Regime Comments as Ballot Stuffing: How Astroturfing Can Deter Protest without Persuasion or Distraction

Michael Masterson*

30 August, 2025

Abstract

I argue that citizens draw conclusions about public opinion from online comments by treating these comments like votes in a poll. This enables governments to ballot stuff by using astroturfing to shape public opinion perceptions. Regimes can deter protests by convincing dissidents that they are alone in their dissatisfaction. This works without requiring astroturfing to persuade citizens to improve their own opinions of the regime, which are often difficult to change. I derive a set of novel empirical implications to distinguish support for my ballot stuffing interpretation from existing accounts that governments intend astroturf posts to persuade or distract. I find support for the ballot stuffing interpretation using both computational and qualitative text analysis methods to examine these implications on a leaked dataset of Chinese astroturf posts. I further conduct an experiment with Chinese respondents that shows astroturfing shapes popularity perceptions.

11,066 Words

One of the most important considerations before taking political action, especially in a risky context, such as protesting against an authoritarian regime, is whether or not citizens

*Assistant Professor at Missouri State University mmasterson@missouristate.edu.

believe that others will join them (Kuran 1991; Wedeen 2015, 151–52; Huang and Cruz 2022). If a citizen thinks many of their fellows will join them, making their movement more likely to succeed, they are more likely to take action. Not only are citizens' beliefs about the regime preferences of other citizens important, but citizens' beliefs about whether other citizens think the regime is popular are also critical to mobilization. If I think that the regime is unpopular but believe that most of my compatriots have incorrectly concluded that the regime is popular, I may assess that they will not protest because they believe others will not turn out, leading to my own decision not to participate (Little 2017). In this way, citizen beliefs about other citizens' preferences and beliefs are at least as important a determinate of collective action as actual level of the regimes popularity, and shaping these beliefs provides an important tool for the regime to control collective action.

I argue that authoritarian regimes, attempt to manipulate citizen beliefs about their fellows' preferences and beliefs in the most direct way possible, by masquerading as ordinary citizens expressing pro-regime preferences. While political scientists have acknowledged the potential of astroturfing to shape perceptions of public opinion, they have tended to view this as only a means to the end (an "underlying tactic") of making astroturfing more effective at persuasion or distraction (Zerback and Töpfl 2022, 401). In contrast, I argue that shaping perceptions of public opinion is the primary strategic goal of astroturfing. This has a direct effect on citizens' propensity to engage in collective action without requiring successful persuasion.

The ability of astroturfing to shape perceptions of public opinion has long been recognized and practiced in marketing. If you want a brand to appear popular, you can post fake comments praising it. Decreasing the probability of anti-regime collective action by shaping citizens' beliefs is a straightforward application of this. If citizens' take these astroturf comments to be genuine expressions of their fellows, they will update their beliefs about how popular the regime is. Research has found that readers update their assessment of public opinion based on internet comments, even though they know these comments are

not representative and are subject to manipulation (Lee 2012, 39, 42–43; Lee and Jang 2010; Zerback and Töpfl 2022). A simple post reading 'approve' on a news article about a new government policy will signal to readers that at least one person feels positively about the regime's policies. In this way, readers who tally posts to assess the weight of online opinion implicitly treat them like votes in a popularity poll. This allows malign actors to use astroturfing to engage in "virtual ballot-stuffing" to shape popular perceptions of public opinion (Walther and Jang 2012).

Further, citizens exposed to astroturfing also know that other citizens read these pro-regime comments, so they will also update their beliefs about whether or not other citizens believe the regime is popular (their secondary beliefs). Only for the first mechanism to function do citizens need to be tricked into thinking the comments are genuine. Astroturfing will affect the secondary beliefs of even citizens who can see through it as long as they believe that at least some of their fellow citizens are fooled (Little 2017, 225). This can deter even incredulous citizens from protesting because they think they are less likely to be joined by their credulous fellow citizens.

Critically, these mechanisms do not require that astroturfing persuades citizens to improve their own opinion of the regime. This matters because regime attitudes are often entrenched and difficult to change (J. Pan, Shao, and Xu 2020, 1–2; Mattingly and Yao 2022, 2–3).

Because astroturf posts rarely contain reasoned arguments, previous research has seen them primarily as a means of distracting citizens from politics (King, Pan, and Roberts 2017; Sobolev 2019). However, for these posts to serve the ballot stuffing function that I propose, they need not contain meaningful arguments, they only need to express a positive opinion. This function does not easily fit into the categories of soft propaganda that seeks to persuade with reasoned argumentation or hard propaganda that seeks to signal the regime's strength by compelling prominent actors to engage in outlandish speech (Huang 2015, 2018; Mattingly and Yao 2022).

I theorize a set of novel empirical implications that can distinguish support for my ballot stuffing interpretation of astroturfing in comparison with existing distraction and persuasion interpretations. To test these implications, I analyze a dataset of known astroturf posts, which leaked from the Internet Propaganda Office of Zhanggong district in Ganzhou city, Jiangxi province, China (King, Pan, and Roberts 2017). I find that the vast majority of these posts, 79.25%, express positivity about the regime or its priorities including: the Chinese Communist Party, government policies, government officials, the state of the economy, and the general state of Chinese society. My measure likely underestimates the amount of such posts because the dataset lacks information about what the posts are responding to and I only include posts that explicitly mention a pro-government topic, meaning a post responding to a government announcement that simply read, "approve!" (赞!) would not count. I also engage in a qualitative reading of these posts and show that their content is more consistent with a ballot stuffing interpretation than alternative understandings of astroturfing.

While this analysis of leaked post is consistent with my theory, it does not reveal whether astroturf posts manipulate citizen beliefs and secondary beliefs about popularity as my theory predicts. To examine support for this mechanism, I conduct a survey experiment with Chinese citizen respondents who read a short description of a state television program and are randomly assigned either to see an astroturf post promoting that program or a control. Respondents are then asked how popular they believe the show is and how popular they believe the average Chinese citizen thinks the show is to measure astroturfing's effect on beliefs and secondary beliefs respectively.

In the following section, I elaborate my ballot stuffing theory of astroturfing, place it in the larger context of research on propaganda, and derive empirical expectations to test it against rival explanations. Next, I explain the dataset I used to test these implications and how I coded it. Third, I analyze the data. Forth, I explain and analyse the survey experiment, and finally I offer concluding comments.

1 The Purposes of Propaganda

Astroturfing is a kind of propaganda and should be analyzed in that the larger context (Han 2015, 106). Authoritarian governments typically use propaganda to persuade citizens of their legitimacy, shape citizens preferences (such as by increasing nationalism or loyalty), or increase their dominance of society by signaling the regime's power (Rosenfeld and Wallace 2024, 265).

The exposure-acceptance model in which the primary purpose of propaganda is persuade citizens with pro-government messages is a classic way of thinking about propaganda (Geddes and Zaller 1989). In this view, the goal of propaganda is to persuade citizens to change their preferences to be more supportive of the regime and its policies (Geddes and Zaller 1989, 334). For persuasive propaganda to succeed, it must make logical arguments and appear credible to citizens (Rosenfeld and Wallace 2024, 266). This can be challenging as findings on whether citizens are persuaded by propaganda are mixed (Rosenfeld and Wallace 2024, 268), and, even studies that find propaganda has persuasive effects often find it is more effective at persuading citizens to support specific policies than to change their overall approval of the regime (J. Pan, Shao, and Xu 2020, 1–2; Mattingly and Yao 2022, 2–3). This suggests that citizen preferences about the regime may be entrenched and difficult to manipulate, highlighting the importance of the regime's ability to use propaganda to dissuade challenges without persuading those who disapprove of it to change their minds.

One way that regime may deter challenges is to signal its power by compelling others to engage in "transparently phony" speech and rituals that go unchallenged (Wedeen 2015, 6, 12; Huang 2015). The distinction between propaganda intended to persuade and propaganda intended not to persuade but to signal the regime's power has come to be refereed to as the difference between 'soft' and 'hard' propaganda (Huang 2015, 2018).

However, the dominant understanding of regime astroturfing is that it is primary to distract from other political content online rather serve as either hard or soft propaganda (King, Pan, and Roberts 2017; Sobolev 2019). This conclusion is based on findings that

astroturfing comments rarely engage in the kind of argumentation we would expect if they were soft propaganda intended to persuaded (King, Pan, and Roberts 2017, 492). Further casting doubt that astroturfing is soft propaganda, multiple studies have found that such posts do not successfully shift online conversations in a more pro-government direction or promote favorable attitudes towards the regime or its policies (Sobolev 2019, 4; Wong and Liang 2023, 1072). For these reasons, even though past research has found most government astroturf posts can be considered "cheerleading" in that they include "positive discussions of valence issues," this research has concluded that "the strategic objective of the regime is to distract and redirect public attention from discussions or events with collective action potential because these posts lack argumentation and tend to avoid controversial topics" and that "cheerleading in directed 50c bursts is one way the government distracts the public (King, Pan, and Roberts 2017, 485, 490, 496)." Further, since the entire purpose of astroturf comments is to appear as genuine grassroots opinion expression (King, Pan, and Roberts 2017, 484; Miller 2016, 2; Sobolev 2019, 7–8), it is unclear how they could serve hard propaganda purposes by showing the government compelling absurd speech. For these reasons, the most prominent study on astroturfing has concluded that "the purpose of 50c activity is to (1) to stop arguments (for which distraction is a more effective than counterarguments) and (2) to divert public attention from actual or potential collective action on the ground (King, Pan, and Roberts 2017, 495–96)."

Without denying that astroturfing posts may sometime distract, I argue that these posts primarily serve a different purpose, namely decreasing the probability of challenges to the regime by changing perceptions of public opinion. Citizens are more likely to believe they will face punishment for challenging the regime and are hence less likely to do so if they believe others support the regime (Kuran 1991). Even if individual citizens' beliefs about public opinion are not directly influenced by astroturfing, they may believe other citizens are influenced by it, which may deter protests. This is known as the "influence of presumed influence" (Huang and Cruz 2022). This is enhanced through the "third person effect" in

which citizens believe others are more influenced by media content than they are themselves (Houston, Hansen, and Nisbett 2011). Studies have found support for this idea by showing that propaganda has a larger influence on citizens' beliefs about other citizens' regime support than on citizens' own regime support and that citizens who believe their fellows support the regime are less likely to engage in anti-regime protests (Huang and Cruz 2022, 1791).

1.1 Astroturfing as Ballot Stuffing

The role of astroturfing becomes clear by focusing on the aspect that sets it apart from all other forms of propaganda, namely that it disguises itself as comments from ordinary citizens (King, Pan, and Roberts 2017, 484; Miller 2016, 2; Sobolev 2019, 7–8). Unlike other forms of media, comments, which ostensibly express the opinions of other citizens, allow readers to make direct inferences about popular opinion (Lee and Jang 2010, 829; Walther and Jang 2012, 3, 9). Consistent with this, research has found that when comments are present on a news article, readers are more likely to use those comments rather than information in the article itself to draw conclusions about public opinion (Lee and Jang 2010, 844). In this case, citizens directly update their beliefs about other citizens' opinions based on online comments (Zerback and Töpfl 2022). For this reason, manipulating these comments is a powerful tool in the regime's arsenal to manipulate citizens' propensity to challenge the regime by manipulating public opinion perceptions. This differs from the soft propaganda view that astroturfing is designed ot increase regime legitimacy (Han 2015, 106). In the ballot stuffing account, astroturfers are more focused on created the appearance of regime support than persuading citizens to support the regime.

It is easy to imagine how this could work. We have all faced the temptation, after reading a tweet or a news article, to scroll down and examine the comments to see how others are reacting. Consistent with this, research suggests readers use comments to draw conclusions about public opinion even when they know these comments are not representative, and readers may still regard online comments as more representative than other information they

have, such as opinions within their social circle (Lee 2012, 42–43; Zerback and Töpfl 2022). One possible reason for this is the availability heuristic. People tend to regard phenomenon that are more easily mentally available, such as those that they are frequently exposed to, as more common (Lee and Jang 2010, 831). Astroturf comments provide "exemplars" that citizens take to be representative of a larger population (Zerback and Töpfl 2022, 402).

Of course, astroturfing would fail to influence citizens' estimates of public opinion if citizens are highly critical of the content they consume online and resistant to this manipulation. However, research suggests that this is not the case. An experimental study of astroturfing found that warning citizens that astroturfing was taking place or about the general strategies used by astroturfers did not reduce the effect of astroturf comments on their perception of public opinion (Zerback and Töpfl 2022, 408–9), suggesting that even in context like China in which users may know astroturfing takes place, it may still affect their public opinion perceptions. In general, citizens tend to use heuristics to evaluate the credibility of online information and rarely verify it (Metzger, Flanagin, and Medders 2010, 416, 425). Resisting the influence of peer exemplars on public opinion perceptions is particularly difficult because many citizens may not consciously realize they are influenced by it. In an experiment in which respondents were shown a manipulated summary of audience responses to key moments in a presidential debate, respondents self-reported that they were not influenced by this information, despite being influenced by it (Fein, Goethals, and Kugler 2007, 182). In one version of the experiment, respondents were informed that the information they received had been manipulated, but this known false information still affected their estimate of what the average national response to the debate was (Fein, Goethals, and Kugler 2007, 185–86). This further suggests that even though citizens know the government astroturfs, this knowledge is unlikely to prevent astroturfing from affecting their estimates of public opinion. In another experiment that examined the effect of online comments on the credibility of Chinese government social media posts, respondents did not respond differently when shown that the comments they were seeing were "recommended comments by account

holder," which meant that the government account curated what comments were shown (Chen and Lu 2024, 12, 15). This further suggests that knowledge of manipulation does not prevent this manipulation from having its intended effect on social media.

Even highly sophisticated citizens who resist the temptation to update the beliefs on the basis of astroturf comments may be deterred from protesting. Just as with normal propaganda, as long as citizens think that at least some of their fellow citizens are deceived, it will affect their behavior on political activities that require coordination such as protest (Little 2017, 225–26, 230). In this way, while governments might prefer to successfully manipulate public opinion perception directly, manipulating secondary beliefs (i.e. beliefs about others' beliefs) provides a fallback mechanism to prevent protests.

So far I have discussed the purpose of propaganda from the top-down perspective of the government, but the astroturfers serving as the state's agents might have their own goals. For example, these workers want to please their bosses by making many posts without writing content that could get them in trouble. They can acheive this by pumping out a large volume of low quality posts that echo acceptable talking points. Previous research has highlighted this as a principal-agent problem that could undermine the persuasiveness of astroturfing (Han 2015, 122–23). However, if the primary purpose of astroturfing is to ballot stuff rather than persuade, the incentive to create a large volume of low quality posts that can shape beliefs about opinion online but lack persuading power is consistent with state goals and may explain why astroturfers do no fear to report the content of their low quality posts to their supervisors (each post in the leaked dataset comes from an email from the astroturfer reporting their post up the chain) and why the state has not adjusted their incentives to foster higher-quality persuasive content.

1.2 Empirical Implications

Below I outline a series of empirical implications that set the ballot stuffing account apart from alternate accounts, including the distraction and persuasion accounts (Han 2015; King,

Pan, and Roberts 2017). As the effect of astroturfing on public opinion has previously been viewed as a means of persuasion or distraction rather than and end in itself (Zerback and Töpfl 2022, 401), this is the first attempt that I am aware of to derive theoretical expectations that would distinguish these motivations for astroturfing. Because the distraction account focuses more on theorizing why the regime might want to use astroturfing posts to distract from collective action than deriving theoretical expectations about the kinds of posts that might be most distracting (King, Pan, and Roberts 2017), I draw on research from psychology to derive novel implications about what we would expect if the regime designs posts to be maximally distracting. However, I recognize that distraction theorists may have a different idea about the likely content of distracting posts, so I make the test especially difficult for my own ballot stuffing explanation by counting any posts that do not fit the ballot stuffing account's expectations as distracting.

While my indicators can distinguish whether the evidence is more consistent with one account than another, they are not causal tests of why posts of are generated. The nature of astroturfing means that researchers must rely on leaks for "ground truth" posts that come from known regime commentators (Miller 2016, 3; King, Pan, and Roberts 2017, 487). Without such a ground truth, researchers may bias their results by assuming posts that fit their theory of astroturfing are indeed astroturf posts, leading to a self-fulfilling examination. However, ground truth datasets also come with limitations that pose problems for causal identification, such as a lack of clarity of how the leaked posts were selected from the universe of astroturf posts. Still the importance of the topic necessitates an attempt to understand these comments. To increase confidence, I have supplemented my empirical analysis below by summarizing the findings of astroturfing research across country contexts on each of the indicators I produce.

Further, while these indicators can show whether the astroturfing posts we observe are more consistent with my ballot stuffing account than alternatives, they do not directly test the assumption that astroturfing posts affect citizen views of public opinion. I choose to focus

my empirical efforts on the posts themselves rather than this assumption primarily because we already have a wealth of research, described in the previous section, that shows that online comments and other exemplars of peer opinion affect citizen views of public opinion, even when citizens know manipulation may be involved (Fein, Goethals, and Kugler 2007; Lee and Jang 2010; Lee 2012; Walther and Jang 2012; Zerback and Töpfl 2022; Chen and Lu 2024). We also already have research that indicates public opinion perceptions are key to collective action (Kuran 1991; Huang and Cruz 2022). For these reasons, the key empirical task is to examine whether the content of astroturf posts are consistent with this goal.

The most basic empirical expectation of the ballot stuffing perspective is that astroturfing posts should express positive views of the government, its officials, and its priorities to create the impression that these are popular. This is distinct from the distraction perspective that expects such posts are "to distract from general negativity, government-related meetings and events with protest potential (King, Pan, and Roberts 2017, 496)." If posts are intended to distract from government activities, they should avoid mentioning these activities. Further, posts with negative affect rather than positive affect might more effectively distract because negative stimuli are better at attracting attention due to a phenomenon know in psychology as negativity bias (Rozin and Royzman 2001; Soroka, Fournier, and Nir 2019). Such distraction posts should further be sensational, surprising, and unpredictable as these are the features most conductive to producing distraction and capturing attention (Davis and McLeod 2003; Parmentier 2014). If the government wants to distract from a discussion of China's economic problems, it might be more effective for astroturfers to bring up a salacious celebrity scandal rather than post positive affirmations about the economy (Roberts 2018, 190–91). Whataboutism is tool that combines sensationalism and negativity that can successful distract from and mitigate the impact of criticism of authoritarian regimes (Bowell 2023; Chow and Levin 2024), so if astroturf posts serve this purpose, we might expect them to use whataboutism by bringing up other countries' failings. In contrast, sensationally negative off-topic posts would not serve a ballot stuffing purpose. To ballot-stuff, posts need not

be sensational, they only need to express a positive opinion about the regime or its priorities. Past research finding that astroturf posts in China tend to have positive emotional valence and rarely contain negative content about other governments or nations is consistent with this expectation (King, Pan, and Roberts 2017, 491).

While the persuasion interpretation of astroturfing would also expect positivity, these approaches have different expectations for the nature of these positive posts. If astroturfing is ballot stuffing, then governments should place a greater emphasis on the quantity than quality of astroturf posts. This is because, as long as a post expresses an opinion, readers will tally it as a vote in favor of that opinion, regardless of whether it makes a compelling argument. The time consumed to construct such an argument might even be counter productive as increasing the time required per post will reduce the total volume of posts. Consistent with the claim that quantity matters for ballot stuffing, researcher have found that the effect of astroturf comments on public opinion perception increases with repeated exposure (Zerback and Töpfl 2022, 402). In contrast, propaganda that is intended to persuade tends to be higher effort and involve high production value media (Mattingly and Yao 2022, 2). Even in the context of internet comments, research has found citizens are more persuaded by high rather than low quality comments, suggesting the government would emphasize comment quality if comments are intended to persuade (Y. Pan and Shu 2022, 2). Indicators of low quality posts include: little to no effort to provide evidence, little to no effort to develop arguments, and repeating or reusing previous comments rather than developing novel comments. The distraction perspective would not expect such repetition as surprise and unpredictability are more effective at producing distraction and capturing attention (Davis and McLeod 2003; Parmentier 2014)

If astroturfing plays a ballot stuffing role, there is no need to specifically engage with and refute comments critical of the regime as simply expressing a pro-government opinion counts as a 'vote' in favor of the government regardless of whether regime criticism is present. Engaging criticism would require astroturfers to devote more time to reading other online

posts and make it more difficult for them to focus on quantity by spamming similar pro-regime content regardless of context. In contrast, the persuasion interpretations predicts that astroturfers should engage and refute regime criticism (Zhang, DiFranzo, and Hendler 2014, 1889).

Indeed, past research has noted that the incentives of astroturfers in China emphasize the production of a great quantity of posts rather than of quality posts that would be more effective in persuasion (Han 2015, 118). Further, research has found that Chinese astroturf posts almost never make arguments (King, Pan, and Roberts 2017, 485), and a leaked dataset of astroturf post from the Thai military found that posts expressed general positivity without providing arguments (Sirikupt 2024, 11). Also consistent with ballot stuffing's emphasis on quantity rather than quality, past research has noted that unlike ordinary citizens, astroturf comments often repeat or retweet each other (Han 2015, 121; Sirikupt 2024, 12–13).

Similarly, astroturfers should ballot-stuff using multiple accounts to create the appearance that more individuals support the regime. While using multiple accounts could be compatible with distraction or persuasion, its necessity to these purposes is less clear. For persuasion, using a single account that cultivates credibility and relationships with its readers might be more effective as persuasion works best when sources of information are few and authoritative and groups are tight-knit (Johnston 2001, 509–10). Empirically astroturfers do use multiple accounts, up to 50 per person (Miller 2016, 12–13). This has been observed across context, including in China, Russia, and South Korea (Miller 2016; Sanovich, Stukal, and Tucker 2018; Keller et al. 2017, 566–67). The most extreme case of this is the use of bots and cyborgs (humans aided by scripting tools) to mass produce posts by the Russian government, which one study found responsible for "by far the largest volume of content" among sources of Russian astroturfing (Sanovich, Stukal, and Tucker 2018, 445–46).

If astroturf posts are intended to ballot stuff, then they must plausibly represent the opinion of genuine citizens to affect beliefs about public opinion. For this reason, astroturf accounts should lose their utility and quickly be abandoned if unmasked. For persuasion and

distraction, this is not necessarily the case. If the goal is to attract attention and distract, a post coming from the government rather than an anonymous online commentator is more likely to attract attention. However, as we might imagine motivations for the government to maintain the anonymity of astroturfing accounts for persuasion or distraction, this implication has more utility in its potential to falsify the ballot stuffing interpretation than to confirm it.

Consistent with the expectation that astroturfing accounts lose their utility when unmasked, the abandonment of these accounts following unmasking has been observed across country contexts. In Russia, when 700 astroturf accounts were leaked, most of the accounts stopped posting within a day (Sobolev 2019, 14). In South Korea, astroturfers quickly stopped activity after the National Intelligence Service's efforts to astroturf in favor of presidential candidate Park Geun-hye leaked (Keller et al. 2017, 565)

If regimes use astroturfing to ballot stuff, we would expect them to use these posts to shape beliefs about popular opinion on issues of priority to the regime. One way of measuring a regime's priorities is to examine how much of its overt propaganda it devotes to a particular issue. For the CCP, nationalism is one of the most important and tightly scripted propaganda issues (Waight et al. 2025, 6). Therefore, we would expect astroturfing posts to promote the governments' priorities on nationalism. Further, the posts should be composed of positive pro-government nationalism to create the impression that more citizens hold this positive view about the government's ideology. There is less reason from a ballot stuffing perspective for astroturfing posts to express negative sentiment towards other nations. This expectation distinguishes the ballot stuffing account from the distraction account, in which astroturfing posts, designed to distract rather than shape opinion or beliefs, should be no more likely to discuss government priorities than the posts of typical citizens, but might use negative nationalists sentiment and whataboutism directed towards other countries.

If control over astroturfing posts is decentralized, as it is in China (Han 2015, 113), the ballot stuffing perspective would expect that local officials with control over astroturfing

posts would use them to promote themselves and their priorities in the same way that the national government does. This could help them deter protests over local issues, giving them a boost in their promotion evaluations that emphasize maintaining stability (J. Wang 2015, 18). In contrast, the distraction perspective would expect that these officials wish to direct attention away from discussion of the local government and its officials and, therefore, should not mention them.

One empirical indicator examined in past literature that cannot differentiate the ballot stuffing explanation from alternatives is whether posts occur in bursts. Past accounts of astroturfing, while cautioning that the reason astroturf posts occur in bursts is unclear, have argued that such bursts are consistent with a distraction account (King, Pan, and Roberts 2017). However, these bursts could exist for a variety of reasons, such as an attempt to influence what topics are trending (Keller et al. 2017, 566), or could just be an artifact of the state-directed coordination of posts. Such bursts have also been observed in traditional media propaganda that is not theorized to have this distraction purpose (Waight et al. 2025, 3). Bursts could also be useful for ballot stuffing purposes because the more temporally proximate the posts are, the more they will affect readers' public opinion impression at a given time. For these reasons, whether or not posts occur in posts cannot distinguish these accounts.

Table 1: Summary of Empirical Expectations for Posts

Theory	Expectations
Ballot stuffing	Express positive affect towards the regime or its priorities
	Emphasize quantity over quality of posts
	Reuse posts to inflate numbers
	Astroturfers use multiple accounts
	Astroturf accounts abandoned when unmasked
	Self-directed nationalism
Distraction	Use sensationalism and negativity to attract attention
	Use surprise and novelty to attract attention
	Whataboutism and other-directed nationalism
	Do not mention regime, its actions, or officials
Persuasion	High-quality pro-government arguments
	Engage and refute criticism of regime

2 Data and Coding

I analyze a dataset of Chinese astroturf posts extracted from leaked emails from the Internet Propaganda Office of Zhanggong in Ganzhou City of Jiangxi Province, China (King, Pan, and Roberts 2017, 485). While this dataset is likely non-representative of regime-commentator posts in some ways, for example it contains more posts about local issues such as the Jiangxi Soviet than I would expect from a nationally representative sample, this is an inevitable trade-off in a world of limited known regime commentator datasets and no way to identify regime commentator posts with certainty in the wild (Miller 2016). This dataset does have the advantage of being the dataset that the distraction interpretation was developed on (King, Pan, and Roberts 2017), so if the ballot stuffing interpretation fairs better on this dataset, it likely better explains state astroturfing in other contexts as well.

The dataset contains 42,828 posts with non-missing content from 01-15-2013 to 11-21-2014. I classify these posts into four categories. I give a brief explanation of each category below, and the appendix contains a complete coding dictionary with examples of each type of post. These categories are neither mutually exclusive nor inclusive, so it is possible for a

post to fit any combination of them.

The first and most important category is `positive affect`. These posts express positive sentiment about a policy, wider policy orientation, or the state of Chinese society. This includes expressing general positive sentiment towards one or more government policies, the Chinese Communist Party (CCP), one or more government officials, the state of modern society, or the economy. Posts that just express general positive sentiment without regard to one of these topics do not count. This differs from King, Pan, and Roberts (2017) 's category of "cheerleading for China" as positive affect is more focused on expressions that would lead readers to make inferences about the writer's opinion of the government or its priorities (499). For example, cheerleading excludes posts that praise government officials and policies, but positive affect includes these posts. Further, cheerleading includes general posts about thankfulness, Chinese culture, or inspirational slogans that positive affect does not unless the post invites the reader to credit the regime for this positivity.

Because ballot stuffing posts should express support for positive nationalist posts rather than negative posts directed at other nations, I create a category to measure each of these. First, self-directed nationalist (`self-nationalism`) posts emphasize national dominance, superiority, loyalty, greatness, or pride as well as a positive identity tie to the nation-state (Ko 2023, 24–25). In contrast, other-directed nationalism (`other-nationalism`) tends to be negative in affective valence and expresses hatred, hostility, or contempt for other nations as well as the idea that other nations are evil (Ko 2023, 25).

Since the ballot stuffing account expects local leaders to take advantage of ballot stuffing to increase the perception that they and their priorities are popular when control over astroturfing is decentralized, I create a `local` category that includes posts that mention events, issues, or individuals that are local to Zhanggong district in Ganzhou City, or Jiangxi Province. If such posts exist in significant number, and the vast majority of such posts are also positive affect posts, this would support the ballot stuffing account.

2.1 Classification Process

First, I trained a native Chinese speaking research assistant on my coding dictionary. In addition to containing definitions and examples of each category, the coding dictionary also instructs coders to use information about the time that the post was created to interpret its content, using King, Pan, and Roberts (2017)'s helpful description of relevant events that occurred during the time period (498). My research assistant and I then iteratively went through several rounds of coding 200 posts each and then coming back and discussing the ones we coded differently. In total, 600 posts were used in this training process, and an additional 200 posts were used to evaluate intercoder reliability once training for a category was complete. This enabled us to achieve an intercoder agreement of at least 90% for each category.[1]

Following this training process, I asked my research assistant to code 3,000 posts according to the codebook. I then used these posts to train a supervised machine learning model to code the remaining posts. Based on the advice in Chang and Masterson (2020), I opt to use a support vector machine model with the loss-function adjusted to penalize misclassification of the rarer class more heavily to classify the posts. I trained a separate model for each category and evaluated the model based on an 80%/20% train-test split. The models performed well not just on overall accuracy but also on precision and recall for each category.[2]

3 Analysis

Supporting the ballot stuffing implication that posts emphasize quantity rather than the quality that the persuasion account expects or the surprise and novelty that the distraction

[1] Krippendorf's α of 0.65 for positive affect, 0.8 for self-nationalism, and 0.89 for local. Krippendorf's α was undefined for the test set for other-nationalism because one of the coders labeled 0 of the 200 posts other nationalism and the other labeled 1 of them.

[2] The positive affect model scored 86.73% accuracy, 93.07% precision, and 90.78% recall. The self-nationalism model scored 90.82% accuracy, 91.15% precision, and 88.43% recall. The other-nationalism model scored 99.83% accuracy, 100% precision, and 75% recall. The local model scored 96.77% accuracy, 94.53% precision, and 90.98% recall.

account expects, nearly half of all posts (21,287 of 42,828) are word for word repetitions of other posts in the dataset. This is without even considering the posts that are slight permutations of others posts, of which there are many. The following subsections examine empirical implications about the content of posts for each category of posts.

3.1 Astroturfing Regime Positivity

Only a very low proportion (less than 5%) of normal Chinese social media posts discuss politics (Chang and Masterson 2020, 402), so the finding that 79.25% of astroturfing posts in the dataset are positive affect posts is striking. This is even more than case because only posts that explicitly mention a topic that reflects on the regime, such as the quality of Chinese society or its economy, the CCP, or Chinese government officials, counts as positive affect. This measure likely underestimates the number of positive affect posts because the dataset contains many posts with content such as "approve of this" (赞一个) that are not coded here as positive affect posts even though in context what is being approved is probably some sort of government activity or priority as many of these posts respond to overt government accounts (King, Pan, and Roberts 2017, 492).

One type of positive affect post expresses general positivity towards the party or the regime. 3,649 posts mention the Party (党), and 3,585 of these are positive affect posts. One such posts reads, "Let us be grateful for our blessings, thank the party, and repay the party and the country for their care." Another writes:

> Since the establishment of new China, the party and government have spared no effort to improve people's lives in Ganzhou, last year even releasing the Some Comments document to work towards making the people live prosperous lives. We are grateful to the party. We ardently love the party. We will definitely seriously study science and culture and make a contribution to the motherland and the people.

Figure 1: Post Frequency

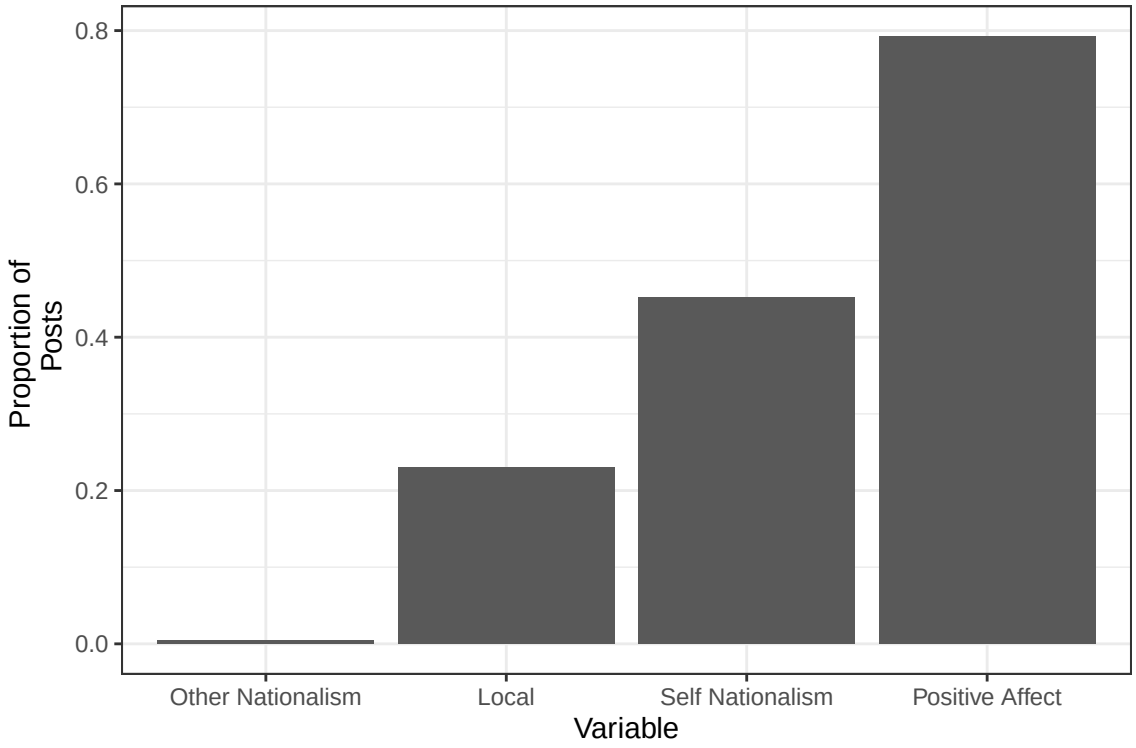

Figure 2: Posts Over Time

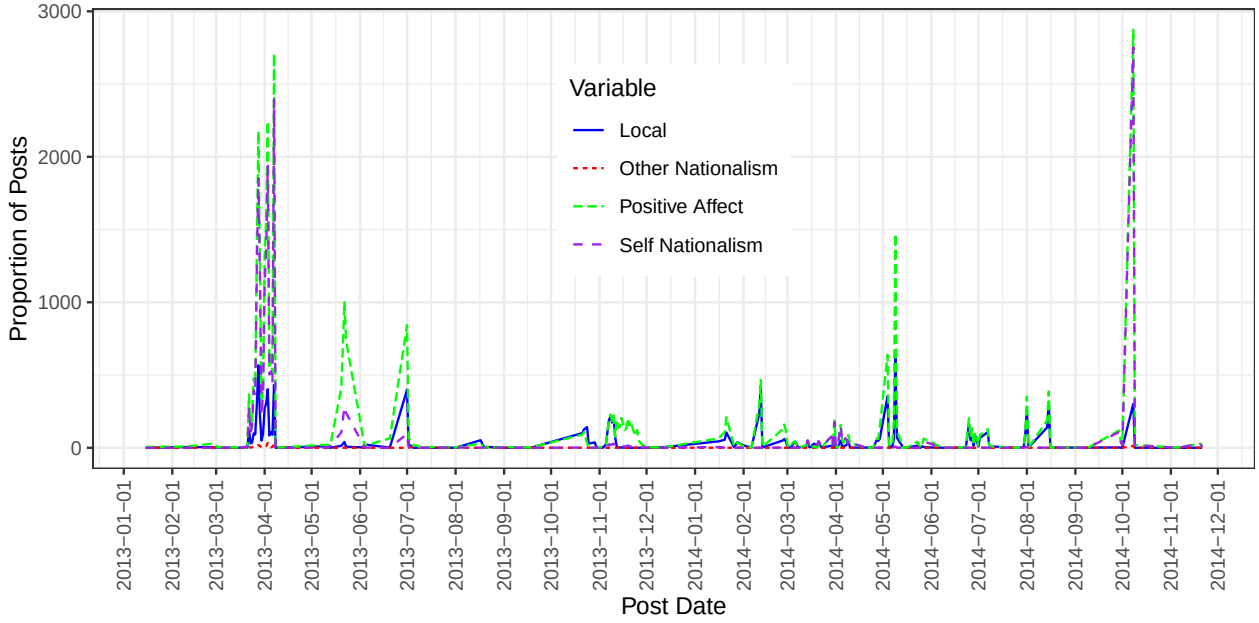

Expressing thanks to the party or the government is a common theme of such posts. For example, one says, "Thank the government for its great efforts to support and help those in need." While these posts are typical in that they do not present detailed and evidence-backed arguments in favor of the party as the persuasion interpretation would expect, they are very clear in expressing a positive opinion about the party and the regime. Illustrative of this opinion expression without argumentation, one post reads, "Forever go with the party. Forever love the party." Further, these posts are not attempts to change the subject to the kinds of salacious but nonpolitical topics that would be most effective at creating distraction.

Another group of positive affect posts explicitly praises CCP meetings. For example, 993 posts mention Chinese Communist Party plenums (全会) and of those, 962 are positive affect posts. A typical such post states that everyone has "energetically organized study of the spirit of the Third Plenum of the Eighteenth Party Congress" and made "its content enter the people's hearts (入心入脑), making their thought and action united with the decisions and dispositions of the CCP Central Committee." Consistent with the ballot stuffing account, this post not only praises government meetings and priorities but tries to give the impression that this praise is widely shared among citizens by saying that "every town and street" has unified with the spirit of the CCP Central Committee. This is typical of the posts that mention party plenums. For example, another states, "The plenum emphasized the success of Reform and Opening at comprehensively deepening reform and providing important experience and necessary long-term support. Most important is persisting with the party's leadership, implementing the party line[…], avoiding the evil path (邪路) of changing flags, and persisting on the road of socialism with Chinese characteristics." While past research has noted bursts of astroturfing activity during party meetings that mentions these meetings, it has argued that these bursts are "consistent with a strategy of distraction… [because] political meetings are periods when government and party officials believe that protests are more likely to take place. During these periods, officials gather and attention is focused on the activities of the regime; as such, successful protests can garner greater attention (King,

Pan, and Roberts 2017, 489)." However, if the desire is to take attention away from the activities of the regime, it would make more sense for astroturfers to attempt to change the subject than to praise these meetings. The content of these posts that not only mention the meetings but do so in a positive context is more consistent with the ballot stuffing rather than the distraction interpretation.

Posts also praise CCP officials. 327 posts mention Xi Jinping and 322 are positive affect posts. One such posts contains a quote from Xi and concludes "Support Secretary Xi." This post occurs word for word in the dataset 3 times. Some party members, seen as models for citizens, are mentioned even more frequently. For example, Gong Quanzhen who joined the CCP in 1952 and is upheld as an inspirational figure by the party is mentioned in 759 posts. While the distraction interpretation might expect that these would be non-political posts that focus on Gong's work as an opera singer, 749 of these posts are positive affect posts that praise at least one regime-attributable aspect of China. The political content of such posts is hardly subtle, for example, one such post states, "Faith and spirit are like a flag that guides us forward and provides unlimited strength. Elder Gong Quanzhen's persistent faith and pure party spirit are also like a flag that moves the peoples' emotions and, gives them confidence and something to ponder." A reason Gong posts may outnumber Xi posts is Gong's connection to Jiangxi, where the dataset originates. It is probable that nationally, more astroturf posts reference Xi than Gong while each province may have a local figure that in that province is referenced more than Xi.

Another set of positive affect posts praise government priorities, concepts, or policies. For example, Xi Jinping began popularizing the concept of the "Chinese Dream" in a speech on November 29, 2012, which is shortly after he became General Secretary and about a month before the dataset begins (Xi 2012). 2,619 posts mention the Chinese Dream, and 2,600 of these posts are positive affect posts. Many such posts clearly attempt to create the impression of widespread support for this initiative. For example, one commentator writes, "The Chinese Dream is the dream of every Chinese person." Following this theme, another

claims, "The Chinese Dream stems from a true sense of the pulse of the vast majority of the people." These posts clearly promote rather than distract from Xi' priority of the Chinese Dream, yet they do so in a way that lacks the argumentation and evidence we would expect if they were intended to persuade. Instead they simply express support for the Chinese Dream and claim that many others support it to, exactly what the ballot stuffing interpretation would expect.

3.2 Astroturfing Nationalism

One clear Chinese government priority is to promote nationalism. Yet, despite the strong presence of other-directed nationalism in the CCP's narrative that emphasizes themes such as China's victimhood and humiliation at the hands of foreign powers (Gries 2005; Z. Wang 2012), the ballot stuffing perspective expects that astroturfing posts should be more likely to contain self-directed nationalism because, unlike other-directed nationalism, its positive valence allows these posts to create the impression of support for the regime. In contrast, other-directed nationalism might be effective if the regime intends to distract from its own problems because of its attention-grabbing negative valance and ability to include whataboutism (Rozin and Royzman 2001; Soroka, Fournier, and Nir 2019; Bowell 2023; Chow and Levin 2024). As show in Figure 1, self-nationalism posts are much more common than other-nationalism posts with the dataset containing a total of 19,386 self-nationalism posts and only 195 other-nationalism posts.

Of these 19,386 self-nationalism posts 18,152 of them are positive affect posts, suggesting self-nationalism posts do not merely promote nationalism but tend to do it in a way that makes clear the writer is also expressing support of the regime. Keeping with the pattern observed in the previous section, while these posts lack the kinds of arguments and evidence that would fit the persuasion interpretation, the political nature of these posts is hardly subtle. For example, one proclaims, "Today we are closer than anytime in history to the goal our forebears tirelessly pursued, the great rejuvenation of the Chinese nation." A common

theme of many of these posts are the praise of martyrs, which a reading of a great many of these posts as well as the context in which the term martyrs is typically used on the Chinese internet makes clear refers to both Chinese nationalist martyrs and CCP martyrs. One such post says, "Today's life did not come easily. Countless heroic martyrs gave their blood for it!" A further posts makes this connection explicit, praising "revolutionary martyrs" (革命先烈) and calling on citizens to stay the course set by the martyrs and to "realize the Chinese dream of reviving the nation" (民族复兴). Rather than present reasoned arguments and evidence, this posts claims that citizens should revive China "to comfort the souls of the revolutionary martyrs in heaven." Of course, because the revolution referred to is the revolution that established the People's Republic of China in 1949, praise of the revolution and its martyrs also suggests approval of the current regime. This is part of a wider strategy in Chinese propaganda to conflate pro-Chinese nationalism with pro-CCP sentiment (Zhao 1998, 289). The regime further takes advantage of special days that make this connection obvious, such as the surge of self-nationalist post around China's National Day on October 1st that celebrates the founding of the PRC (see Figure 1). Another time period in which the regime exploits this association to create a surge of self-national posts is around the Qingming festival that commemorates the dead in early April.

Even among the 195 other-nationalist posts, 188 also contain self-nationalism, and 165 of them are positive affect posts, suggesting that they are more part of an attempt to stoke pro-regime nationalism than to distract from domestic politics with whataboutism. Indeed, even in posts coded as other-nationalism because they discuss topics such as China's past humiliation, the nation that humiliated China is rarely specifically mentioned, placing the focus much more on venerating the Chinese martyrs of these conflicts than on demonizing the other side. In fact Japan, the most likely candidate to be demonized for past humiliation, and the United States, perhaps the most likely target for whataboutism, are only mentioned 111 and 24 times respectively in the dataset.

A typical other-nationalist post says, "Even as time flows and memories are indirect, the

blood and smoke of more than seventy years ago, the humiliation and tears of more than seventy years ago, the pain and bloody battles of more than seventy years ago, seem not to have dissipated from our beloved land." Another humiliation post that does not explicitly name the humiliator that was posted on 10/08/2014 (around National Day) says, "How many benevolent soldiers offered their heads and hot blood just to realize national independence, protect against humiliation, make the country rich and strong, and provide for the people's well-being?" Even the posts that explicitly mention Japan put more of a focus on the role of the CCP in fighting Japan than on Japanese atrocities. One such post says:

> Looking back at history, we cannot forget the first shot at Marco Polo Bridge. We cannot forget the first great victory at Pingxingguan. We even more cannot forget the first counterattack of the great battle of the Hundred Regiments Campaign [the CCP's most significant offensive against Japan in WWII]. From [the Manchurian Incident on] 9/18/1931 to the surrender of Japanese imperialism, How many Chinese sons and daughters shed their last drop of blood on this land? How many Chinese offered their wholehearted sincerity for this land? The fifteen-year-old Liu Hulan [a CCP member] raised her hand and swore an oath in front of the five-star red flag without any regrets; sister Jiang [referring to Jiang Zhuyun, a CCP member] and her war comrades with tears in their eyes embroidered the five-star red flag with wholehearted loyalty. In the Second Battle of Shanghai, the 29th army fought bravely in a desperate situation, the five heroes of Langya Mountain jumped off the cliff without hesitation, astonishing the world.

Even though this post contains some other nationalism in its reference to Japan, the focus is even more on self-nationalism in highlighting the heroism of model Chinese with particular attention to figures and symbols, such as the five-star red flag, associated with the CCP. Also noteworthy is that unlike other Chinese nationalist propaganda, while this post mentions Japan it does not dwell on, or even mention, Japanese atrocities (Z. Wang 2012).

In fact the Nanjing Massacre, perhaps the most well-known and commonly mentioned of these in regime propaganda, is only mentioned in a single post in the entire dataset. Further, rather than a graphic, emotion-provoking description of this event, the post simply mentions that the Standing Committee of the National People's Congress voted to make September 3rd Victory Day for the war and December 13 a memorial day for the victims of the Nanjing Massacre. This is consistent with the finding above that most other-nationalist posts are also self-nationalist and positive affect. Overall, the evidence suggests these nationalists posts are part of a strategy to spread positive sentiment about the regime rather than to distract with negative sentiment about other countries.

3.3 Astroturfing Local Government Positivity

912 posts reference the "Two Meetings" (两会) of the Ganzhou's People's Congress and Political Consultative Committee, which took place during the dataset, and 865 of these are positive affect posts. Far from trying to distract from government meetings, astroturf posts are promoting them. Typical posts express non-argumentative support of the local government and its priorities. For example, one such post states, "Support the Zhanggong District government's reform and innovation." Others express general confidence about the direction of their locality (and implicitly the government running it), such as "Tomorrow Ganzhou will get better and better." Other posts are more explicit in their praise of the Ganzhou city government. For example, one claims, "The Ganzhou government's work is more and more transparent. It has many new measures to promote development in people's lives. I hope the Zhanggong District will develop faster and faster and the happiness index will soar." Note that this post, just like the others, fails to provide detailed evidence, which might contribute to persuasion, suggesting it plays more of a ballot stuffing than a persuading role. Other posts may use self-nationalism with a local spin. One such post, supporting the Jiangxi governments priority to develop the Jiangxi Soviet area states, "revive the development of the Jiangxi Soviet Region! I will always ardently love the party and forever go with the party!

Revolutionary martyrs will never be forgotten!" Another similar post proclaims, "Carry forward the will of the martyrs and revitalize and develop southern Jiangxi."

These posts also praise local officials. For example, there are 279 posts about Shi Wenqing who was the CCP Secretary of Ganzhou during the dataset. Of those posts, 252 are positive affect posts. One such post says, "Take practical action and win the people's hearts. Gannan leader Shi Wenqing entered villages and homes to listen to opinions with a kind and sincere attitude, asking detailed and profound questions. One local proclaimed 'the good cadre from the Soviet have returned!' " Another praises Shi for "listening to the opinions of netizens" in order to "develop party and mass line education activities, raise the Soviet spirit, and revive Gannan district." That local officials who have influence over local astroturfing would use it to increase popular perceptions of their and their policies' popularity, helping them deter protests that could derail their record on social stability and chances for promotion, is consistent with the ballot stuffing account.

4 Survey Experiment

This survey experiment is designed to test whether astroturf posts manipulate beliefs about popularity and beliefs about beliefs about popularity (secondary beliefs) as the theory predicts. Because citizens in authoritarian countries that are known to astroturf might be more wary of astroturfing, and, therefore, less likely to update their beliefs about popularity on the basis of social media posts, it is important to conduct this survey experiment with Chinese respondents rather than respondents in a democratic country that might be more credulous due to less experience with astroturfing.

Unfortunately, the current political climate in China makes it impossible to ask directly about citizens' beliefs about regime popularity. Instead, I examine whether astroturf posts can alter citizen beliefs about the popularity of the state television show *Legal Report* (今日说法). This still allows an examination of whether astroturfing can alter citizen beliefs

about popularity. It also has the advantage that citizens know the government wishes to promote this show, so they might be wary of social media posts that promote this show in the same way they might be wary of posts that promote government priorities generally. My analysis of the leaked astroturf posts also shows that they often promote even relatively minor issues and priorities of the regime, such as discouraging people from using fireworks, so the astroturf posts used in the study to promote a specific the regime goal of increasing *Legal Report* viewership is consistent with the way real astroturf posts also promote smaller goals of the regime rather than always attempting to directly manipulate beliefs about the regime as a whole.

I choose the television show *Legal Report* for a few reasons. First, this show is broadcast nationally and well known enough that citizens have a sense of its popularity (*Legal Report* has been on air since 1999). Second, this show airs every day, making it salient to respondents in a way that seasonal shows that take breaks between seasons might not be. *Legal Report* also promotes the known government priority of rule of law (法治), which in a Chinese context has more of a meaning of rule by law than the liberal conception of rule of law in which the law constrains the state. The promotion of the regime priority of rule of law is consistent with the astroturf posts in the dataset of which 172 mention 'rule of law.' The show is very open about its goal of promoting rule of law. For example, one of the first sentences about the show in the short official description of it on China Central Television's (CCTV) website announces that it "upholds the process of recording China's progress in rule of law and emphasizes spreading legal knowledge, supervising law enforcement, promoting the enactment of laws, and serving the people" (CCTV 2025).

In the experiment, respondents read a brief description of the Chinese television show *Legal Report* that is presented as a Weibo post. This description says:

> Since 1999, CCTV has shown *Legal Report* (今日说法). This program covers legal cases in China and promotes rule of law (法治). This show currently broadcasts on CCTV-1 everyday at 12:25 PM.

Respondents are randomly assigned to three treatment conditions that are shown as the top reply to the original post. In the first control condition, no reply is shown. In the spam control condition, citizens see a reply that reads, "Please visit my online store for a good deal!" This condition is included to control for the possibility that any reply affects beliefs about popularity. In the astroturf treatment conditions, the reply reads, "[I] approve of this show!" (赞这个节目！).

Respondents are then asked three outcome questions. The first asks them what percentage of Chinese citizens they think watch *Legal Report* at least once a week to measure astroturfing's effect on beliefs about popularity. The second outcome question asks citizens how popular they believe the average citizen thinks *Legal Report* is to measuring astroturfing's effect on secondary beliefs about popularity. My theory expects that the percentage of estimated viewers should be higher for both outcome questions in the astroturf treatment condition as compared with either of the control conditions. The third question asks them how positively they feel about *Legal Report* on a Likert scale to examine whether, contrary to my theory but consistent with persuasion accounts, astroturfing persuades them to increase their preference for the show.

5 Conclusion

Without denying that governments may sometimes astroturf to persuade or distract, I have argued that the primary purpose of astroturfing posts is to manipulate perceptions of public opinion. This takes advantage of what sets astroturfing apart from other forms of propaganda, namely that these posts purport to contain the opinions of ordinary citizens. While others have noted astroturfing can shape opinion perceptions, I contribute theoretically by arguing that this is the goal of astroturfing rather than merely a means to persuade or distract. By convincing citizens that support for the regime and its priorities are high, the regime can deter collective action by convincing dissidents that few of their fellows share their views,

making any such collective action unlikely to achieve its goals. I call this view of astroturfing as primarily a means to shape public opinion perceptions, the ballot stuffing perspective as it supposes that readers treat comments like votes when accessing public opinion, an assumption backed by previous research (Fein, Goethals, and Kugler 2007; Lee and Jang 2010; Lee 2012; Walther and Jang 2012; Chen and Lu 2024). From the government's perspective, ballot stuffing has a distinct advantage over persuasion strategies in that it does not require convincing citizens to change their own preference about the regime. Past studies have shown that such preferences are difficult to manipulate even in the best circumstances with high-production value propaganda videos, so short astroturfing comments are likely unsuitable for this goal (J. Pan, Shao, and Xu 2020, 1–2; Mattingly and Yao 2022, 2–3).

Further, I develop a wide variety of empirical implications (summarized in Table 1) to evaluate whether astroturfing activity is more consistent with this ballot stuffing interpretation or previous interpretations that argue astroturfing is primarily intended to distract from political issues that might inspire collective action or to persuade citizens to change their preferences about the regime. I find strong support for the ballot stuffing perspective by analyzing a dataset of leaked astroturf posts from China, using both computational and qualitative text-analysis methods to test these implications.

While I summarized the findings of previous research with regard to my empirical implications in the "Empirical Implications" section, future research should directly test these implications using ground truth datasets of leaked astroturf posts from other countries to examine whether the results are consistent across contexts. Another area for future research to explore is whether the new technology of large language models might allow states to create a wider variety of posts that better respond to their online context or perhaps even resolve the trade-off between quantity and quality of posts faced by human astroturfers. Finally, while I have argued that the primary purpose of astroturfing is to manipulate public opinion perceptions, I do not deny that in some cases regimes may use astroturfing to distract or persuade. Future research should examine what leads regimes to choose among these tactics.

6 References

Bowell, Tracy. 2023. "Whataboutisms: The Good, the Bad and the Ugly." *Informal Logic* 43 (1): 91–112.

CCTV. 2025. "《今日说法》Legal Report." *CCTV.com*. https://tv.cctv.com/lm/jrsf/.

Chang, Charles, and Michael Masterson. 2020. "Using Word Order in Political Text Classification with Long Short-Term Memory Models." *Political Analysis* 28 (3): 395–411. https://doi.org/10.1017/pan.2019.46.

Chen, Bin, and Shuning Lu. 2024. "In Punishment We Trust: Analyzing Characteristics and Credibility of Rumor-Debunking Messages on Chinese Social Media." *The International Journal of Press/Politics*, 19401612241291811.

Chow, Wilfred M, and Dov H Levin. 2024. "The Diplomacy of Whataboutism and US Foreign Policy Attitudes." *International Organization* 78 (1): 103–33.

Davis, Hank, and S Lyndsay McLeod. 2003. "Why Humans Value Sensational News: An Evolutionary Perspective." *Evolution and Human Behavior* 24 (3): 208–16.

Fein, Steven, George R Goethals, and Matthew B Kugler. 2007. "Social Influence on Political Judgments: The Case of Presidential Debates." *Political Psychology* 28 (2): 165–92.

Geddes, Barbara, and John Zaller. 1989. "Sources of Popular Support for Authoritarian Regimes." *American Journal of Political Science*, 319–47.

Gries, Peter Hays. 2005. *China's New Nationalism: Pride, Politics, and Diplomacy*. First paperback print. A Philip E. Lilienthal Book.". Berkeley, Calif.: Univ. of California Press.

Han, Rongbin. 2015. "Manufacturing Consent in Cyberspace: China's 'Fifty-Cent Army'." *Journal of Current Chinese Affairs* 44 (2): 105–34.

Houston, J Brian, Glenn J Hansen, and Gwendelyn S Nisbett. 2011. "Influence of User Comments on Perceptions of Media Bias and Third-Person Effect in Online News." *Electronic News* 5 (2): 79–92.

Huang, Haifeng. 2015. "Propaganda as Signaling." *Comparative Politics* 47 (4): 419–44.

———. 2018. "The Pathology of Hard Propaganda." *The Journal of Politics* 80 (3): 1034–38.

Huang, Haifeng, and Nicholas Cruz. 2022. "Propaganda, Presumed Influence, and Collective Protest." *Political Behavior*, 1–24.

Johnston, Alastair Iain. 2001. "Treating International Institutions as Social Environments." *International Studies Quarterly* 45 (4): 487–515.

Keller, Franziska, David Schoch, Sebastian Stier, and JungHwan Yang. 2017. "How to Manipulate Social Media: Analyzing Political Astroturfing Using Ground Truth Data from South Korea." In *Proceedings of the International AAAI Conference on Web and Social Media*, 11:564–67.

King, Gary, Jennifer Pan, and Margaret E Roberts. 2017. "How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, Not Engaged Argument." *American Political Science Review* 111 (3): 484–501.

Ko, Jiyoung. 2023. *Popular Nationalism and War*. New York: Oxford University Press.

Kuran, Timur. 1991. "Now Out of Never: The Element of Surprise in the East European

Revolution of 1989." *World Politics* 44 (01): 7–48.

Lee, Eun-Ju. 2012. "That's Not the Way It Is: How User-Generated Comments on the News Affect Perceived Media Bias." *Journal of Computer-Mediated Communication* 18 (1): 32–45.

Lee, Eun-Ju, and Yoon Jae Jang. 2010. "What Do Others' Reactions to News on Internet Portal Sites Tell Us? Effects of Presentation Format and Readers' Need for Cognition on Reality Perception." *Communication Research* 37 (6): 825–46.

Little, Andrew T. 2017. "Propaganda and Credulity." *Games and Economic Behavior* 102: 224–32.

Mattingly, Daniel C, and Elaine Yao. 2022. "How Soft Propaganda Persuades." *Comparative Political Studies* 55 (9): 1569–94.

Metzger, Miriam J, Andrew J Flanagin, and Ryan B Medders. 2010. "Social and Heuristic Approaches to Credibility Evaluation Online." *Journal of Communication* 60 (3): 413–39.

Miller, Blake. 2016. "Automated Detection of Chinese Government Astroturfers Using Network and Social Metadata." *Available at SSRN 2738325*.

Pan, Jennifer, Zijie Shao, and Yiqing Xu. 2020. "The Effects of Television News Propaganda: Experimental Evidence from China." *Available at SSRN 3579148*.

Pan, Yating, and Zhan Shu. 2022. "How Can Pro-Regime Online Comments Influence Chinese Netizens' Attitudes and Opinion Expressions? Normative or Informational Social Influence?" *International Journal of Public Opinion Research* 34 (4): edac036.

Parmentier, Fabrice BR. 2014. "The Cognitive Determinants of Behavioral Distraction by Deviant Auditory Stimuli: A Review." *Psychological Research* 78 (3): 321–38.

Roberts, Margaret E. 2018. *Censored: Distraction and Diversion Inside China's Great Firewall*. Princeton, New Jersey: Princeton University Press.

Rosenfeld, Bryn, and Jeremy Wallace. 2024. "Information Politics and Propaganda in Authoritarian Societies." *Annual Review of Political Science* 27.

Rozin, Paul, and Edward B Royzman. 2001. "Negativity Bias, Negativity Dominance, and Contagion." *Personality and Social Psychology Review* 5 (4): 296–320.

Sanovich, Sergey, Denis Stukal, and Joshua A Tucker. 2018. "Turning the Virtual Tables: Government Strategies for Addressing Online Opposition with an Application to Russia." *Comparative Politics* 50 (3): 435–82.

Sirikupt, Chonlawit. 2024. "Drowning Out Dissent: The Thai Military's Quest to Fabricate Popular Support on Twitter." *The International Journal of Press/Politics*, 19401612241279158.

Sobolev, Anton. 2019. "How Pro-Government 'Trolls' Influence Online Conversations in Russia." *Working Paper*. https://www.wpsanet.org/papers/docs/2019W-Feb-Anton-Sobolev-Trolls-VA.pdf.

Soroka, Stuart, Patrick Fournier, and Lilach Nir. 2019. "Cross-National Evidence of a Negativity Bias in Psychophysiological Reactions to News." *Proceedings of the National Academy of Sciences* 116 (38): 18888–92.

Waight, Hannah, Yin Yuan, Margaret E Roberts, and Brandon M Stewart. 2025. "The Decade-Long Growth of Government-Authored News Media in China Under Xi Jinping." *Proceedings of the National Academy of Sciences* 122 (11): e2408260122.

Walther, Joseph B, and Jeong-woo Jang. 2012. "Communication Processes in Participatory Websites." *Journal of Computer-Mediated Communication* 18 (1): 2–15.

Wang, Juan. 2015. "Managing Social Stability: The Perspective of a Local Government in China." *Journal of East Asian Studies* 15 (1): 1–25.

Wang, Zheng. 2012. *Never Forget National Humiliation: Historical Memory in Chinese Politics and Foreign Relations*. Contemporary Asia in the World. New York: Columbia University Press.

Wedeen, Lisa. 2015. *Ambiguities of Domination: Politics, Rhetoric, and Symbols in Contemporary Syria*. University of Chicago Press.

Wong, Stan Hok-Wui, and Jiachen Liang. 2023. "Attraction or Distraction? Impacts of Pro-Regime Social Media Comments on Chinese Netizens." *Political Behavior* 45 (3): 1071–95.

Xi, Jinping. 2012. "Achieving Rejuvenation Is the Dream of the Chinese People." *Speech*, November. https://www.neac.gov.cn/seac/c103372/202201/1156514.shtml.

Zerback, Thomas, and Florian Töpfl. 2022. "Forged Examples as Disinformation: The Biasing Effects of Political Astroturfing Comments on Public Opinion Perceptions and How to Prevent Them." *Political Psychology* 43 (3): 399–418.

Zhang, Qingpeng, Dominic DiFranzo, and James A. Hendler. 2014. "Social Networking on the World Wide Web." In *Encyclopedia of Social Network Analysis and Mining*, edited by Reda Alhajj and Jon Rokne. New York, NY: Springer New York. https://doi.org/10.1007/978-1-4614-7163-9_104-1.

Zhao, Suisheng. 1998. "A State-Led Nationalism: The Patriotic Education Campaign in Post-Tiananmen China." *Communist and Post-Communist Studies* 31 (3): 287–302. http://www.sciencedirect.com/science/article/pii/S0967067X98000099.