regimes whose ability to retain office are threatened might be less concerned about the negative effects of promoting conspiracy theories in the long term and thus benefit more from the short term benefits of spreading any single conspiracy theory. This implies the reverse relationship too: we expect that autocracies will generally oversupply conspiracy theories relative to democracies facing the same levels of threat.

We expect to find variation in the promotion of conspiracy theories by autocracies. We expect that governments that feel relatively secure in the long-term will tend to undersupply conspiracy theories. Governments that feel threatened in the long-term will tend to oversupply conspiracy theories. This oversupply will manifest in two ways. First, the baseline supply of conspiracy theories by the state will be higher. Two, the oversupply will manifest in more short-term responsiveness to threatening events. In this expectation, we align with Rozenas and Stukal (2019) who finds that deflection in Russia's state-controlled media "is used more intensely in politically sensitive times (elections and protests)" (982).

Regimes of all types will consider the costs of promoting conspiracy theories and reduce their conspiracy theory supply when these costs increase, all else staying equal. However, for governments that feel especially threatened in the long- and short-term, conspiracy theories may continue even when the cost to international cooperation is high, because increased threats make conspiracy theory promotion less costly in the short term.

While our theory predicts the quantity and timing of conspiracy theories, we have less to say about the content of these conspiracy theories. Previous studies suggest that governments might promote conspiracy theories with coherent content aimed to craft a narrative of events that promotes national cohesion Gray (2010a, 126). Alternatively, state-supplied conspiracy theories may be jumbles of incoherent content designed to divert attention from events that present the regime in an unflattering light (1999; 2020). Without a strong expectation about the content of the conspiracy theories that state-controlled media oversupplies, we will examine them deductively to see whether they follow a pattern of coherent narrative building or

incoherent distraction.

4 The Supply of Conspiracy Theories in Egypt, 1998-2018

We test these arguments by examining conspiracy theories in the largest state-controlled newspaper in Egypt between 1998 and 2018. We select Egypt for scientific and practical reasons. In order to test the effects of threat and regime type on state supply of conspiracy theories, we need variation in both key variables. During the time period we examine, Egypt has significant variation in both threat levels and regime type, including an unexpected and rapid period of democratization followed by an equally sudden return to autocracy. Egypt is also intrinsically important as the largest media market in the Arab world, and the home of some of the most important newspapers in the region. More practically, the digitization of key Egyptian newspapers facilitates our data collection and our familiarity with Egyptian politics enables us to analyze these materials qualitatively.

We measure the prevalence and content of conspiracy theories published by the main regime newspaper, *al-Ahram*, from 1998 to 2018. We focus on print media because it is stable over the twenty-year period we investigate and the lines of government influence are clear. Conspiracy theories also circulate on television and social media in contemporary Egypt, but rapid changes to these platforms make it impossible to make meaningful comparisons across the multiple regime transitions of Egypt's last twenty years.

We use these data to estimate the degree to which the state-controlled newspaper under- or oversupplies conspiracy theories as the Egyptian government changes regime types and experiences varying levels of threat. Our concepts of under- and oversupply require a counterfactual estimate of how many conspiracy theories the state-controlled newspaper would have published if it were independent. This cannot be measured directly, so we estimate this counterfactual quantity using conspiracy theories published in Egypt's main independent newspaper, *Al-Masry Al-Youm* since its founding in 2004. During the period we examine starting in 2005, these two newspapers produced 1,064,169 articles, which we analyze comprehensively. At a high level, our research strategy is to use a statistical text analysis model to detect conspiracy theories in each newspaper and then estimate models predicting the supply of conspiracy theories in each.

***Al-Ahram* and *Al-Masry Al-Youm* in the Egyptian Media Space**

Al-Ahram is Egypt's flagship newspaper. Founded in 1875, *Al-Ahram* was independent until 1960 when it was nationalized by Gamal Abdel Nasser. Its influence within Egypt was largely uncontested until the emergence of an independent newspaper, *Al-Masry Al-Youm* in 2004. We do not have reliable statistics on subscribers or circulation for either newspaper, but data from Google suggests that online searches for the newcomer exceeded searches for *Al-Ahram* by 2007. The relative ranking of these competitor papers has fluctuated since but both remain prominent: as of March 2021, Alexa.com rates *Al-Masry Al-Youm* as the 44th most popular website in Egypt while *Al-Ahram* is 31st.

Al-Ahram is accused by independent Egyptian journalists of promoting conspiracy theories at the behest of the government, especially in the wake of violent events that challenge the regime.³ Conspiracy theories in *Al-Ahram* are often attributed to “informed sources” and anonymous security sector officials, though these sources are viewed with skepticism.⁴

Al-Masry Al-Youm is the most comparable independent alternative to *Al-Ahram*. Its independence derives from its financial backing by Salah Diab, a powerful businessman with a background in oil (Al-Azm, 2015). Still, it has faced some constraints in authoritarian Egypt. Government approval was required to establish the newspaper and anecdotal evidence suggests that *Al-Masry Al-Youm* has occasionally come under government pressure (Peterson, 2011). Despite this, we find no evidence that *Al-Masry Al-Youm* would print conspiracy theories at the behest of the government. There are numerous accounts of direct intervention in *Al-Ahram*'s editorial process (Hammond, 2005); we are not aware of comparable examples in the case of *Al-Masry Al-Youm*. Additionally, regime pressure across both newspapers would bias our results against finding undersupply or oversupply, since the difference in that case would be zero; the effects we measure are likely an underestimate of the true difference.

Detecting Conspiracy Theories in Arabic Newspaper Text

We collect all of *Al-Ahram* since 1998 and all of *Al-Masry Al-Youm* since 2005 from the newspapers' websites. Our comparative analysis of the two papers focuses on the period between 2005 and 2018. We are missing data for *Al-Ahram* on 22 days because the issues were not archived on the Internet correctly, but this

³<https://dailynewsegypt.com/2016/10/02/al-ahram-blindly-parrots-governments-forces-evil-claim/>, Accessed 3/21/2021.

⁴<https://dailynewsegypt.com/2016/09/19/548755/>, Accessed 3/21/2021.

missingness is not concerning given the long timespan of our data. After removing duplicate articles (which appeared on the website for multiple days but only appeared in print once⁵), we analyze 826,765 articles from *Al-Ahram* and 484,100 from *Al-Masry Al-Youm*.

We attempt to comprehensively identify all of the conspiracy theories that appeared in both newspapers between 1998 and 2013. Reading over one million news articles is infeasible, so we turn to a statistical methods for classifying text (Grimmer and Stewart, 2013), adapted for Arabic. At a high level, our classification process is a simple application of supervised learning: we labeled a small number of texts by hand and then train an algorithm to mimic the hand-coding. This broadly succeeded, and our final classifier achieves 97.5 percent accuracy, but the full process was complicated. We include some details here and more in the appendix.

Upon inspecting the articles, we found that conspiracy theories are fairly rare. Randomly sampling articles to label by hand resulted in too few conspiracy theories for us to adequately train a classifier. We instead exploit the fact that writers in Egyptian newspapers almost always use some variant of the word “conspiracy” when introducing a conspiracy theory. We selected 18 key words that our initial reading revealed were frequently used to introduce conspiracy theories; these translate as “conspiracy,” “plot,” “machinations,” “collude,” “collusion,” and “intrigue” (and plural forms and spelling variations, where appropriate). We identified the 32,956 articles containing these terms and concentrated our hand-coding efforts there.

We sampled 1,500 articles containing our key words and hand-coded whether each paragraph in these articles contained a conspiracy theory or not. We settled on identifying conspiracy theories at the paragraph-level after substantial trial and error. Sentences proved too short and entire articles were often too wide-ranging in their scope. We developed coding rules (see the appendix) and then hired two native Arabic-speakers from North Africa to code each of the 1,500 articles according to our rules. They coded which paragraphs in each article contained conspiracy theories and who the perpetrators and victims in these conspiracy theories were. Our coders largely agreed when identifying conspiracy theory paragraphs (95.9% agreement) and we adjudicated discrepancies.

We use this hand-coded data to train a statistical classifier to identify conspiracy theories in the remaining articles. We settled on a simple workhorse model from machine learning, the random forest classifier

⁵There were 25,178 articles in *Al-Ahram* that were duplicated on multiple days, and only two for *Al-Masry Al-Youm*. We spot checked the digital archives against print editions to confirm that these appeared in the print edition only on the first day.

(Breiman, 2001), because it performed well and did not overburden our computational resources.

To train the classifier, we preprocessed the hand-labeled paragraphs using standard tools for cleaning and stemming Arabic text. We used the 1,990 most frequently occurring terms as the predictors in our classifier. We also include indicator variables for the broad type of article: news, opinion, culture, or sports. We then trained our random forest model using k-fold cross-validation, a process in which parts of the data are iteratively held out, model tuning parameters are optimized to the remaining data, and then the accuracy of the resulting models is estimated by predicting the held out data. We also held out 10 percent of the training data that was never used in any of the classifier training, so we have a true estimate of the accuracy of our model. In this held-out 10 percent, we correctly identify conspiracy theory paragraphs 78 percent of the time, and correctly identify non-conspiracy theory paragraphs 99 percent of the time, for an overall accuracy of 97.5 percent.

We apply the random forest model to classify conspiracy theories in the remaining 449,297 paragraphs in articles that contain our key words, resulting in a total of 30,473 conspiracy theory paragraphs across the two newspapers. This is our preferred measure of conspiracy theories and is what we report in most figures and models, but our results are robust to several alternative measures that we mention below and describe in the appendix.

What is in these conspiracy theory paragraphs? Summarizing 30,000 paragraphs, some of which contain truly far-fetched accounts, is challenging. We describe some of these conspiracy theories qualitatively below. For now, we consider one important feature of each conspiracy theory: who is claimed to have conspired against whom? Accurately identifying the alleged perpetrators and victims in all 30,000 paragraphs is not feasible with either close reading (because of time constraints) or automated text analysis (substantial variety in the language describing perpetrators and victims requires human judgment). Instead we focus on the training set, which is representative of all 30,000 paragraphs but is small enough to read. As our coders classified articles, we had them record the perpetrator(s) and victim(s) in each conspiracy theory in the original language of each article. We then combed through this list ourselves to identify the various ways each actor is described and consolidate a list of the perpetrators and victims in the conspiracy theories in the training set. Not every coded text identified perpetrator(s) or victim(s), but this analysis allows us to characterize those that do.

Figure 1 plots each of the entities that appears in the training set conspiracy theories more than five times. Each entity is plotted (in red text) on the x-axis according to the proportion of mentions in which that entity is described as a perpetrator of a conspiracy by *Al-Ahram*, with the y-axis indicating the proportion of *Al-Ahram* conspiracy theories that mention that entity (also indicated by font size). The blue arrows point from the *Al-Ahram* values to the *Al-Masry Al-Youm* values, showing the difference between them.




Figure 1: Entities plotted in red text by how often they are perpetrators (x-axis) and how often they appear (y-axis) in the training set articles from *Al-Ahram*. Blue arrows indicate how often the same entities appear and how often they are perpetrators in *Al-Masry Al-Youm*. The length and direction of the arrows show the difference between the newspapers. See the Appendix for details about these entities.

Figure 1 shows that the entity Egypt (which includes the “Egyptian people”) is the most frequently mentioned entity in both newspapers and is virtually always a victim. Egyptian government entities, which we tabulate separately, are mentioned almost as frequently, but are evenly split between mentions as a victim and perpetrator. America, Israel, and the Muslim Brotherhood are frequent perpetrators in *Al-Ahram*, with “terror,” the West, Britain, Qatar, and France, appearing somewhat less frequently but always as perpetrators too. Roughly 18 percent of conspiracy theories referenced “shadowy forces” or other vague, unspecified

perpetrators, which we combine as “nebulous” perpetrators.

The blue arrows reveal that conspiracy theories in *Al-Masry Al-Youm* emphasize some of these entities differently. *Al-Masry Al-Youm* mentions the Egyptian government substantially more often, and mentions America, Israel, the West, terrorists, Britain, and France at about half the rates of *Al-Ahram*. This gives a sense already that *Al-Ahram* is interested in de-emphasizing the role of the Egyptian Government as a perpetrator of conspiracies, and emphasizing conspiracies purportedly perpetrated by the West and by Islamists.

To get a sense of which major types of conspiracy theories appear in these newspapers, we used k-means clustering to cluster the conspiracy theories in the training set based on similar combinations of perpetrators and victims. This method is exploratory, so we examined many possible clustering solutions — more than we can fully report here — to get a sense of which types of conspiracy theories are most prevalent. Many of the conspiracy theories cluster together according to which main actor is purportedly attacking the Egyptian people: the Egyptian Government (13%), America (11%), nebulous forces (8%), Israel (7%), the Muslim Brotherhood (6%), terrorists (5%), and a range of other actors in unusual combinations (15%). However, roughly a quarter of the conspiracy theories do not fit this pattern; instead they are a hodgepodge of assorted perpetrators and victims that appear infrequently. This, combined with the very large number of entities that we have already lumped into a miscellaneous category in our coding, suggests that there is a wide range of conspiracy theories in the newspaper beyond those in which Egypt is the victim. We suspect that these appear for different reasons, and that when the government newspaper oversupplies conspiracy theories, they are disproportionately the conspiracy theories alleging plots against Egypt.

The Content of Conspiracy Theories

The daily count of conspiracy theory paragraphs printed in each newspaper is our primary outcome variable in our quantitative analysis. To show how these counts have fluctuated over time, we plot the moving average for each newspaper in red (*Al-Ahram*) and blue (*Al-Masry Al-Youm*) in Figure 2.

To contextualize these trends, we also plot two variables in the background that are important predictors of conspiracy theories in our models below. First, we indicate who the executive of the Egyptian state is at any given time using alternating gray and white background shading. Second, we plot the moving average of violent events in Egypt on each day as reported in the ACLED data set (Raleigh et al., 2010).

In combination, these measures operationalize the key theoretical variables that we expect to influence

government supply of conspiracy theories. The literature to which we contribute argues that regime type is important, which often gets bluntly coded in quantitative analysis as democracy versus autocracy. If we applied such a blunt coding to Egypt over this time period, we might mark the period of Muhammad Morsi's rule (2012-2013) as democracy and the other periods as autocracy. However, focusing on Egypt allows us to consider that the political sphere in Egypt may have had a different feel under each executive: autocracy under al-Sisi is widely seen as more repressive than under Mubarak's last years, or under the Supreme Council of the Armed Forces (SCAF) which seized power after the 2011 revolution. Thus we mark political eras in Egypt during our period of observation based on the dates of executive leadership change.

Our argument is that autocratic governments do not uniformly supply conspiracy theories at all times; they over-supply conspiracy theories when they feel threatened and under-supply them when feeling secure. We cannot directly measure the private perceptions of threat or security by Egyptian executives over the last 15 years. Public statements projecting security or insecurity may be insincere, and direct measures of private feelings are not possible, especially retrospectively. Instead, we proxy for the threat faced by the Egyptian government using the counts of events that the government is likely to find threatening: attacks, battles, and protests. The ACLED data set is widely cited as one of the best available sources for these sorts of event counts, so we present it here, but we employ various other measures in our analysis below to be confident that our results are not dependent on this particular data set. During the 20-year period from 1998 to 2018, ACLED reports 3,883 protests, 1,811 riots, 1,909 battles, 1,718 explosions or remote violence, 1,058 instances of violence against civilians, and 661 "strategic developments." The ACLED data may undercount smaller events, such as localized protests, but the major events most likely to threaten the government appear to be accurately counted.

From Figure 2, we can see several correlations that are central to the regression results we report below. First, conspiracy theories in both newspapers are generally responsive to events; when there are more protests, attacks, and battles happening in Egypt, there are more conspiracy theories in the newspaper. Second, *Al-Ahram* generally supplies fewer conspiracy theories than *Al-Masry Al-Youm* prior to 2011, and generally supplies more afterward, with the exception of the year when the Muslim Brotherhood is in power. Third, relatively few violent political events happen during Mubarak's rule prior to the end of 2010, reinforcing our inference that the Mubarak government felt relatively secure most of the time. The exception




Figure 2: Average counts of conspiracy theories in each newspaper over time. The solid lines indicate the moving average. Dashed lines indicate the average during each administration.

is at the beginning of 2011 as the revolution kicks off, which was accompanied by a spike in conspiracy theories in *Al-Ahram*. By contrast, the governments led by Mansour (with al-Sisi in the background) and then by al-Sisi directly have experienced levels of violent political events that, while declining since a high in 2013, remain much higher than the status quo prior to 2011. We interpret this graph to be consistent with the claim that al-Sisi government feels threatened to a degree that Mubarak government did not.

Regression Approach and Results

To test whether oversupply of conspiracy theories is a result of threatened autocratic government in Egypt, we turn to a regression analysis of these same data. Our outcome variable in these regressions remains the count of conspiracy theory paragraphs in each newspaper on each day. Because our outcome variable is measured in counts, we use a generalized linear model with a negative binomial link that is appropriate for modeling (potentially over-dispersed) counts.

To test whether threat leads to oversupply, we proxy threat using the sum of all violent political events in ACLED over the previous 7 days. We interact this variable with an indicator variable for each newspaper, expecting that *Al-Ahram* will be more responsive to increases in ACLED event counts than *Al-Masry Al-*

Youm.

We include several control variables that might themselves explain variation in conspiracy theories or confound the relationship between the events-newspaper interaction and conspiracy theories. Both newspapers have gotten longer over time, so we control for the number of paragraphs and articles in each issue of the newspaper. A fair number of articles in both newspapers are drawn from international news agencies and wire services, such as Reuters and the Associated Press. We suspect that the supply of conspiracy theories is different in these articles than those that are written entirely by local journalists, so we control for the count of paragraphs and articles for which one or more of these news agencies are listed as contributors in the byline. Finally, we control for the day of the week, expecting that social norms and journalistic practices that vary throughout the week will affect the supply of conspiracy theories.

In some models, we add year-month fixed effects, which are indicator variables for each calendar month between 2005 and 2018. These year-month fixed effects control for any variation in conspiracy theory supply that lasts for an entire calendar month. Any model with these fixed effects is likely to understate the full correlation between events and conspiracy theories, but including them gives us confidence that the correlations we observe between events and conspiracy theories are due to daily variation, rather than simply longer-term trends.

We present our main regression results graphically in Figure 3, with a full table in the Appendix. The coefficients of our negative binomial generalized linear models show the direction and statistical significance of each of our covariates, but the substantive magnitude is not directly interpretable, so we use the technique from King, Tomz and Wittenberg (2000) to simulate the change in conspiracy theories associated with changes in our key independent variables.

We find that *Al-Ahram* oversupplies conspiracy theories relative to *Al-Masry Al-Youm* in response to the same events. The two panels of Figure 3 show predicted counts of conspiracy theories in each newspaper as the count of ACLED events in the last 7 days increases from zero to the maximum of 151. The histogram of ACLED events appears on the x-axis of each panel. The left panel, corresponding with Model 1 in Table 1 of the Appendix, shows that when the count of events is low, the state-controlled newspaper supplies detectably fewer conspiracy theory paragraphs on average: 2.2 per day for *Al-Ahram* and 2.7 per day for *Al-Masry Al-Youm* with zero events in the last 7 days. The predicted values cross when ACLED events reach




Figure 3: *Al-Ahram oversupplies conspiracy theories relative to Al-Masry Al-Youm in response to the same events.*

28, and beyond that, *Al-Ahram* prints substantially more conspiracy theories than *Al-Masry Al-Youm*. Near the maximum, when ACLED records 150 events in the last week, *Al-Ahram* prints 17.9 conspiracy theory paragraphs per day, 2.3 times as many as *Al-Masry Al-Youm* following the same events.

The right panel shows that a similar relationship, smaller in magnitude, exists when we account for the baseline level of conspiracy theories in each calendar month by adding year-month fixed effects. In this model, the predicted number of conspiracy theories in each newspaper with zero events in the the last seven days are statistically indistinguishable from each other. In the month we have chosen to simulate effects — July 2013 — the baseline rate of conspiracy theory paragraphs is about 4.4, but this baseline changes with each month. As the count of events increases, the more conspiracy theories appear in *Al-Ahram*, rising to 7.8 when there have been 150 events in the last 7 days. In contrast, the predicted number of conspiracy theories in *Al-Masry Al-Youm* stays constant, or even declines slightly, as the number of events increases.

These two models predicting the count of conspiracy theories, with and without year-month fixed effects, provide our main quantitative conclusions. First, both newspapers are responsive to events, but differentially so. *Al-Ahram* is highly responsive to events when we compare across months and within each month. When the count of events changes from its minimum to its maximum, *Al-Ahram* increases its supply of conspiracy theories approximately sevenfold, from 2.5 to 18. This positive relationship holds even when we allow each

month to have its own individual baseline rate of conspiracy theories, which means that *Al-Ahram* responds to increases in events by supplying more conspiracy theories regardless within the time-frame of a month. This result is not simply because some broad periods of Egyptian politics lend themselves to conspiracy theories more than others; this shows day-to-day responsiveness by the state-controlled newspaper. *Al-Masry Al-Youm* is far less responsive. When events increase from zero to 150, it does roughly triple its supply of conspiracy theories from 2.7 to 7.7 per day. However, when we condition on year-month fixed effects, there is no evidence that *Al-Masry Al-Youm* responds to daily changes in events by increasing conspiracy theories.

Our results are robust to a very large number of other modeling and measurement alternatives: Poisson regression; linear regression; aggregating conspiracy theories at the article level; counting only conspiracy theories that are endorsed by the journalist; including the data from *Al-Ahram* back to 1998; estimating separate regressions for each type of the ACLED event individually; using the count of ACLED events on the same day only; using the one-day lag of ACLED events; using the logged ACLED events in the last 7 days to account for skewness; using the logged count of ACLED events on the same day only; proxying threat with ACLED death counts instead of event counts; using a dichotomous variable “ACLED crisis” defined as 1 if ACLED events are in the top 75th percentile and 0 otherwise; a dichotomous variable “ACLED crisis 2” defined as 1 if ACLED events are in the top 90th percentile and 0 otherwise; proxying threat with the Global Terrorism Database terrorism event counts or death counts (LaFree and Dugan, 2007); omitting combinations of control variables; adding indicator variables for each regime; and a few more. For brevity, the details are in the Appendix.

Several additional analyses deserve more attention because they are revealing about the supply of conspiracy theories. First, we find that the moving average of ACLED events over the last year is a robust predictor of conspiracy theory oversupply by *Al-Ahram*, whether we include or exclude the moving average of ACLED events over the last 7 days.⁶ This result supports our argument that long-term threat affects the supply of conspiracy theories in addition to short-term threat.

After completing the analysis of *Al-Ahram* and *Al-Masry Al-Youm*, we became concerned that our findings might be dependent on idiosyncrasies of the two newspapers, or that *Al-Masry Al-Youm* might not be

⁶When we include both the last week and last year of ACLED events in the same model, we calculate the moving average of the last year excluding the last week to avoid any mechanically-induced collinearity.

as independent as we believe. With advice from Egyptian colleagues,⁷ we collected 775,126 articles (from 3,230 days) from the independent newspaper *Al-Shuruq*, and 75,196 articles (from 878 days) from the next-most-prominent state-controlled newspaper, *Al-Gomhuria* (see the Appendix for details). We were unable to collect either of these newspapers completely so we do not include them in our main analysis, but our results remain substantively similar when we add the available data.

Next, we divide the data by periods, estimating the same specification in model 1 but for only the dates corresponding to the tenure of each executive between 2005 and 2018. We plot the key predictions in Figure 4. If each of the five governments we examine — Mubarak, the SCAF, Morsi, Mansour, and al-Sisi — use the same conspiracy theory strategy, then we expect each of the five panels in Figure 4 to look similar to the top left of Figure 3. Instead, we find that *Al-Ahram* only begins oversupplying conspiracy theories relative to *Al-Masry Al-Youm* in 2013 when al-Sisi and the military take over in a coup and installed Mansour as president. Under both the dictatorship of Mubarak and the elected government of Morsi, both papers supply more conspiracy theories when ACLED records more events, but there is no detectable difference in the slopes of these curves. Under the interim caretaker government SCAF, there is again no difference in the slopes of the curves, but both trend downwards, meaning that as more events happen, both newspapers actually decrease the supply of conspiracy theories. It is after the coup that the trends emerge that drive our results above. Since 2013, *Al-Ahram* supplies more conspiracy theories in response to violent events while *Al-Masry Al-Youm* does not. The predicted differences are smaller than those in model 1 because the Mansour and al-Sisi eras have a substantially higher baseline of conspiracy theories than the Mubarak era. Still, the oversupply in the post-coup periods is statistically significant and substantively important.

To better visualize exactly when oversupply of conspiracy theories by *Al-Ahram* is happening, we generate estimates of oversupply from model 1. We calculate the daily under- or oversupply in *Al-Ahram* by calculating predicted values from the model using the covariate values for *Al-Ahram* on a given day but setting the *Al-Ahram* indicator variable to zero, as if it were the independent newspaper. We then subtract the observed value of conspiracy theories in *Al-Ahram* that day from the predicted value if *Al-Ahram* were independent and use this as our daily estimate.

We plot these values in Figure 5, with the x-axis indicating the date and the y-axis indicating the esti-

⁷Many thanks to the students and faculty at Cairo University who recommended we analyze *Al-Shuruq*.




Figure 4: *Al-Ahram oversupplies conspiracy theories when Mansour and al-Sisi are in office, but not when Mubarak, the SCAF, and Morsi are in office.*

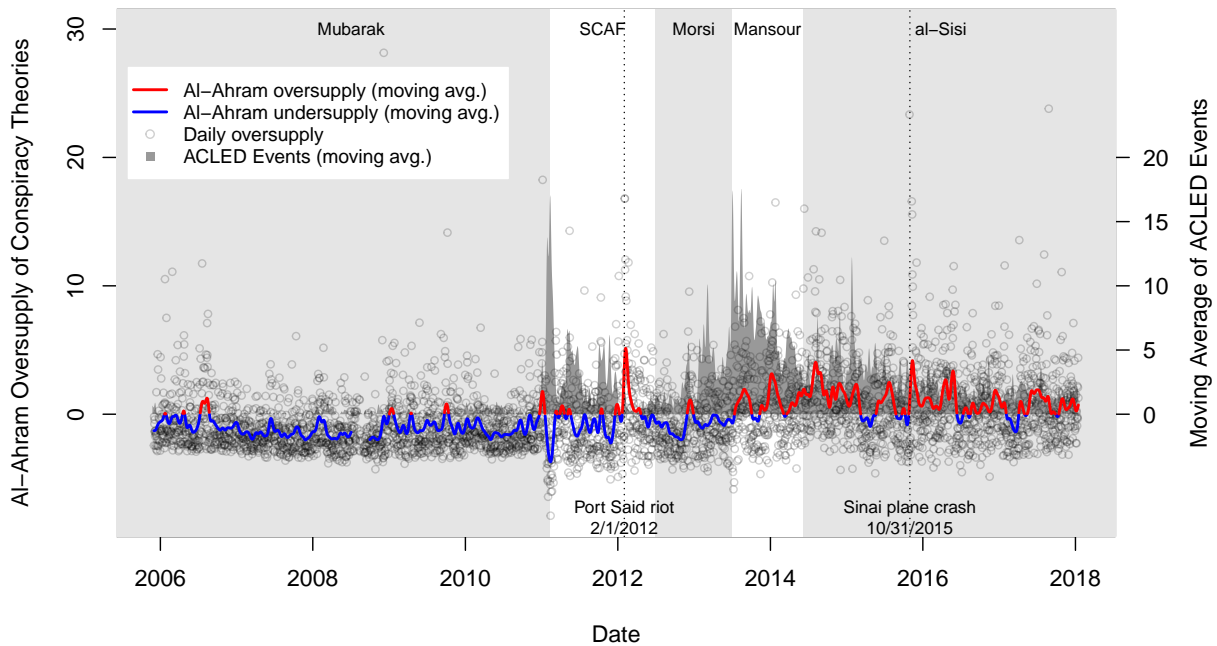


Figure 5: *Oversupply of conspiracy theories in Al-Ahram, based on counterfactual estimates from a negative binomial regression model of how many conspiracy theory paragraphs would have appeared if Al-Ahram had been Al-Masry Al-Youm on each day.*

mated undersupply (negative values) or oversupply (positive values) by *Al-Ahram*. As before, we indicate the regime changes with alternating white and gray regions, and plot the moving average of ACLED events in the background. It is hard to see trends in the daily oversupply estimates because there is substantial variability. To clarify the trends, we calculate a moving average of the under/oversupply estimates, using

blue for periods of undersupply and red for periods of oversupply.

We find that under Mubarak, the state-controlled newspaper undersupplied conspiracy theories nearly the whole time. One interesting exception is that *Al-Ahram* begins to oversupply conspiracy theories just at the onset of the 2011 revolution, but then flips back to the most extreme undersupply that we observe in the entire time period right after Mubarak is deposed. We take this oversupply to be Mubarak's last-ditch effort to divert attention as his grip on power weakened, and the subsequent undersupply to be the SCAF saying "nothing to see here, folks." Our conclusion is that in the absence of threat the reputational costs of printing conspiracy theories in the official paper far exceeded any short term gains for the vast majority of the Mubarak government's rule in the period observed in our dataset.

Undersupply remains the norm while the SCAF is in power, except for an extreme period of oversupply that followed the Port Said soccer riot, an extreme embarrassment for the regime.⁸ Undersupply continues after Morsi takes office. This year is the only period of democracy that we observe, so this provides some evidence that democracy decreases government supply of conspiracy theories. If we code Egypt as a democracy during the Morsi era and autocratic during the others, then the indicator variable for democracy is statistically and substantively correlated with fewer conspiracy theories in *Al-Ahram* when we add that variable to model 1 and calculate a triple interaction of ACLED events \times *Al-Ahram* \times democracy. However, Figure 5 makes it clear that this is because *Al-Ahram* oversupplies during the autocracy of Mansour and al-Sisi, but not Mubarak or the SCAF. We do not observe enough periods of democracy to say for sure whether it inhibits government supply of conspiracy theories, but we can conclude that autocracy alone is not a sufficient condition for oversupply.

The oversupply of conspiracy theories after the 2013 coup is immediate, substantial, and persistent. There have only been a couple of months during the last 5 years of our data when *Al-Ahram* is not oversupplying conspiracy theories, and they are immediately followed by large upticks, most notably following the Sinai plane crash in which a Russian airliner was downed by terrorists and the Egyptian government scrambled to deflect criticism over the security breach.

The changes in oversupply that we observe mirror other trends in the Egyptian media space. With the resignation of Mubarak and the first free and fair parliamentary and presidential elections since indepen-

⁸The Appendix includes a qualitative assessment of this incident.

dence in 1922, Egyptian journalists enjoyed two remarkable years of relative openness (Peterson, 2011), before the authoritarian government of al-Sisi clamped down again after the coup. This aligns with our conclusion that neither threat nor authoritarianism alone is a sufficient explanation; it is authoritarians feeling threatened who are most likely to print conspiracy theories.

Qualitative Analysis of Oversupplied Conspiracy Theories

To learn more about the Egyptian government's goals, we qualitatively examine some of the conspiracy theories that the state-controlled media prints during periods of oversupply. Our goal is to inductively learn what types of conspiracy theories the state-controlled newspaper promotes on days when oversupply is high, so we sample qualitative cases using our quantitative estimates of conspiracy theory oversupply, shown in Figure 5. We identify the days on which predicted oversupply by *Al-Ahram* is above the 80th percentile and select a random sample of thirty days from that set of 881. This sampling strategy, in which our cases depend intimately on our regression model, leverages the strengths of nested analysis (Lieberman, 2005) and the integrative mixed-methods tradition (Seawright, 2016). With these thirty randomly selected days, two of us evaluated every *Al-Ahram* article with a conspiracy theory, summarizing the nature of the alleged conspiracy and noting whether any of the incidents or actors matched those involved in events from the ACLED data in the previous two weeks.

In the 141 articles we examined, references to perpetrators, victims, and incidents from the prior two weeks of ACLED were present only 13% of the time for the perpetrator, 6% of the time for the victim, and 9% of the time for the incident. The vast majority of *Al-Ahram's* conspiracy theories on these days have little connection to recent events, suggesting that the state intends to deflect rather than craft a narrative about what is happening.

When state-controlled media does reference recent events in conspiracy theories, it is in response to exceptionally pivotal events. October 12, 2011 was in our sample, which was preceded by the "Maspero Massacre," in which several dozen peaceful (mostly Coptic) protesters were killed and hundreds wounded when security forces attacked them in front of the Maspero television building. Four of the seven articles with conspiracy theories referred to the incident, but declined to provide a detailed account or name a perpetrator responsible for the attack. Instead, the articles vaguely allege the presence of evil forces threatening the Egyptian people. For example, one article quotes a religious figure who "confirmed the existence of

collusion by domestic and foreign elements, that aims to push Egypt toward a state of anarchy” and “called on all Egyptian people, both Muslim and Christian, to exercise self restraint, stay calm, and not participate in any demonstrations or sit-ins.”

At least half of the conspiracy theories in *Al-Ahram* we randomly selected for close reading appear formulated to promote unity by invoking vague fears. The formula consists of (1) a reference to (usually undefined) domestic or foreign perpetrators, (2) an undefined plot to divide the Egyptian people by sowing chaos, and (3) a call for unity, calm, and fortitude. *Al-Ahram* casts the Egyptian government as the only protection from these plots, portraying leaders as stalwart in the face of threats others could not face. They also portray these leaders as astute, being wise to threats that others are too naive to see. And, lastly, they rail against dissenters, often explicitly. For instance, an article from February 24, 2016, two years after the military coup that brought Egypt’s current president, Abdel Fatah al-Sisi to power, notes that “those who see al-Sisi’s words merely as a method to stay in power or to strengthen the grip of the security apparatus—respond with disregard and contempt to his talk of conspiracies and plots, when in actuality the facts of the matter demonstrate that 99% of what happens on the ground in Egypt is not a coincidence at all, and those who deny the existence of a conspiracy against Egypt and the region are either naive, ignorant, or themselves a participant in the conspiracy, perhaps without even being aware.”

5 Conclusion

When and why do states promote conspiracy theories? In this paper, we have offered new concepts, theory, and evidence to answer this question.

We introduce the concepts of oversupply and undersupply of conspiracy theories by states. Conspiracy theory oversupply happens when state-controlled media spreads more conspiracy theories than if it were independent; undersupply happens when state-controlled media spreads fewer. Several important studies have examined changes in the number of conspiracy theories in state-controlled media over time, but none, to our knowledge, have considered counterfactually what state-controlled media would produce if it were not state-controlled. If we are correct that independent media also promote conspiracy theories, than it is not merely the appearance of conspiracy theories in state-controlled media that needs explanation. Rather, our goal is to explain why state-controlled media sometimes promote more or fewer conspiracy theories than

we would expect if they were independent.

To explain the over- and undersupply of conspiracy theories by states, we take two theoretical factors that are considered separately in the literature—regime type and diversionary incentives—and bring them together in a unified theory. We agree with previous studies that states face diversionary incentives to promote conspiracy theories when they are under threat, but these incentives are offset by the risk that doing so will backfire. With these costs in view, we argue that states in which regime elites do not feel particularly threatened will undersupply conspiracy theories in state-controlled media. Once regime elites begin to feel threatened, they will switch to oversupplying conspiracy theories. Our argument applies to democracies as well as autocracies—regime elites in democracies can feel threatened as well—but autocrats face fewer incentives to preserve the integrity of the public sphere when their survival is threatened.

Most empirical studies of state conspiracy theory promotion focus solely on the dynamics of autocratic politics, often under a single stable regime. We test our argument in Egypt where rapid shifts in regime type allow us to test whether political institutions matter while holding the media market constant. We measure the supply of conspiracy theories in millions of Egyptian newspaper articles over twenty years, which allows us to make comparisons across time and between newspapers that were not previously possible. We find that conspiracy theories appear in both state-controlled and independent media on a daily basis, but typically in low numbers. In line with our theory, we find that state-controlled media begins to oversupply conspiracy theories as events threaten the regime; most of this oversupply has happened under the most recent Egyptian government which clearly feels substantial threat. Our qualitative examination of these conspiracy theories reveals that they typically encourage societal solidarity and stability in the face of vague threats rather than crafting a narrative about specific recent events.

We find suggestive evidence that democracy matters: the elected Muslim brotherhood government faced substantial threat but undersupplied conspiracy theories in the state-controlled newspaper nonetheless. However, many in Egypt might dispute that the Muslim Brotherhood government was genuinely democratic, and a single-country analysis cannot fully test our argument.

Our theoretical and conceptual contributions suggest several fruitful avenues for future exploration. For scholars of conspiracy theory supply, our approach could be extended to study the supply of conspiracy theories by states across the Arab World and beyond. More studies could compare state-controlled and

independent media in places where both operate. In places like the United States where media independence is strong, our methods suggest new ways to analyze the supply of conspiracy theories by non-state actors across the political spectrum.

How can scholars of conspiracy theory demand build on our study? To date, most experimental studies do not consider variation in the supply of conspiracy theories, and those that note its importance mostly control it away (Ryan and Aziz, 2020). We can imagine a research agenda that more closely mimics the wide variation in supply in the real world and experimentally varies it to learn how conspiracy theory supply affects demand.

Beyond questions of conspiracy theories, the study of Arabic-language news media in political science is in its infancy (Alrababa'h, 2019), and exciting new developments are in the works that show the promise of a text-as-data approach (Alrababa'h, 2021). We encourage scholars of political communication to consider what we can learn about politics in the Arab world with the new data and tools our study provides.

References

- Al-Azm, Sadik J. 2015. Broadcasting and Businessmen in Egypt: Revolution is Business. In *Arab Media Moguls*, ed. Donatelle Della Ratta, Naomi Sakr and Jakob Skovgaard-Petersen. London: I.B. Tauris chapter 1.
- Arababa'h, Ala'. 2019. "Quantitative Text Analysis of Arabic News Media." *MENA Politics Newsletter* 2(2):37–41.
- Arababa'h, Ala'. 2021. Manufacturing threats: Diversionary discourse and autocratic survival PhD thesis Stanford University.
- Arababa'h, Ala and Lisa Blaydes. 2020. "Authoritarian media and diversionary threats: lessons from 30 years of Syrian state discourse." *Political Science Research and Methods* p. 116.
- Berinsky, Adam J. 2017. "Rumors and health care reform: Experiments in political misinformation." *British journal of political science* 47(2):241–262.
- Black, Jeffrey. 2008. "Egypt's press: More free, still fettered." *Arab Media & Society* 4:1–13.
- Breiman, Leo. 2001. "Random forests." *Machine learning* 45(1):5–32.
- Bruder, Martin, Peter Haffke, Nick Neave, Nina Nouripanah and Roland Imhoff. 2013. "Measuring individual differences in generic beliefs in conspiracy theories across cultures: Conspiracy mentality questionnaire." *Frontiers in Psychology* 4(April):225.
- Bueno de Mesquita, Bruce, Alastair Smith, James D. Morrow and Randolph M. Siverson. 2003. *The Logic of Political Survival*. Cambridge, MA: MIT Press.
- Cassese, Erin C, Christina E Farhart and Joanne M Miller. 2020. "Gender differences in COVID-19 conspiracy theory beliefs." *Politics & Gender* 16(4):1009–1018.
- Dixon, Ruth M. and Jonathan A. Jones. 2015. "Conspiracist ideation as a predictor of climate-science rejection: An alternative analysis." *Psychological Science* .
- Einstein, Katherine Levine and David M. Glick. 2014. "Do I think BLS data are BS? The consequences of conspiracy theories." *Political Behavior* 37(3):679–701.
- Enders, Adam M, Steven M Smallpage and Robert N Lupton. 2020. "Are all birthers conspiracy theorists? On the relationship between conspiratorial thinking and political orientations." *British Journal of Political Science* 50(3):849–866.
- Goertzel, Ted. 1994. "Belief in conspiracy theories." *Political Psychology* 15(4):731–742.
- Gray, Matthew. 2010a. *Conspiracy theories in the Arab world: Sources and politics*. New York: Routledge.
- Gray, Matthew. 2010b. "Revisiting Saddam Hussein's political language: The sources and roles of conspiracy theories." *Arab Studies Quarterly* 32(1):28–46.
- Grimmer, Justin and Brandon M Stewart. 2013. "Text as data: The promise and pitfalls of automatic content analysis methods for political texts." *Political analysis* 21(3):267–297.

- Hammond, Andrew. 2005. *Pop culture arab world!: Media, arts, and lifestyle*. Santa Barbara, CA: ABC-CLIO.
- Hobbs, William R and Margaret E Roberts. 2018. "How sudden censorship can increase access to information." *American Political Science Review* 112(3):621–636.
- Hofstadter, Richard. 1965. *The Paranoid Style in American Politics*. Vintage.
- Huang, Haifeng. 2015. "Propaganda as signaling." *Comparative Politics* 47(4):419–444.
- Huang, Haifeng. 2018. "The pathology of hard propaganda." *The Journal of Politics* 80(3):1034–1038.
- King, Gary, Michael Tomz and Jason Wittenberg. 2000. "Making the most of statistical analyses: Improving interpretation and presentation." *American journal of political science* pp. 347–361.
- LaFree, Gary and Laura Dugan. 2007. "Introducing the global terrorism database." *Terrorism and political violence* 19(2):181–204.
- Lazer, David MJ, Matthew A Baum, Yochai Benkler, Adam J Berinsky, Kelly M Greenhill, Filippo Menczer, Miriam J Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild et al. 2018. "The science of fake news." *Science* 359(6380):1094–1096.
- Lieberman, Evan S. 2005. "Nested Analysis as a Mixed-method Strategy for Comparative Research." *American Political Science Review* pp. 435–452.
- Lucas, Christopher, Richard A Nielsen, Margaret E Roberts, Brandon M Stewart, Alex Storer and Dustin Tingley. 2015. "Computer-assisted text analysis for comparative politics." *Political Analysis* 23(2):254–277.
- Nyhan, Brendan and Thomas Zeitzoff. 2018. "Conspiracy and Misperception Belief in the Middle East and North Africa." *The Journal of Politics* 80(4):1400–1404.
- Oliver, J. Eric and Thomas J. Wood. 2014. "Conspiracy theories and the paranoid style(s) of mass opinion." *American Journal of Political Science* 58(4):952–966.
- Peisakhin, Leonid and Arturas Rozenas. 2018. "Electoral effects of biased media: Russian television in Ukraine." *American journal of political science* 62(3):535–550.
- Peterson, Mark Allen. 2011. "Egypt's media ecology in a time of revolution." *Arab Media & Society* 14:1–16.
- Popper, Karl R. 1945. *The open society and its opponents*. London: Routledge.
- Radnitz, Scott. 2016. "Paranoia with a purpose: conspiracy theory and political coalitions in Kyrgyzstan." *Post-Soviet Affairs* 32(5):474–489.
- Radnitz, Scott and Patrick Underwood. 2017. "Is belief in conspiracy theories pathological? A survey experiment on the cognitive roots of extreme suspicion." *British Journal of Political Science* 47(1):113.
- Raleigh, Clionadh, Andrew Linke, Håvard Hegre and Joakim Karlsen. 2010. "Introducing ACLED: an armed conflict location and event dataset: special data feature." *Journal of peace research* 47(5):651–660.

- Rozenas, Arturas and Denis Stukal. 2019. "How autocrats manipulate economic news: Evidence from Russias state-controlled television." *The Journal of Politics* 81(3):982–996.
- Ryan, Timothy and Amanda Aziz. 2020. "Is the Political Right More Credulous?: Experimental Evidence Against Asymmetric Motivations to Believe False Political Information."
- Seawright, Jason. 2016. *Multi-method Social Science: Combining Qualitative and Quantitative Tools*. Cambridge University Press.
- van Prooijen, Jan-Willem, Karen M Douglas and Clara De Inocencio. 2018. "Connecting the dots: Illusory pattern perception predicts belief in conspiracies and the supernatural." *European journal of social psychology* 48(3):320–335.
- Wedeen, Lisa. 1999. *Ambiguities of domination: Politics, rhetoric, and symbols in contemporary Syria*. Chicago: University of Chicago Press.
- Wedeen, Lisa. 2019. *Authoritarian Apprehensions: Ideology, Judgment, and Mourning in Syria*. University of Chicago Press.
- Williamson, Scott and Mashail Malik. 2020. "Contesting Narratives of Repression: Experimental Evidence from Sisi's Egypt." *Journal of Peace Research* .
- Wood, Michael J. 2015. "Some dare call it conspiracy: Labeling something a conspiracy theory does not reduce belief in it." *Political Psychology* pp. n/a–n/a.
- Wood, Michael J., Karen M. Douglas and Robbie M. Sutton. 2012. "Dead and alive: Beliefs in contradictory conspiracy theories." *Social Psychological and Personality Science* 3(6):767–773.
- Yablokov, Ilya. 2015. "Conspiracy Theories as a Russian Public Diplomacy Tool: The Case of Russia Today (RT)." *Politics* 35(3-4):301–315.