

The Supply of Conspiracism in State-Controlled Media*

Gabriel Koehler-Derrick[†] Richard A. Nielsen[‡] David Romney[§]

March 9, 2024

Abstract

When, and why, do governments promote conspiracy theories and use conspiracist language in official media? We build on claims that autocrats use misinformation for diversionary purposes by showing that the level of threat a regime faces affects conspiracism: a mix of “classic” conspiracy theorizing and conspiratorial language. Governments facing threats may attempt to stave them off by promoting conspiracy theories in state media. By contrast, secure governments are not as prone to conspiracism because it can be politically costly in the long-term. We test our arguments by examining conspiracism in Egypt’s print media between 2005 and 2018. When the government faced threats, the state-controlled newspaper prints more conspiracism than its independent counterpart. This relationship is moderated by changes in Egypt’s government: the state newspaper supplied less conspiracism during a brief period following relatively free and fair Presidential elections.

Abstract word count: 138

*Thanks to Ala’ Alrababa’h, Lisa Blaydes, Melani Cammett, Daniel Carpenter, Fotini Christia, Moustafa Hefny, Amaney Jamal, Brendan McElroy, Killian Clarke, Peter Henne, and audiences at the AALIMS conference, the University of Cairo, the University of Colorado Boulder, the University of Vermont, and the Middle East Studies Association for feedback and suggestions. Sabah Hamamou generously provided a copy of her book. Amani Selmi, Mehdy El-loumi, Elisha Cohen, and Elizabeth Parker-Magyar provided excellent research assistance. This research was made possible by a grant from the Carnegie Corporation of New York, supporting Richard Nielsen as an Andrew Carnegie Fellow.

[†]Assistant Professor, New York University, Abu Dhabi, gkd225@nyu.edu

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

[§]Assistant Professor, Brigham Young University, david.romney@byu.edu

1 Introduction

On August 27, 2013, Egypt’s most prominent state-run newspaper, *Al-Ahram*, published shocking allegations on its front page: “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 Patterson, 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.” The article prompted a stern response from Ambassador Patterson in which she rebuked the Egyptian government for spreading misinformation.² Did the conspiracy theory appear at the behest of the government? It seems likely, based on evidence from the subsequent trial of former President Morsi, which included transcriptions of testimony given by former President Hosni Mubarak alleging a plot by Hamas and Hizbullah to break into prison to free Mohammed Morsi and other members of the Muslim Brotherhood (El-Ghobashy, 2021, 231). While the former President Mubarak made no mention of Ambassador Patterson, his claim otherwise mirrors those published in *Al-Ahram*.

This blatant government promotion of a conspiracy theory angering an important ally is puzzling. The key to understanding this puzzle lies in the domestic threats confronting the Egyptian government. Two months earlier, Abdel Fatah al-Sisi ousted Egypt’s first democratically-elected President, Mohamed Morsi. On August 14th, 2013, security forces violently repressed those protesting the takeover, killing 1,000 demonstrators in Cairo’s Raba’a square and roiling the Egyptian public further. Against this backdrop of threat, the al-Sisi government turned to conspiracy theories and conspiracist language in the hopes of distracting angry Egyptian citizens. The costs of any long-term damage to the US-Egypt relationship paled in comparison to the immediate domestic threat of revolution.

In this paper, we offer theory and evidence to advance our understanding of *conspiracism* in

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

²Patterson’s response is no longer on the US Embassy website so we reproduce it in full in Appendix A.

official news media. We use the term “conspiracism” to refer to a mix of “classic” conspiracy theorizing and conspiratorial language that has similar qualities but is less clear on specifics; what Rosenblum and Muirhead (2019, 2) describe as the “new conspiracism” or “conspiracy *without* the theory.” There is little consensus about why states promote conspiracism. While some point to the psychology of individual political leaders or political culture (Gray, 2010*b*), the prevailing hypothesis is that autocratic regimes supply more (Gray, 2010*a*), in part to divert attention from potential challenges to the regime (Alrababa’h and Blaydes, 2020). We build on this insight by specifying when threats create these diversionary incentives, and propose that variation in threat, interacted with regime type, is an important cause of state conspiracism.

We first develop the concept of government supply of conspiracism: when a government intentionally alters the quantity and quality of conspiracy theories and conspiratorial language circulating in society. Governments can change both *what* and *how* they communicate to citizens: inventing detailed conspiracy theories, as in the case of Egypt’s accusations against Patterson, and altering the framing of government communication to invoke ideas of conspiracy without specifics. Changes in both quantity and quality of conspiracism are reflected in the language governments use.

Past studies have typically assumed that any conspiracism in state-controlled media appears at the behest of the state, but we argue that journalists and editors have agency, even at state-run media, and sometimes promote conspiracism for reasons apart from state intervention. To account for this, we conceptualize situations where a government intentionally increases the supply of conspiracism as *oversupply*, and situations where a government intentionally decreases the supply of conspiracism as *undersupply*. While the possibility of oversupply is the main focus of previous studies, we argue that undersupply can be just as strategic and important.

We develop theory and evidence about the oversupply and undersupply of conspiracism in state-controlled media. We build on the idea that autocratic governments face incentives to circulate conspiracism to divert from poor performance (Rozenas and Stukal, 2019) and manipulate domestic politics (Alrababa’h and Blaydes, 2020), but 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 conspiracism for governments. In the absence of threats, the costs of conspiracism are likely to outweigh the benefits, leading official media to undersupply conspiratorial language relative to independent media when the regime feels secure. However, as threats increase, regime elites are likely to oversupply conspiratorial language, prioritizing short-term gains over potential long-term costs. Regime type matters: more democratic governments have less media control and may pay higher costs for promoting conspiracism.

We test these arguments by examining the state supply of conspiracism in Egypt between 1998 and 2018. We collect all of the articles in Egypt's main state-controlled newspaper, *Al-Ahram* and examine the prevalence, timing, and content of conspiratorial language using qualitative coding and automated text analysis methods (Lucas et al., 2015). We compare *Al-Ahram* to the independent newspaper *Al-Masry Al-Youm*, which experts on Egyptian media agree is the most comparable privately-owned newspaper to *Al-Ahram* (Elmasry, 2012a). This comparison allows us to determine when journalists at the state-controlled newspaper are oversupplying and undersupplying conspiracism relative to journalists at a comparable independent newspaper (although *Al-Masry Al-Youm* still faces some constraints operating under authoritarianism (Cooper, 2008)).

Dramatic, unexpected developments in Egyptian politics over the last two decades offer an opportunity to estimate the effects of political threat on the state supply of conspiracy theories using observational data. We cannot experimentally manipulate Egypt's threats or institutions, but events in Egypt have proven unpredictable to both its rulers and outside observers, giving us some confidence that events and institutional change are driving conspiracism and not the reverse. Readers skeptical of our causal interpretation may nevertheless be interested in our description of conspiratorial language used by the Egyptian government over the last twenty years.

We find that the state-controlled newspaper undersupplies conspiracism relative to the independent newspaper when regime stability is not threatened. As destabilizing events increase, the

regime oversupplies conspiracism: at the high point of 150 threatening events in the past week, our model predicts that the state newspaper will print 18 conspiracist paragraphs compared to 8 in the independent newspaper. The degree of oversupply varies even as autocracy remains constant—most of the oversupply occurred under the autocratic Abdel Fatah al-Sisi regime, while the autocratic government of Hosni Mubarak undersupplied because it faced fewer threats. The democratically-elected Morsi government also undersupplied conspiracism, but our conclusions about regime type are tentative because the depth and quality of democracy was limited, and other characteristics of the Morsi government could explain this as well.

After examining state conspiracism quantitatively, we analyze it qualitatively using a nested analysis (Lieberman, 2005). We sample thirty days when our regression model predicts that the Egyptian government oversupplied conspiracism and find that, more often than not, conspiratorial language in the official newspaper are vague. Conspiratorial descriptions of specific perpetrators, victims, or incidents are often only tenuously connected to actual recent events. Conspiracism seems designed to increase a sense of threat and foreboding, portraying the state as the only defense against chaos and division.

Beyond Egypt, our theory offers a possible explanation for the supply of conspiracism by other states, both autocratic and democratic, though future research is needed to test and refine our arguments in other contexts. Our approach opens new possibilities for measuring and predicting the supply of conspiracism by various actors in other countries and contexts.

2 Propaganda and Conspiracism

Following Oliver and Wood (2014), we define conspiracism as political discourse with three key features: (1) “locat[ing] the source of unusual social and political phenomena in unseen, intentional, and malevolent forces,” (2) interpreting events “in terms of a Manichean struggle between good and evil,” and (3) asserting that “mainstream accounts of political events are a ruse” (953). In line with prior studies, our definition does not rely on whether a particular conspiracy theory is true.

Conspiracism is a form of information manipulation, alongside flooding, distraction, censorship, and misinformation (Lorentzen, 2014; Cirone and Hobbs, 2023; Roberts, 2018; Alrababa'h and Blaydes, 2020; Peisakhin and Rozenas, 2018; Wang and Huang, 2021). All can be considered propaganda: “the management of collective attitudes by the manipulation of significant signals” (Lasswell, 1927, 627). A growing body of literature examines the strategic use of propaganda in authoritarian contexts (Peisakhin and Rozenas, 2018; Rozenas and Stukal, 2019; Wang and Huang, 2021), but we argue that conspiracism presents opportunities and challenges that are distinct from other strategies in a regime’s misinformation toolkit (Keremoğlu and Weidmann, 2020).

In contrast to positive propaganda, such as fabricating good economic news (Rozenas and Stukal, 2019) or a leader’s extraordinary abilities (Wedeen, 2019), conspiracy theories are typically negative and invoke a threat. This makes them riskier for governments to deploy, because promulgating narratives that a regime is susceptible to conspiracies might suggest weakness. Yet conspiracism in the media has advantages over other forms of information manipulation. Media can sustain authoritarian rule by supporting regime priorities with content that caters to the interests of readers (Stockmann and Gallagher, 2011), and given the prevalence of belief in conspiracy theories (Oliver and Wood, 2014; Nyhan and Zeitzoff, 2018), readers are likely to find them interesting. Conspiracism appears more organic than other forms of hard propaganda, and may spread more easily through word of mouth, websites, and social media, much like political rumor (Berinsky, 2017; Huang, 2015*b*; Schon, 2020). Because of this seeming “bottom-up” quality, citizens who can recognize blatant misinformation may struggle to tell whether conspiracy theories are crafted by the government or not. When other forms of propaganda are unbelievable (Wedeen, 2019), conspiracy theories may garner belief because of the human desire to make sense of complex events (Oliver and Wood, 2014). In contrast to distraction or flooding, conspiracy theories can better shape citizens beliefs to match a regime’s desired narratives, creating desired actions instead of mere passivity. Thus, conspiracism has unique benefits and challenges for a regime and may follow a different strategic logic than other propaganda strategies.

We organize our discussion around three reasons states might supply conspiracism: genuine belief, symbolic power, and diversion. Some heads of state genuinely believe they are 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 always irrational. Autocrats may use state-controlled media to publicize their fears to build support and justify responses.

Alternatively, regimes might cynically promote conspiracism to craft narratives that enhance the state's symbolic power (Gray, 2010*a*, 133-134). State-promoted conspiracism does 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). Conspiracism can also communicate political priorities that assist coordination among political elites (Radnitz, 2016).

Finally, manufactured conspiracism 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 conspiracism 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 to show that "the Assad regime long sought to focus public attention on forces external to the regime, consistent with a logic of diversionary threat" (2). If governments justify repression by citing threats of conspiracy, then the diversionary use might follow the logic explained by Carter and Carter (2022, 671) who "expect propaganda-based threats of repression to be used sparingly: when the regime is most concerned about mass protests."

There is reason to believe that this diversionary conspiracism might work in the short term.

Experimental evidence shows that support for the Egyptian government's repressive policies increases when it accuses its opponents of conspiring to commit violence (Williamson and Malik, 2020). Violent political rhetoric has been shown in other contexts to increase support for political violence (Kalmoe, 2014). Even when diversionary conspiracism is not believed, it could prop up authoritarian governments by sowing fear. Huang (2015a) shows that Chinese students exposed to state propaganda remain dissatisfied with the government but become unwilling to openly dissent because they believe the regime is strong. Wedeen (2019, Chapter 4) similarly argues that the conspiracy theories circulating during the 2011 Syrian uprising and civil war bolstered the regime.

Autocratic diversion is the most prevalent explanation, but it faces theoretical challenges. Research suggests that even autocrats pay a price when they indulge in too much conspiracy promotion (Huang, 2018). Conspiracism may bolster the government in the short-term but make it look weak in the long-term. And blaming foreign conspirators can backfire diplomatically. What is missing from the existing literature is a clearer picture of when the threats facing a regime will outweigh the costs.

3 Costs, Threats, and Supply of Conspiracy Theories

Before we can theorize how the costs and benefits of conspiracism affect government strategies, we first refine our concept of supply to isolate conspiracism that is a result of government strategy. Differentiating conspiracism supplied by the state from conspiracism supplied by other actors is important because it probably has different causes.

Oversupply and Undersupply of Conspiracy Theories

Assuming that any conspiracism in state-controlled media appears at the behest of the government ignores the agency of editors and journalists. Journalists report on complex, destabilizing events and, like other people, may turn to conspiracism to make sense of them. Secretive or violent political events heighten uncertainty and anxiety, making journalists (and readers) especially prone

to conspiracism when straightforward explanations are elusive. Thus, we expect the news media, both state-controlled and independent, to sometimes supply conspiratorial language in response to unexplained events even without government manipulation.

We conceptually differentiate conspiracism supplied by the state from conspiracism by other actors, even when it appears in state-owned media. To guide this differentiation, we consider the counterfactual question: how would the supply of conspiracism in the state-controlled media be different if it were not state-controlled? This difference can be both quantitative, affecting the amount, and qualitative, affecting the content.

States can artificially increase or decrease conspiracism in the news media, and both may be strategically useful. If the state circulates quantitatively more conspiracism than independent media, we call this *oversupply*. Conversely, if state-controlled media circulates less conspiracism, this is *undersupply*. Conceptualizing the oversupply and undersupply helps us avoid assuming that all conspiracism is a result of the same causal process and alerts us to the possibility that governments might strategically undersupply conspiracism in official media outlets to avoid costs. It is to these costs that we now turn.

Costs of Promoting Conspiracism

Previous scholarship tends to suggest that governments, especially autocratic ones, supply conspiracism because of its effectiveness at neutralizing dissent and distracting citizens from poor political performance. Neutralizing dissent, both from citizens and from competing elites, is the central challenge of authoritarian rulers. Although “violence is an ever present and ultimate arbiter of conflicts in authoritarian politics,” (Svolik, 2013, 2), regimes often struggle to know who to repress (Gandhi and Przeworski, 2007; Gandhi, 2010), so non-violent ways of neutralizing opposition are an attractive alternative. The puzzle, then, is if the primary effect of spreading conspiracism is to increase support for government policies among some citizens and neutralize dissent among others, why wouldn’t governments just go back to the well over and over again?

We argue that existing theories do not incorporate the costs of conspiracism. While the up-front costs may be low — conspiracism is easy to produce, easy to disseminate, and does not need to be logically coherent to be effective (Wood, Douglas and Sutton, 2012) — the long term costs are high. Hard propaganda can worsen citizens’ opinions of the government, even when quelling open dissent. Over time, Huang (2018) argues, “by eroding the legitimacy of the state and public satisfaction, it may aggravate the regime’s long-term prospects” (1038). An over-reliance on conspiracism makes a government less credible about real threats. Wang and Huang (2021) show that when governments deny unfavorable information by labeling it “fake news,” citizens reduce their belief in future denials by the government if the unfavorable information is revealed to be true. They conclude that “false denials have both immediate and lasting effects on government credibility and can erode citizen satisfaction with the government.”

Information manipulation can backfire by encouraging citizens to seek out true information (Hobbs and Roberts, 2018; Gläβel and Paula, 2020), and obvious conspiracism might encourage citizens to seek alternatives to state media. For example, the Syrian government’s promotion of outlandish claims means “[s]tate-controlled newspapers in Syria are widely considered to be functional tablecloths, rather than respected records of current events” Wedeen (1999, 2). Other information manipulation can similarly backfire. Most broadly, although sowing confusion and paranoia among citizens can be a short-term ploy, it is not a recipe for long-term economic growth or effective public policy. And outside of a state’s borders, promoting conspiracism invites international criticism and complicates diplomatic relationships (Alrababa’h and Blaydes, 2020).

These costs vary by regime type. While they apply in some respect to all regimes, we argue that they are especially costly as governments become more democratic. Political elites in democracies face incentives to satisfy relatively large portions of their populations to stay in office, which encourages greater custodianship of public goods (Bueno de Mesquita et al., 2003). For a government attempting to retain office with good policies, promoting conspiracism is of limited use as a general strategy and has substantial risks. In democracies with strong political parties, parties have longer

time horizons than individual politicians, making them less likely to incur long-term costs for short term gains. Politicians and bureaucrats in democracies also face greater international opprobrium for promoting misinformation because freedom of the press is a democratic norm.

Threats to Government Survival

The factor that makes governments turn to promoting conspiracism in the face of these costs is the perception that their rule is threatened. This perception of threat is most commonly in response to events that are reasonably seen as threatening. Leaders fear that their rule will end because of war, armed attacks, coups, popular uprisings, protests, and strikes, and the political conditions that lead into these events, such as poor economic performance and underprovision of public goods (Bueno De Mesquita and Smith, 2023). Perception of threat does not need to be rational, and some leaders maybe more paranoid than others. Our simplifying assumption is that there are certain events and conditions that make all leaders fearful, and our argument is that these explain a great deal of the variation in conspiracy theory oversupply.

When deciding to undersupply or oversupply conspiracism, a government weighs the costs against the benefits. When a government feels secure, the costs of conspiracy promotion generally outweigh the benefits, and the government will undersupply. We argue that this calculus holds even in authoritarian regimes. Secure autocrats seeking to project an image of control and assurance are unlikely to promote conspiracism that undermines this image and provide little benefit. While even the strongest autocrat faces some opposition, if the threat of overthrow is not acute, it is not worth promoting conspiracism. Overuse when not it is not needed undermines its effectiveness when it is.

As threats become more frequent and acute, the government faces incentives to promote conspiracism to aid survival in the short term, even if there are long-term consequences. As noted in the literature, the short-term benefits of oversupplying conspiracism can be large, and might make the difference to regime survival. As events become especially complex and uncertainty grows, the

government may also get away with more far-fetched claims, because they may be more credible to both domestic and international audiences.

While short-term variation in threat is most important, government threat perceptions may shape supply in the long term as well. When leaders face relatively few threatening events over long periods of time, they may conclude that threats to their rule are low. This low baseline of threat perception may make some leaders slower to promote conspiracism even when they do face severe threats to their rule. In contrast, leaders that have faced many threatening events in the past, or come to power during tumultuous times, will likely maintain a higher baseline feeling of threat, even after threatening events subside.

While our theory predicts the quantity and timing of conspiracism, we have less to say about the content. We expect governments to use conspiratorial language that they think will be most effective at strengthening support and neutralizing opposition, and these strategies can vary over time even within the same state and regime (Alrababa'h and Blaydes, 2020). The literature suggests two broad strategies: messaging and misdirection. In a messaging strategy, governments might use conspiracist language with coherent content aimed to craft a particular narrative of events that promotes national cohesion Gray (2010*a*, 126). Alternatively, a government following a misdirection strategy might promote 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 conspiracism that state-controlled media oversupplies, we will examine conspiracist language deductively for evidence of messaging and/or misdirection.

Democracies face many of the same threats as autocracies, but we argue that they are less likely to respond by oversupplying conspiracism. The costs are likely higher for democracies, and this pushes the threshold higher for how threatened a democratic leader must feel before turning to conspiracism. Institutional features of democracy also make democrats less likely to promote conspiracy theories. Accountability to a large selectorate should incentivize leaders to provide more public goods (Bueno de Mesquita et al., 2003), including a media environment that is not domi-

nated by government misinformation. In most instances, the costs of losing office are less dire than those of autocrats, so events that threaten removal from office are less existentially threatening. Democracy also may reduce the control that a leader can have over the media, especially in established democracies, though elites can still communicate in ways that are not state-sanctioned. And using propaganda to entrench domestic power may be less attractive to leaders who have internalized democratic norms of deliberation and majority rule. Of course politicians in democracies may still turn to conspiracism for short term political gain, for all the same reasons as autocrats. Our prediction is that they will do so less, not that they won't do it at all.

Hybrid regimes, new democracies, and incomplete transitions to democracy present a variety of conditions that may make conspiracism more or less attractive to a regime. To the extent that the conditions of democracy we have just described hold, then we expect governments in these categories to also face higher costs for conspiracism, and thus generally turn to conspiracism less. Conversely, a new democratic regime that finds itself inheriting an official state media apparatus may face the temptation to use that media for its benefit, especially if democratic norms of media freedom have not yet been internalized and democratic institutions are not entrenched. We expect the state supply of conspiracism in these settings to follow the logic of costs and benefits we have outlined, but given the variety of institutional arrangements and incentives, whether the costs or benefits will dominate depends on the case.

Thus, regime type helps explain broad tendencies in conspiracism supply, but is not a complete explanation by itself. In addition to regime type, the degree of threat perceived by the government will impact the supply of conspiracism by the government. We expect to find variation in the promotion of conspiracism by autocracies: autocracies that feel secure will tend to undersupply, while autocracies that feel threatened will oversupply. Others have recognized the importance of threats. For example, Rozenas and Stukal (2019) finds that deflection in Russia's state-controlled media "is used more intensely in politically sensitive times (elections and protests)" (982). What we add is a clearer sense of how threats stack up against the costs of promoting conspiracies and

how costs and benefits interact with regime type.

4 Government Supply of Conspiracist Language in Egypt

We test how threats and changes to institutions affect the supply of conspiracism in state-controlled media by turning to data from Egypt. Despite the generality of our theoretical arguments, we follow the majority of previous studies in this area that draw on just one or two countries for empirical evidence (Peisakhin and Rozenas, 2018; Hobbs and Roberts, 2018; Rozenas and Stukal, 2019; Alrababa'h and Blaydes, 2020). Investigating a single country presents trade-offs: we gain the ability to understand conspiracism deeply in context at the expense of uncertainty about how our findings might extend elsewhere. Like others, we choose this approach because conspiracist language is rich in contextual meaning and getting that right is a first-order concern. Future research might extend elements of our approach to multiple countries.

We select Egypt for scientific and practical reasons. In order to test the effects of threat and regime type on state supply of conspiracism, we need variation in both key variables. During the time period we examine, Egypt experienced 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 language skills and fieldwork experience allow us to analyze these materials quantitatively and qualitatively.

We measure conspiracism using the prevalence and content of conspiracist language published by the main state-owned 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. Conspiracist language also circulates on television and social media in contemporary Egypt (Armbrust, 2019, Chapter 9), 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 leave an investigation of conspiracism in Egyptian social media for future research.

Our goal is to explain variation in the conspiracist language printed in *al-Ahram* at the behest of the state, rather than for other reasons. Attributing conspiracist language to government intervention requires either intimate knowledge of the process through which a government pressures journalists to choose their words, or comparison to other media that are not state-controlled. We cannot directly trace the process by which each instance of conspiracist language appears in the newspaper, so we compare conspiracism in *al-Ahram* to Egypt's main independent newspaper, *Al-Masry Al-Youm*. When more conspiracist language appears in *al-Ahram* than *Al-Masry Al-Youm*, we consider this to be oversupply of conspiracism by the Egyptian government. Conversely, when more conspiracist language appears in *Al-Masry Al-Youm*, we infer that the government is under-supplying conspiracism. At a high level, our research strategy is to detect conspiracist language in each newspaper and then estimate regression models predicting the supply of conspiracy theories in each, which we interrogate qualitatively with nested analysis.

Because our inferences rely crucially on these newspapers, the next subsection considers the degree to which they match our concepts of state-controlled and independent media.

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. It is not the only state-owned newspaper, and there is heterogeneity among the others, as in other contexts (Stockmann, 2012), but *Al-Ahram* is unique because of its broad circulation and influence. *Al-Ahram's* editor-in-chief is directly appointed by the Egyptian government serves as a regime mouthpiece, with a majority of domestic stories citing government officials as their main source (Cooper, 2008, 5). Numerous accounts allege direct intervention in *Al-Ahram's* editorial process (Hammond, 2005, and see Appendix B). There are allegations that it promotes conspiracy theories from "informed sources" and anonymous

security sector officials, especially in the wake of events that challenge the regime.³

Al-Ahram's status as the most important daily newspaper in Egypt was uncontested until the emergence of *Al-Masry Al-Youm* in 2004. *Al-Masry Al-Youm* derives relative independence from its financial backing by Salah Diab, a powerful businessman with a background in oil (Al-Azm, 2015). *Al-Masry Al-Youm*'s editors are independently appointed and change fairly frequently, and we have not uncovered accounts of direct government intervention in its editorial process. Still, it has faced constraints in authoritarian Egypt. *Al-Masry Al-Youm* and its owner have occasionally come under government pressure (Peterson, 2011), and the pressure has become more acute in the aftermath of the 2013 coup which brought Egypt's current president Abdel Fatah al-Sisi to power.

Our direct comparison of *Al-Ahram* and *Al-Masry Al-Youm* follows the precedent set by more than a decade of research. Since *Al-Masry Al-Youm*'s founding in 2004, studies of print media in Egypt, whether quantitative or qualitative, have directly compared the two newspapers to gain insights on censorship, editorial autonomy, and public support for independent media (Black, 2008; Cooper, 2008; Elmasry, 2012a; Peterson, 2011). The basis for this comparison is straightforward: no other daily in Egypt compares to either paper in terms of circulation, audience, or reputation (Badr, 2021, 225). Subscriber and circulation data are not available 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, the web analysis service *Alexa* rated *Al-Masry Al-Youm* as the 44th most popular website in Egypt while *Al-Ahram* was 31st. While there are also other official newspapers in Egypt, *al-Ahram* is considered the "flagship" of the official press. *Al-Masry Al-Youm* has been described as "Egypt's largest and most reliable private newspaper" (Ketchley and El-Rayyes, 2021, 294). As one ethnographic study argued the two newspapers are, "leaders in their respective ownership categories and, to a considerable extent, representative of the ownership categories to

³<https://dailynewsegypt.com/2016/10/02/al-ahram-blindly-parrots-governments-forces-evil-claim/> and <https://dailynewsegypt.com/2016/09/19/548755/>, Accessed 3/21/2021.

which they belonged” (Elmasry, 2012b, 122). Still, to alleviate concerns that our results are reliant on just two data sources, we include data from two other Egyptian newspapers, one state-owned and one independent, in a robustness check.

While *al-Masry al-Youm* provides the most plausible counterfactual for what *Al-Ahram* might publish without government intervention, we can hardly claim that it is free from political constraints (Cooper, 2008, 8). This does not necessarily mean it is not independent. After all, the pressure on *al-Masry al-Youm* is a result of the paper’s willingness to defy the regime. Still, in autocracies, it is rare that any media outlet is truly beyond any state influence, so empirically, we must make do with the most independent media there is. While this may lead to some bias when measuring oversupply and undersupply, the bias is likely conservative: if there are gaps between state-controlled media and partially independent media, we would expect even larger gaps with fully independent media. If *Al-Masry Al-Youm* were no more independent than *Al-Ahram*, then we would presumably see no differences between the supply of conspiracy theories.

Detecting Conspiracy Theories in Arabic Newspaper Text

We collect all of *Al-Ahram* since 1998 and all of *Al-Masry Al-Youm* since 2005, up until 2018 (comprising 826,765 articles in *Al-Ahram* and 484,100 in *Al-Masry Al-Youm*). Reading over a million articles would take at least 20 years, so we turn to statistical text classification. Our approach 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. Our process was complicated by the fact that conspiracist language is much rarer than prior studies relying on anecdotes suggested. We initially began coding a simple random sample of articles, but we stopped because the training set had so little conspiracist language that subsequent classification algorithm would have failed. Instead, we turned to a key word approach because our reading revealed that articles promoting conspiracism are forthright, almost always employing some variant of the word “conspiracy.” We selected 18 key words that our initial reading revealed were frequently used to introduce conspiracism (forms

the hand-coded data to assess accuracy; we correctly identify conspiracist paragraphs 78% of the time, and correctly identify non-conspiracist paragraphs 99% of the time, for an overall accuracy of 97.5% (specificity = .990, sensitivity = .787, F1 = .877). We then apply the trained 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 conspiracist paragraphs across the two newspapers.

To validate our measure, we returned to close reading. First, we read a random sample of 500 paragraphs from all articles, looking for conspiracist paragraphs our key word approach missed. We found just one conspiracist paragraph and it was also successfully identified using our key word classification approach. We also read paragraphs the algorithm labeled conspiracist in our qualitative analysis below, which provided an additional opportunity to evaluate the quality of the automated classification. We find no significant classifier errors.

To ensure we were not missing conspiracist language because of our key words, we undertook an additional classification task in the spirit of King, Lam and Roberts (2017) seeking to identify other paragraphs that ought to have been included. Although our classifier is trained on articles with key words, the approach of King, Lam and Roberts (2017) suggests that we could grow our list of key words and indeed identify misclassified conspiracist paragraphs by iteratively applying our classifier to additional texts that do not contain our key words. Paragraphs that are scored highly by the classifier are worth considering for inclusion in our count of conspiracist paragraphs, even if they don't contain our key words. We applied our classifier to the 1.2 million articles without our key words and identified just a few hundred additional conspiracist paragraphs which make no difference to our estimates below (see Appendix F.4).

We also consider that newspapers might report conspiracy theories to criticize, rather than promote them. The effects on readers may be same; Berinsky (2017, 241) finds that “attempting to quash rumors through direct refutation may facilitate their diffusion.” Moreover, mentioning a conspiracy theory neutrally, or even critically, can be a strategy for promoting it while maintaining plausible deniability. Still, we asked our coders to evaluate how each conspiracy theory in the

training set was framed and find that 54% are endorsed by the author, 17% criticized, and 29% presented neutrally. Our results are robust to including only the conspiracist paragraphs that we estimate are endorsed by the author. Taken together, these results build our confidence that there is no significant set of false negatives or false positives lurking in our data.

Several facts about the nature of conspiracism in Egyptian news media emerge from closely reading our training set. First, Egypt is the victim in about half; many highlight Egypt at an institutional level (e.g. “the Egyptian state,” “the Egyptian economy,” and “the Egyptian national character”) while others emphasize Egyptians themselves (e.g. “the Egyptian people” or “the sons of Egypt”). Second, approximately 45% of conspiracism have clearly alleged perpetrators, victims, and actions, while the rest is vague. Third, when perpetrators are mentioned, they are most often the United States (25%) and/or Israel (20%).

Measuring Threats and Regime Type

Our argument is that both undersupply and oversupply of conspiracism by the Egyptian government can be explained by the level of threat it faces and whether its institutions are more or less autocratic. Both variables present challenges for measurement.

If we could measure the degree of threat perceived by the Egyptian government, this would be our ideal variable for predicting government supply of conspiracist language. Elite perceptions of threat are, for the most part, unobservable, and inferring them from public statements is prohibitively difficult, at least with any granularity and accuracy. Instead, we assume that leaders feel threatened when they face destabilizing political events, which are more readily observable. Theories of authoritarian rule highlight two main categories of threat: the risk of overthrow from sustained protests or attack and the risk of coups (Svolik, 2013, 4). Protest may not be universally threatening to all authoritarian regimes (Lorentzen, 2013), but observers agree that it has been a major source of threat in Egypt (El-Ghobashy, 2021). To measure the intensity of these threats to the Egyptian government, we use counts of significant attacks, battles, and protests in Egypt, as

reported in the ACLED data set (Raleigh et al., 2010).

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.” We treat these counts with some caution. Clarke (2021) shows that ACLED undercounts peaceful, localized, and rural protest events in Egypt between 2012-2013. There are also concerns that ACLED might miss more protests in Egypt prior to 2011. If our study were focused on protest dynamics instead of threat, these biases might make ACLED unusable. For the purpose of measuring threat to the government, however, the fact that ACLED is primarily accurate for larger, more violent events of all types makes it useful. We find similar results with other measures of threat, including Clarke’s improved measure of protest, and terrorism threats from Global Terrorism Database (LaFree and Dugan, 2007).

For the most part, measuring Egypt’s degree of autocracy is straightforward: all governments between 1998 and 2018 have been autocratic, with the exception of Muhamad Morsi’s presidency from 2012–2013. Autocracy under al-Sisi, whether ruling through interim president Adly Mansour (2013–2014) or directly, (2014 onward) is seen by many as more repressive than Mubarak’s last years (until 2011), or the period from 2011-2012 when the Supreme Council of the Armed Forces (SCAF) ruled. Yet despite variation, all of these regimes fit standard definitions of autocracy.

Classifying the Morsi government is more difficult. Morsi was democratically elected in an election that was widely declared to be relatively free and fair.⁵ There was more press freedom during Morsi’s rule. While *Al-Ahram* remained under state ownership, the Morsi government had less informal control because the newspaper remained staffed with supporters of the former regime (see Appendix B), and Morsi did not forcibly replace them as an autocrat might. Yet some are uncertain whether the administration should be viewed as democratic because the Muslim Brotherhood, Morsi’s party, had questionable ideological attachments to the principles of democracy. In

⁵V-Dem’s polyarchy measure of democracy jump dramatically in 2012 (see Appendix Figure 15, and V-Dem’s clean election index scores 2012 as the best in Egypt’s history since 1900 <https://www.v-dem.net>, accessed 10/19/2021).

any case, democracy was hardly consolidated during Morsi's rule. Norms of democratic power-sharing were not established, and unlike most democracies, Morsi faced imprisonment and death when he lost power. Anticipating this possibility likely heightened the threats he felt in office.

On balance, we consider Egypt to be more democratic under Morsi than at any other time, but the ambiguities make our findings about the effects of regime type frustratingly uncertain. Alternative research designs offer few remedies. Cross-national comparisons with other countries are unavoidably confounded by country-specific factors, and instances of staunch autocracies transitioning to full-fledged democracy in a short amount of time are rare. With these limitations in mind, the sudden transition of Egypt to holding relatively free elections in 2012 after decades of autocracy is about the best case for analysis we can hope to get in the real world.

The Quantitative Supply of Conspiracist Language in Response to Threats

We turn first to explaining variation in the quantity of conspiracist language supplied by the Egyptian government between 1998 and 2018. We examine the content qualitatively in the next section.

We seek to explain the variation in conspiracist paragraphs over time, which we plot in Figure 1 for *Al-Ahram* (red) and *Al-Masry Al-Youm* (blue). We show both the (smoothed) moving average in solid lines, and the average for each government in dashed lines. We also plot our key theoretical variables: dark gray shading indicates the moving average of significant contentious political events and alternating light gray and white shading indicates the timing of transitions from one government to the next, with text labels denoting the head of state.

Figure 1 visually displays several correlations that persist in our regression results. First, as more protests, attacks, and battles occur in Egypt, there is more conspiracist language in both newspapers. Second, Mubarak faced very few threatening events relative to other administrations, suggesting that he may have perceived a lower level of threat than his successors. Third, the supply of conspiracist language in *Al-Ahram* changes dramatically with regime changes and appears broadly correlated with the number of threatening events.




Figure 1: Average counts of conspiracist paragraphs in each newspaper. The solid lines indicate the moving average. Dashed lines indicate the average during each administration.

We estimate these correlations more formally with regression. Our outcome is the count of conspiracist paragraphs in each newspaper on each day, so we use a generalized linear model with a negative binomial link appropriate for (potentially over-dispersed) counts. We proxy threat to the government 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*.

Without the ability to manipulate threatening events, we adopt a strategy of conditioning on likely confounders to make inferences about the effect of threatening events on conspiracism supply. The unexpected and timing of these threatening events makes it plausible that they caught Egypt’s government by surprise; we argue that they are plausibly exogenous after we condition on just a small number of potential confounders. We include several control variables that might confound the relationship between the events-newspaper interaction and conspiracy theories. Some articles in both newspapers are drawn from international news agencies which are unlikely to pro-

mote conspiracism, so we control for the count of paragraphs and articles from these sources. We also control the total number of paragraphs and articles in each issue of the newspaper, regardless of source, because both newspapers get longer over time. Finally, we control for the day of the week, expecting that journalistic practices might vary day by day. In some models, we add year-month fixed effects, which are indicator variables for each calendar month between 2005 and 2018. These models restrict the variation to one-month periods, showing that government supply is responsive to daily events. These fixed effects also address concerns that ACLED is not comparable over long time frames; it is at least comparable in the time-frame of a month.

Rather than including changes of regime in our primary specification, we instead estimate oversupply and undersupply in response to threats and these controls, and then observe qualitatively how periods of oversupply align with changes in regime. This is largely for convenience — interacting regime changes with the event-newspaper interaction results in a triple interaction that is complicated to interpret. The results are substantively the same when we fit that more complex model. We cannot claim that changes in regime or regime type are plausibly exogenous, because governments likely use conspiracism to try to prevent regime change when facing unexpected threats.

We present our main regression results graphically in Figure 2, with a table in Appendix E. The model coefficients are not directly interpretable, so we plot changes in predicted counts of conspiracist paragraphs. Figure 2 shows our main result: *Al-Ahram* supplies more conspiracist paragraphs than *Al-Masry Al-Youm* in response to the same events. We graphically present predicted counts of conspiracist paragraphs in each newspaper as ACLED events increase from zero to the maximum of 151, along with the histogram of ACLED events on the x-axis. The left panel, corresponding with the numerical results reported in Table 2 of the Supplement, shows that when there are no threatening events, the state-controlled newspaper supplies fewer conspiracist paragraphs: 2.2 per day for *Al-Ahram* and 2.7 per day for *Al-Masry Al-Youm*. The predicted values cross when there have been 28 ACLED events in the last week. When ACLED records 150 events, *Al-Ahram* prints 17.9 conspiracist paragraphs per day, 2.3 times as many as *Al-Masry Al-Youm*.

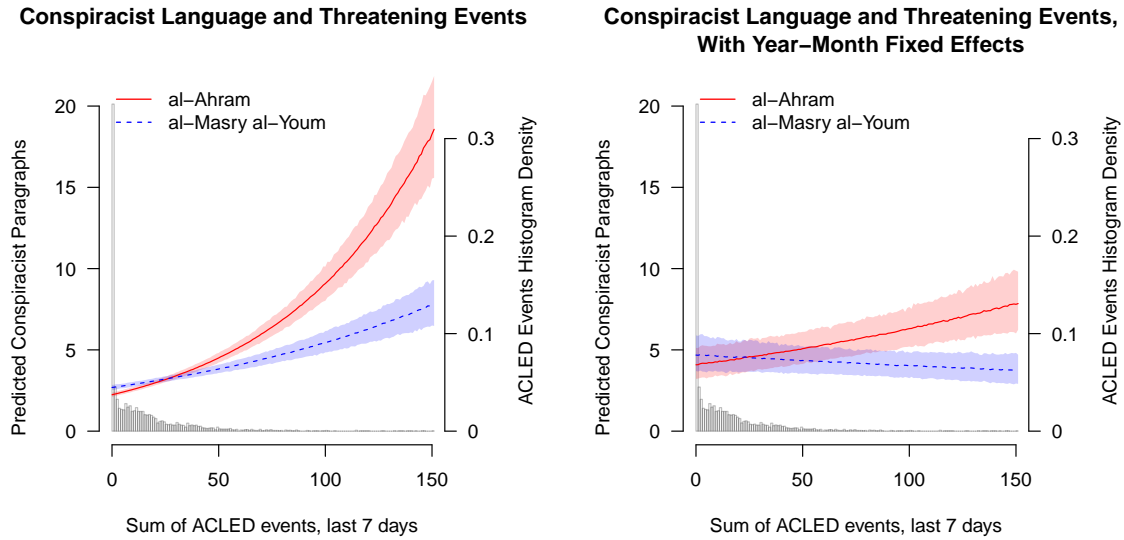


Figure 2: *Al-Ahram oversupplies conspiracist paragraphs relative to Al-Masry Al-Youm, given the same events.*

The right panel shows predicted conspiracist paragraph counts with year-month fixed effects. *Al-Ahram* is still responsive to threatening events, even within a single month; we choose July 2013 for presentation, but changing the month just moves the baseline up or down. As the count of events increases, conspiracist paragraphs in *Al-Ahram* increase from from 4.4 to 7.8 when there have been 150 events in the last 7 days. In contrast, the predicted number of conspiracist paragraphs in *Al-Masry Al-Youm* remains constant or declines slightly as the number of events increases. Figure 2 shows that *Al-Ahram* diverges most from *Al-Masry Al-Youm* when ACLED records at least 50 events in the past week, which is true for 226 out of 4,429 days, or 5.1%.

To better understand the timing of oversupply by *Al-Ahram*, we generate estimates from model 1. We calculate the daily undersupply 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 conspiracist paragraphs 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 3, with the x-axis indicating

the date and the y-axis indicating the estimated undersupply (negative values) or oversupply (positive values) by *Al-Ahram*. We show a moving average of the under/oversupply estimates, using blue for periods of undersupply and red for periods of oversupply.

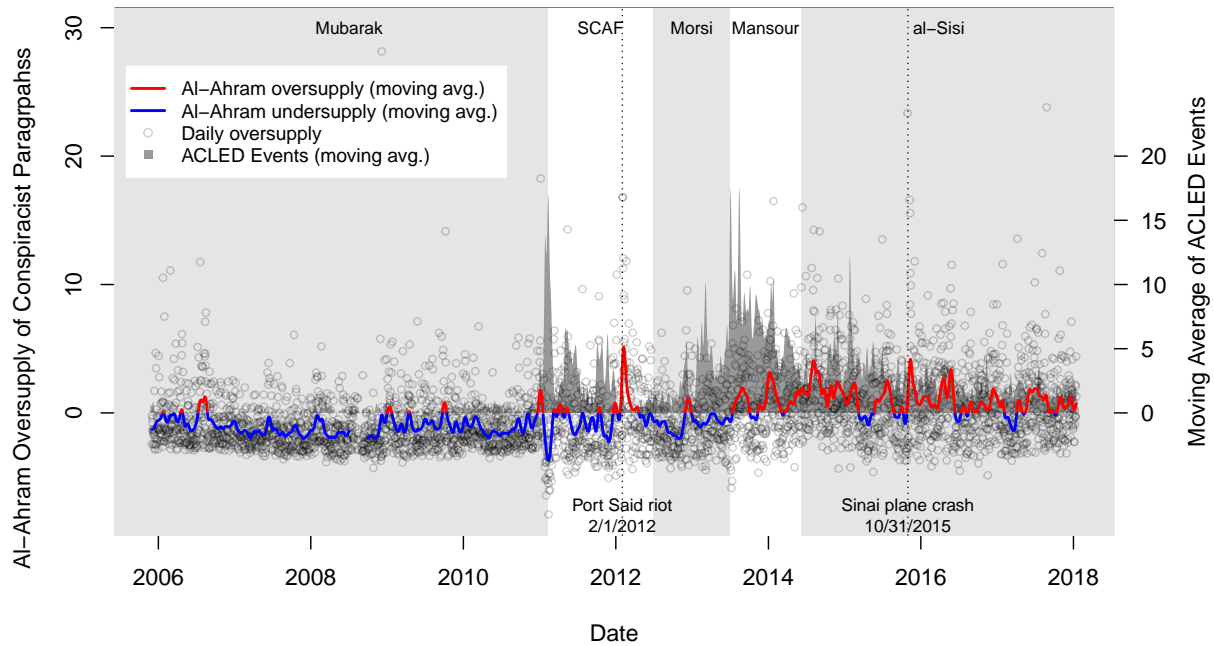


Figure 3: *Oversupply of conspiracist paragraphs in Al-Ahram, based on counterfactual estimates from a negative binomial regression model of how many conspiracist paragraphs would have appeared if Al-Ahram had been Al-Masry Al-Youm on each day.*

The state-controlled newspaper consistently undersupplied conspiracism under Mubarak. One exception is that *Al-Ahram* begins to oversupply conspiracism 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. This oversupply was Mubarak’s last-ditch effort to divert attention as his grip on power weakened, and the undersupply reflects attempts by *Al-Ahram*’s journalists to cover revolutionary events impartially after Mubarak’s departure temporarily decreased government influence over the newspaper (see Appendix B). Undersupply remains the norm under the SCAF, except for a period of dramatic oversupply after the Port Said soccer riot, an embarrassing event for the interim regime. The democratically-elected Morsi government undersupplied

conspiracism despite facing increasing numbers of threatening events. This is consistent with the argument that more democratic governments supply less conspiracism, but this result could also be the result of factors idiosyncratic to the Morsi administration.

Al-Ahram has consistently oversupplied conspiracist language under the Mansour and al-Sisi governments. Large upticks follow particularly threatening events, such as the downing of a Russian airliner in the Sinai by terrorists. Our data suggest that the government mobilizes a particular set of *Al-Ahram* authors to promote conspiracism after threatening events. Most newspaper articles list one or more authors; we observe that a relatively small number of authors — approximately 175 — produce the vast majority of conspiracist language, and these authors are statistically more likely to write when the number of ACLED events was high the previous day.

These results show that regime type alone is not sufficient to explain variation in state conspiracism, which we confirm by re-estimating the same specification in model 1 separately for the tenure of each executive. Although Mubarak, the SCAF, and Mansour/Al-Sisi governments were all autocratic, the general approach of Mubarak was to undersupply conspiracist language, while the general approach of Al-Sisi (who effectively controlled Mansour) is to oversupply. This supports our argument that unthreatened autocrats are likely to undersupply conspiracism.

Our results are robust to a large number of other modeling and measurement alternatives: Poisson regression; linear regression; aggregating conspiracist language at the article level; counting only conspiracist language that is endorsed (rather than criticized) by the reporting 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;

omitting combinations of control variables; adding indicator variables for each regime; and a few more. Details are in Appendix F.

Additionally, we consider the possibility that our results are spurious because of limitations or biases in the ACLED data. One alternative is to focus on terrorism, which is a subset of the events we think might be threatening that is highly observable, and for which alternative data sources are available. We measure terror threats in Egypt using event counts from the Global Terrorism Database (LaFree and Dugan, 2007), which have a 0.52 correlation with our ACLED measure. Our results are robust to this alternative measure of threat (see Appendix F).

New protest data hand-coded by Clarke (2021) from *Al-Masry Al-Youm* offers a significant improvement over ACLED for measuring protests in Egypt between January 2012 and July 2013. Several measures of protest can be constructed from these data; we prefer a measure that counts the sum of ongoing anti-regime protests as this best captures protests that threaten the government. Our analysis with ACLED finds that *Al-Ahram* generally undersupplies conspiracist language in 2012-2013, and we find similar evidence of undersupply using Clarke's measure (see Appendix F). We cannot test whether our findings from the Mubarak, Mansour and al-Sisi governments hold because of the temporal limits of Clarke's data.

One other robustness check deserves mention. To allay concerns that our findings might be dependent on idiosyncrasies of the two newspapers, we collected 775,126 articles (from 3,230 days) from another independent newspaper *Al-Shuruq*, and 75,196 articles (from 878 days) from a less-prominent state-controlled newspaper, *Al-Gomhuria* (see the Appendix F). Our results remain substantively similar with these additional newspapers.

Our conclusions from these models are as follows. We find evidence that both newspapers print more conspiracist language in response to threatening events, but differentially so. The official paper, *Al-Ahram*, is far more responsive to threatening events than the independent *Al-Masry Al-Youm*. These differences are substantively large; when the count of events changes from its minimum to its maximum, *Al-Ahram* increases its supply of conspiracist paragraphs approximately sevenfold,

from 2.5 to 18 per day. The oversupply in *Al-Ahram* responds to day-to-day changes in threatening events, suggesting government intent. This oversupply of conspiracist language is concentrated in times when autocratic Egyptian governments faced large numbers of threatening events, and is strongest since the 2013 coup. By contrast, the authoritarian government of Mubarak preferred to undersupply conspiracism from the mid-aughts until it faced extreme crisis at the end of 2010. Neither threat nor authoritarianism alone is a sufficient explanation. Rather, when authoritarian governments feel threatened, they are more likely to print conspiracist language.

Qualitative Analysis of Oversupplied Conspiracism

To infer more about the Egyptian government's goals, we qualitatively examine some periods of oversupply. Our goal is to inductively learn what types of conspiracist language the state-controlled newspaper promotes when it oversupplies, so 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 conspiracist paragraph, summarizing the alleged conspiracy and noting incidents or actors matching those recorded by ACLED 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* conspiracist language on these days have little connection to recent events, suggesting that the state intends to misdirect rather than send a message about what is happening.

State-controlled media does reference recent events in conspiracist language following some pivotal events. Our sample included October 12, 2011, immediately following the "Maspero Massacre," in which several dozen peaceful protesters, primarily Coptic Christians, were killed and

hundreds wounded when security forces attacked them in front of the Maspero television building. Four of the seven articles with conspiracist language referred to the incident, but primarily allege the presence of evil forces threatening the Egyptian people without providing a detailed account or naming a specific plot or perpetrator. 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 conspiracist language in *Al-Ahram* we randomly selected for close reading is formulated to promote unity by invoking vague fears. The formula consists of (1) vague references to 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 these nefarious threats.⁶ They also describe Egypt’s leaders as astute, vigilant to threats that others are too naive to see. Lastly, these articles often explicitly criticize dissenters, especially those who are skeptical of the regime’s conspiracism. For instance, on February 24, 2016, two years after the coup that brought al-Sisi to power, an article 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 engage in conspiracism? We offer new concepts, theory, and evidence to answer this question. We introduce the concepts of oversupply and undersupply of conspiracism

⁶This is resonant with efforts to restore the “prestige of the state,” documented by El-Ghobashy (2021, Chapter 6).

by governments. Conspiracism oversupply happens when state-controlled media spreads more conspiracist language than if it were independent; undersupply happens when state-controlled media spreads less. Several important studies have examined changes in the amount of conspiracist language 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. Independent media also promote conspiracism, so it is not merely the appearance of conspiracist language in state-controlled media that needs explanation. Rather, we need to explain why state-controlled media sometimes promote more or less than independent media.

To explain the over- and undersupply of conspiracism by states, we take two theoretical factors that are considered separately in the literature—regime type and diversionary incentives—and bring them together. We agree with previous studies that states face diversionary incentives to promote conspiracism when they are under threat, but these incentives are offset by the risk that doing so will backfire. We argue that official media will undersupply conspiracism in the absence of threats. As threats increase, the official media is much more likely to oversupply conspiracism. Our argument applies to some democracies as well as autocracies—regime elites in democracies can feel threatened as well—but autocrats are more likely to control a media outlet and face fewer incentives to preserve the integrity of the media.

Most empirical studies of state conspiracism focus solely on the dynamics of autocratic politics, often under a single stable regime. In Egypt, rapid shifts in regime type allow us to test whether political institutions matter while holding the media market constant. We measure the supply of conspiracist language in more than one million Egyptian newspaper articles over twenty years, which allows comparisons across time and between newspapers that were not previously possible.

In line with our theory, we find that state-controlled media begins to oversupply conspiracist language as events threaten the regime; most of this oversupply has happened during Egypt's most recent autocratic administrations. Our qualitative examination of this conspiracism reveals that it typically encourages societal solidarity and stability in the face of vague threats rather than craft-

ing a narrative about specific recent events or actors. During Egypt's brief period of democratic rule, from 2012–2013, we find the opposite pattern: the state-owned newspaper undersupplies conspiracism despite high numbers of threatening events. This suggests that democracies may undersupply, but our inferences are necessarily tentative.

The tentative nature of our findings about regime type point to an important limitation of our study. While our use of observational data has many strengths, we are unable to vary the timing or values of any of the key causal variables in our theory. We cannot manipulate either regime threat or regime type, so the statistical correlations between these factors and conspiracism could be due to confounding factors. Manipulating either of these variables strikes us as impossible, but perhaps future scholars will find situations where threat or regime type have varied randomly to allow causal inference with less threat of confounding.

Our theoretical and conceptual contributions suggest avenues for future exploration. For scholars of political communication, our approach could be extended to study the supply of conspiracism by states across the Arab World and beyond. Data limitations force us to analyze data from a single country; future studies could compare state-controlled and independent media in other countries, within and outside the Middle East region. In places like the United States where media independence is strong, our methods suggest new ways to analyze the supply of conspiracist language by non-state actors across the political spectrum. Considering applications to established democracies reveals a conceptual shortcoming of our work: that our notions of over- and undersupply are premised on the presence of some form of state-controlled media. Analyzing over- and undersupply of conspiracism in countries without state-controlled media will require modifications to our approach, but we argue that the general framework still applies in countries such as the United States where public officials sometimes spread conspiracy theories through other means.

Our study also has implications for scholars of political psychology. To date, most experimental studies do not consider variation in the supply of conspiracy theories, and those that note this variation work to control it away (Ryan and Aziz, 2021). We can imagine a research agenda that

more closely mimics the wide variation in supply in the real world and experimentally varies timing and intensity to learn how conspiracy theory supply affects demand.

Beyond questions of conspiracism, the study of Arabic-language news media in political science is in its infancy, and exciting new research shows 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.
- Alrababa'h, Ala'. 2021. Manufacturing threats: Diversionsary discourse and autocratic survival PhD thesis Stanford University.
- Alrababa'h, Ala' and Lisa Blaydes. 2020. "Authoritarian media and diversionsary threats: lessons from 30 years of Syrian state discourse." *Political Science Research and Methods* p. 1–16.
- Armbrust, Walter. 2019. *Martyrs and Tricksters: An Ethnography of the Egyptian Revolution*. Princeton University Press.
- Badr, Hanan. 2021. Egypt: A Divided and Restricted Media Landscape after the Transformation. In *Arab Media Systems*. Ebsco open editions ed. Berlin: .
- 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.
- Bueno De Mesquita, Bruce and Alastair Smith. 2023. "Institutional Change as a Response to Unrealized Threats: An Empirical Analysis." *Journal of Conflict Resolution* 67:1032–1062.
- Bueno de Mesquita, Bruce, Alastair Smith, James D. Morrow and Randolph M. Siverson. 2003. *The Logic of Political Survival*. Cambridge, MA: MIT Press.
- Carter, Erin Baggott and Brett L Carter. 2022. "When Autocrats Threaten Citizens with Violence: Evidence from China." *British Journal of Political Science* 52(2):671–696.
- Cirone, Alexandra and William Hobbs. 2023. "Asymmetric flooding as a tool for foreign influence on social media." *Political Science Research and Methods* 11(1):160–171.

- Clarke, Killian. 2021. "Which Protests Count? Coverage Bias in Middle East Event Datasets." *Mediterranean Politics* pp. 1–27.
- Cooper, Kenneth J. 2008. "Politics and priorities: Inside the Egyptian press." *Arab Media & Society* (06).
- El-Ghobashy, Mona. 2021. *Bread and Freedom: Egypt's Revolutionary Situation*. Stanford University Press.
- Elmasry, Mohamad Hamas. 2012a. "Journalism with Restraint: A Comparative Content Analysis of Independent, Government, and Opposition Newspapers in pre-Revolution Egypt." *Journal of Middle East Media* 8(1).
- Elmasry, Mohamad Hamas. 2012b. "Producing News in Mubarak's Egypt: An analysis of Egyptian newspaper production during the late Hosni Mubarak era." *Journal of Arab & Muslim Media Research* 4(2):121–144.
- Gandhi, J. and A. Przeworski. 2007. "Authoritarian Institutions and the Survival of Autocrats." *Comparative Political Studies* 40(11):1279–1301.
- Gandhi, Jennifer. 2010. *Political Institutions Under Dictatorship*. New York: Cambridge University Press.
- Gläbel, Christian and Katrin Paula. 2020. "Sometimes less is more: Censorship, news falsification, and disapproval in 1989 East Germany." *American Journal of Political Science* 64(3):682–698.
- 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.
- 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.
- Huang, Haifeng. 2015a. "Propaganda as signaling." *Comparative Politics* 47(4):419–444.
- Huang, Haifeng. 2015b. "A war of (mis)information: The political effects of rumors and rumor rebuttals in an authoritarian country." *British Journal of Political Science* FirstView:1–29.
- Huang, Haifeng. 2018. "The pathology of hard propaganda." *The Journal of Politics* 80(3):1034–1038.
- Kalmoe, Nathan P. 2014. "Fueling the Fire: Violent Metaphors, Trait Aggression, and Support for Political Violence." *Political Communication* 31(4):545–563.

- Keremoğlu, Eda and Nils B Weidmann. 2020. "How Dictators Control the Internet: A Review Essay." *Comparative Political Studies* 53(10-11):1690–1703.
- Ketchley, Neil and Thoraya El-Rayyes. 2021. "Unpopular Protest: Mass Mobilization and Attitudes to Democracy in Post-Mubarak Egypt." *The Journal of Politics* 83(1):291–305.
- King, Gary, Patrick Lam and Margaret E Roberts. 2017. "Computer-Assisted Keyword and Document Set Discovery from Unstructured Text." *American Journal of Political Science* 61(4):971–988.
- LaFree, Gary and Laura Dugan. 2007. "Introducing the global terrorism database." *Terrorism and political violence* 19(2):181–204.
- Lasswell, Harold D. 1927. "The Theory of Political Propaganda." *American Political Science Review* 21(3):627–631.
- Lieberman, Evan S. 2005. "Nested Analysis as a Mixed-method Strategy for Comparative Research." *American Political Science Review* pp. 435–452.
- Lorentzen, Peter. 2014. "China's Strategic Censorship." *American Journal of Political Science* 58(2):402–414.
- Lorentzen, Peter L. 2013. "Regularizing Rioting: Permitting Public Protest in an Authoritarian Regime." *Quarterly Journal of Political Science* 8(2):127–158.
- 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.
- Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S Corrado and Jeff Dean. 2013. "Distributed representations of words and phrases and their compositionality." *Advances in neural information processing systems* 26.
- 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.
- Radnitz, Scott. 2016. "Paranoia with a purpose: conspiracy theory and political coalitions in Kyrgyzstan." *Post-Soviet Affairs* 32(5):474–489.

- 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.
- Roberts, Margaret. 2018. *Censored: Distraction and Diversion Inside China’s Great Firewall*. Princeton University Press.
- Rosenblum, Nancy L and Russell Muirhead. 2019. *A Lot of People Are Saying: The New Conspiracism and the Assault on Democracy*. Princeton University Press.
- Rozenas, Arturas and Denis Stukal. 2019. “How autocrats manipulate economic news: Evidence from Russia’s state-controlled television.” *The Journal of Politics* 81(3):982–996.
- Ryan, Timothy J and Amanda R Aziz. 2021. “Is the political right more credulous? Experimental evidence against asymmetric motivations to believe false political information.” *The Journal of Politics* 83(3):1168–1172.
- Schon, Justin. 2020. “How narratives and evidence influence rumor belief in conflict zones: Evidence from Syria.” *Perspectives on Politics* pp. 1–14.
- Seawright, Jason. 2016. *Multi-method Social Science: Combining Qualitative and Quantitative Tools*. Cambridge University Press.
- Soliman, Abu Bakr, Kareem Eissa and Samhaa R El-Beltagy. 2017. “Aravec: A Set of Arabic Word Embedding Models for Use in Arabic NLP.” *Procedia Computer Science* 117:256–265.
- Stockmann, Daniela. 2012. *Media Commercialization and Authoritarian Rule in China*. Cambridge University Press.
- Stockmann, Daniela and Mary E. Gallagher. 2011. “Remote Control: How the Media Sustain Authoritarian Rule in China.” *Comparative Political Studies* 44(4):436–467.
- Svolik, Milan W. 2013. *The Politics of Authoritarian Rule*. Cambridge University Press.
- Wang, Chengli and Haifeng Huang. 2021. “When “fake news” becomes real: The consequences of false government denials in an authoritarian country.” *Comparative Political Studies* 54(5):753–778.
- 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., 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.