Catching Fire: How National Humiliation Spreads Hostile Foreign Policy Preferences on Chinese Social Media

Michael Masterson

13 July, 2023

Abstract

Research shows that emotions alter preferences, which are central to many models of political choice. However, how can theories that explain individual-level preference change explain policy outcomes, which usually take place at higher levels of analysis? I outline three competing approaches to this question, ultimately building on the third approach, which argues that emotions can spread through emotional contagion within identity groups. This implies emotions can broadly shift preferences towards or away from conflict during crises. It also explains how identities, which are relatively constant over the medium term, can suddenly become fault lines in a conflict. I find supporting evidence using a large (more than 1.6 billion posts), nationally representative data set of Chinese social media posts. Users' national humiliation posts lead their followers to become more likely to make posts about national humiliation, suggesting contagion, and about using military force and maintaining territorial disputes, suggesting preference change.

'Going viral,' in an equally evocative phrase, can be expressed as 'catching fire' (火了) in Chinese internet slang. In August and September 2012, at the same time social media

posts about China's dispute with Japan over the Senkaku/Diaoyu Islands were catching fire (Figure 2), Japanese cars literally caught fire as nationalist Chinese protesters set them alight (Gao 2012). Protests erupted in over 200 cities in the largest anti-Japanese protest since relations between China and Japan were normalized in 1972, with at least 68 cities experiencing demonstrations numbering 1,000 or more participants (Wallace and Weiss 2015, 404–5, 413). Gries, Steiger, and Wang (2016) have even argued that "popular nationalism expressed both online and in street demonstrations played a crucial role in escalating the Chinese party-state's response to the 2012–2013 Diaoyu Islands dispute" (275).

What explains this kind of sudden shift in feelings, and how does it influence preferences? One possible answer is that actors' preferences vary along with their emotions (McDermott 2004). This has broad implications for international relations theories that explain the actions of states or individual leaders based on their preferences (Fearon 1995; Moravcsik 1997; Lake and Powell 1999).

However, acknowledging that preferences vary along with emotions poses at least two challenges. First, are these changes significant enough to overcome the loss of parsimony from including them in theory? Second, even if we can explain how emotions affect the preferences of individuals, can this constitute an explanation of outcomes at higher levels of analysis, such as the state, where policy is made (Powell 2017, S274)?

I draw on group-based emotions theory from psychology to provide a partial answer to these questions (Mackie and Smith 2017, 659; Mackie, Devos, and Smith 2000), arguing that public expressions of group-based emotions, meaning emotions that are experienced by individuals on behalf of self-assigned identity categories (Goldenberg et al. 2016, 118–19), create emotional contagion by leading fellow group identifiers to become more likely to experience and express these emotions. I suggest several mechanisms for this, including the desire to conform to group emotional norms and the amplification of emotions individuals believe will help their group succeed in conflict (Barsade and Knight 2015, 26; Goldenberg et al. 2016, 123). Significantly for the political implications of this process, it is not only

the emotion that spreads but the policy preferences that go along with it.

I choose Chinese social media as the empirical context to examine this issue. Social media is a particularly important site to examine the impact of emotional narratives on policy preferences because it is a key network through which narratives spread (Halperin 2014, 74; Hall and Ross 2015, 848; Kertzer and Zeitzoff 2017). Further, Chinese leaders pay close attention to social media activity and receive daily briefings on it. While popular pressure is unlikely to shape Chinese foreign policy during normal times, it is more likely to make a difference during a crisis (Weiss 2019). Section 2.2 addresses potential concerns about censorship and regime commentators.

I choose national humiliation as the emotion to examine both because this group-based emotion is important in the Chinese context (Wang 2012; Callahan 2010) and because its affect on foreign policy preferences, namely that it increases preferences for hostile foreign policies by decreasing sensitivity to the cost of these policies, is well understood (Masterson 2022).

I use a nationally representative data set of Sina Weibo posts, which contains over 1.6 billion posts created from 2009 to 2014. Sina Weibo is a Twitter-like social media platform. Like other such platforms, users sometimes discuss politics, but most posts focus on other topics, such as entertainment news and daily life. I use supervised machine learning, meaning algorithms which are trained based on human-coded sets of posts, to label Weibo posts based on whether they contain national humiliation narratives and whether they express support for hostile foreign policies.

I find that users' posts about national humiliation increase the number of posts among their followers about national humiliation on the following day. I also find that these followers become more likely to make posts supporting a variety of hostile policies: using military force, maintaining disputed territorial claims, and erecting trade barriers. These findings hold even when controlling for media reports, international conflict events, and user and post characteristics. This suggests group-based emotions spread through emotional contagion

on social media and that this spread also shapes policy positions among group members. These findings cast light on the role of group-based emotions in otherwise mysterious shifts in public opinion and issue salience that have the potential to influence policy.

This paper is structured as follows. Section 1 lays out three competing perspectives on emotions and preference change in international relations and builds on the third one, which argues that emotions shape preferences broadly across identity groups. This section then derives hypotheses about expressions of national humiliation and support for hostile polices, and discusses case selection. Section 2 addresses data collection and coding. Section 3 descriptively analyzes the data. Section 4 formally tests the hypotheses and addresses some limitations that provide opportunities for further research. The final section offers concluding comments.

1 Emotional Contagion and Foreign Policy Preferences

Regardless of approach or paradigm, preferences play an important role in theories of international relations (Waltz 1986, 85; Fearon 1995; Wendt 1999, 123). However, in most approaches that adopt rationalist underpinnings, such as neorealism, neoliberalism, and bargaining theory, preferences formation is treated as exogenous and preferences are taken as fixed (Moravcsik 1997, 519).

However, researchers have produced a wealth of evidence that individuals' emotional states affect their preferences (Halperin et al. 2011; Hatemi et al. 2013; Zeitzoff 2014; Renshon, Lee, and Tingley 2017; Barnhart 2020). This kind of preference change differs from persuasion and socialization because it can happen quickly rather than best taking place "over long periods of time" and because the effect of emotions on preferences may dissipate just as quickly once passions die down (Johnston 2001, 510). This volatility of emotional influences on preferences is important because it gives emotional explanations the potential to account for sudden shifts in preferences during crises that could explain how

identities, such as national, ethnic, or religious identities, which have been held passively over relatively long periods of time, can result in sudden political movements or violence.

Yet, emotional accounts of political preference change also face serious challenges. First, is the variation they explain important enough to offset the loss in parsimony from incorporating emotions? Part of this question involves how widely and quickly emotions can change preferences. If emotions only change preferences for a few individuals, then their significance is in doubt. Further, if their effects take too long to spread, then the emotions of those first aroused could dissipate by the time the effect reaches others, decreasing the likelihood the emotion will result in widespread protests or political change.

Second, even if emotions cause individual preference change, can these emotions aggregate to influence policy, which is typically made by more than one individual? This recurring issue has been labeled the "aggregation problem" (Powell 2017, S274). The question the aggregation problem raises is not simply whether emotions spread between humans. Instead, the aggregation problem raises the question of whether the *political* effects on decision making also travel along with these emotions and, if so, under what conditions?

There are three possible responses to these challenges. The first is to bracket emotions because they either do not aggregate or their effect is so weak or idiosyncratic that incorporating emotions sacrifices parsimony without a commiserate return in explanatory power (Powell 2017). The second is to say that the emotions of leaders matter because they directly affect decision making and negotiation (Renshon and Lerner 2012; Wong 2019), but emotions spread too slowly or inconsistently to create popular pressure for policy change or for popular emotions to reach leaders. The third approach argues that emotions can spread broadly and quickly across identity groups through emotional contagion, enabling emotions to influence leaders' decisions both directly as leaders experience them and indirectly through popular pressure (Sasley 2011; Hall and Ross 2015).

Rationalist accounts generally take the first approach. In a classic text on rationalist approaches to international relations, Lake and Powell (1999) make a "methodological bet" that

focusing on the strategic environment rather than actors' preference changes will produce better explanations (Lake and Powell 1999, 19). In particular, they wish to avoid reliance on "untheorized" preference changes to explain behavior, suggesting that the weaknesses of current accounts of preference change, rather than rejection of the possibility of preference change, motivates this choice. Recently, Powell (2017) expanded this argument, cautioning that the aggregation problem poses major limitations to applying behavioral insights to international relations and that resolving this problem will not be easy because "bottom-up, first-image theories have so far proved very difficult to construct" (S275).

The second and third possibilities both agree that emotions spread and influence political judgements but differ on the conditions necessary for this to take place. One difference is whether emotions can spread broadly and quickly enough for the public to effectively exert pressure on policy makers. An important consideration in assessing this difference is how often the conditions necessary for emotional contagion to take place are satisfied. If only individuals who have face-to-face encounters experience the emotions of their counterparts intensely enough to be politically mobilized, then this would be more consistent with emotional contagion only acting within small groups of decision makers.

Much of the work on emotional contagion does focus on processes that require direct interaction, such as physical imitation and reactions to tone of voice (Nummenmaa et al. 2008, 572; Mackie and Smith 2017, 660), which suggests that non-face-to-face contagion, which lacks these mechanisms, might be too weak to affect political judgements. International relations theories have also posited the importance of face-to-face interactions to transmitting emotions, for example: in judging sincerity in negotiations (Hall and Yarhi-Milo 2012), in empathizing with counterparts' emotions to understand their intentions (Holmes 2013), and in triggering emotional events that affect decision-making during negotiations (Keys and Yorke 2019, 1238).

A final approach, relies on a common identity to shape the way that self-identifying members of an identity group respond to an emotional trigger, which does not need to

be experienced in person (Sasley 2011, 460; Mercer 2014). A foundational idea in this perspective is that individuals have identities at both personal and group levels, which have psychological rather than merely social consequences (Turner 1988, 2, 45–46). In other words, the identities that individuals hold affect how they perceive and interpret events, which affects the emotions they experience (Mackie, Devos, and Smith 2000). This is because they consider the event's consequences for the group rather than just its consequences for themselves as individuals (Mackie, Smith, and Ray 2008, 1870).

Others have built on this theory to explain variation within individuals who all identify with a group. For example, strong group identifiers are more likely to have intense reactions to group-based emotional triggers than weak identifiers (Mackie, Smith, and Ray 2008, 1868–69; Smith, Seger, and Mackie 2007). This implies that those who strongly identify with their nation are more likely to have emotional reactions and preference shifts. Further, different groups can be activated in different situations, such as exposure to group symbols, that can trigger different emotional responses (Mackie, Smith, and Ray 2008, 1870).

The self-categorization approach assumes that emotions have the same feeling and behavioral consequences at the group level as at the individual level (Mackie, Smith, and Ray 2008, 1873). This allows the application of studies, such as experiments examining the effect of personal emotions on preferences and decision making, to make predictions about the effect of experiencing those same emotions when they are triggered with respect to group identities.

Which group is salient to appraise a particular event is determined by accessibility, how much an individual values and uses a group identity, and fit, how well the identity accounts for perceived differences among individuals (Hogg and Reid 2006, 12). While the same person can experience different emotions depending on whether a group or individual identity is salient (Mackie, Smith, and Ray 2008, 1873), in practice, personal and group-level emotions tend to work together (Keltner and Haidt 1999, 514).

This approach also provides a unique mechanism to explain emotional contagion. When

individuals respond to emotional cues about how other group members are feeling, they engage in "emotional self-stereotyping" and adopt what they perceive to be the average emotional response among their fellow group identifiers (Mackie, Smith, and Ray 2008, 1871). The self-categorization approach largely breaks down the distinction between private acceptance in which an individual 'truly feels' an emotion and external compliance in response to perceived social pressure by arguing that both are part of the same process (Abrams et al. 1990, 98). Self-categorization theory generally holds that individuals internalize group norms (Hogg and Reid 2006, 13). While these effects may be strongest in face-to-face interactions, they can also occur when prototypical group members' emotional expressions are transmitted through media (Smith, Seger, and Mackie 2007, 443).

Despite the assumption that emotional expressions tend to be internalized, self-categorization theory also explains how emotional expressions can matter politically, even if some group members express emotions that do not match their felt emotions. This is because those expressions may shape other group identifiers' perception of what the group's average emotional response is. This is made more probable because group members interpret their fellow group members' behavior stereotypically if that behavior is ambiguous (Hogg and Reid 2006, 10–11). For example, when an online post about national humiliation creates ambiguity about whether the poster feels this emotion or is merely choosing to express it, fellow group identifiers are more likely to adopt the stereotypical interpretation that the emotion is genuinely felt. Hogg and Reid (2006) introduce the concept of "pluralistic ignorance" to describe situations in which individuals incorrectly conclude that most other identifiers accept group norms that they personally reject. Further, individuals not only respond to their perception of their group's average emotion but also polarize their emotional response away from outgroups and favor pro-norm deviation from the group emotional average (Hogg and Reid 2006, 18; Hogg, Abrams, and Brewer 2017, 527). For example, this would favor more extreme anti-Japan emotional reactions among Chinese identifiers to a prominent Japanese politician visiting Yasukuni Shrine.

This approach implies that emotional influences have the potential to be much more explosive. For example, an identity group could react largely similarly to an event that they learn about through the news or on social media. This could lead to a sudden shift in conflict preferences during a crisis that could create political pressure for leaders to respond. It also greatly increases the opportunities for leaders, who themselves identify with the relevant group, to be exposed to contagion and have their preferences directly affected.

My theory builds on this third approach. I draw on research on group-based emotions and appraisal theory to strengthen the theoretical case that political emotional contagion cascades across identity groups and does not require face-to-face interaction. Group-based emotions are triggered when an individual learns of events that are perceived as relevant for the group as a whole (Goldenberg et al. 2016, 121).

Theoretical mechanisms that could lead to non-face-to-face emotional contagion include: the construction of group norms about what emotions are appropriate to express in particular situations (Barsade and Knight 2015, 26), social comparison, in which individuals take cues from others about which emotions are appropriate to express (Festinger 1954), and exposure to emotional words triggering facial muscles related to emotional expression (Foroni and Semin 2009). Further, when individuals are given information that a fellow group member has a particular emotional reaction to an event, they become more likely to share that reaction (Mackie and Smith 2017, 659). Conforming increases bonding with fellow group members and is an inherently positive experience that can even overcome aversion to experiencing negative emotions (Goldenberg et al. 2016, 124).

While emotional responses to triggers may be temporary, group-based emotions are tied to identity categories that are more enduring. This increases the chance that they reoccur over time when individuals contemplate particular international disputes, which raises the chance that an individuals' preferences will be consistent with the preferences they have in an emotionally hot rather than cool state. However, if an individual were able to avoid emotional arousal and revisit the issue, they might be able to minimize the impact of their

previous emotion on their preferences (McDermott 2004).

I further draw on appraisal theory to explain how the emotional frames that group members encounter influence their emotional experience. According to appraisal theory, the cognitive meaning individuals attribute to events (appraisal) plays an important role in determining what emotion individuals experience (Frijda 2007, 97). Framing from fellow group members affects how group identifiers appraise events (Halperin 2014, 69). For example, if a group member frames an event as nationally humiliating, then other group identifiers are more likely to appraise this event as humiliating. Halperin (2014) notes that there is a "huge potential" for propagating information framed to trigger group-based emotions through social media (74).

In nationalist politics, the pressure on individuals to conform should be even stronger than in other forms of social groups. When individuals believe that their group-based emotion will help their group in conflict, they tend to experience that emotion more intensely (Goldenberg et al. 2016, 123). For example, citizens may have reason to believe that expressions of national humiliation will benefit China in international bargaining by sending a signal of resolve to its counterparts (Weiss 2014).

Introducing these mechanisms to international relations is critical provide a theoretical basis for the expectation that emotions cascade across identity groups. In particular Hall and Ross (2019) encourage scholars to attend to the construction of group-based emotions through political discourse, the way the salience of these constructions changes over time, and the effect of these constructions on encouraging particular political positions (1369). Without a strong theoretical account of how group-based emotional transmission takes place, the assertion that individuals within an identity group experience similar emotions risks reducing to a way of smuggling in the old strategy of ignoring the aggregation problem by personifying the group. Such accounts seem to appeal to "ethereal" or "spiritual" mediums through which emotions spread, and this perception has limited the application of group-based theories (McDermott 2014, 562). As long as some combination of the non-face-to-face

mechanisms discussed above are sufficient to allow emotional influences on political decisions to spread, then the third account of emotional contagion can be sustained along with its implications for broad and rapid preference change. I leave it to future work to establish whether some of these mechanisms are more important than others.

To produce hypotheses that will distinguish between the second and third accounts of political emotional contagion, I focus on a context in which face-to-face interactions cannot take place: text-based social media posts. Finding an effect of emotional contagion within an identity group in this context would provide important evidence for the third approach to emotional contagion.

My account of emotional contagion on social media begins with individual group members spontaneously sharing content laden with group-based emotional frames. This occurs because group members are intrinsically motivated to share group-based emotional experiences (Goldenberg et al. 2016, 129). Other group members then encounter these posts and experience the forces leading to emotional convergence within groups discussed above. As a result, they are more likely to share such content on social media themselves.

> **H1: When users express a group-relevant emotion on social media, identity group members who follow them become more likely to express this emotion in the future.**

If emotions spread from users to followers, then followers should also become more likely to take the policy positions that these emotions motivate.

> **H2: When users express a group-relevant emotion on social media, other identity group members who follow them become more likely to express support for foreign policies consistent with that emotion's influence on preferences**

1.1 National Humiliation and China

Among the group-relevant emotions linked to foreign policy preferences, national humiliation is arguably the most important. Researchers have pointed to "humiliation" as a motivation for Russia's 2014 annexation of Crimea (Larson and Shevchenko 2014). Humiliation has also been linked to hostile foreign policy preferences in the United States during the Vietnam War (Steinberg 1996), French imperial territorial conquest (Barnhart 2020), and Chinese escalation during foreign policy disputes (Callahan 2010, 96; Wang 2012). Further, the mechanism through which humiliation alters conflict preferences, decreasing sensitivity to the cost of conflict (Masterson 2022), is well understood, which enables me to create a variety of policy preference change predictions to test.

This previous research on the microfoundations of humiliation's effect on conflict preferences also helps address a potential concern that exposure to emotional expression signals users to express preferences they already have rather than causing preference change. Experimental work has found that humiliation *does* change individuals preferences over conflict, even in contexts in which conflict poses monetary costs to respondents (Masterson 2022). For this reason, while a change in expressed policy preferences following exposure to humiliation content could be read as either preference change or merely a strategic choice to express existing but previously unexpressed preferences (see section 4.3 on the challenges of separating evidence for expressed and felt emotion outside of a laboratory context), the preference change interpretation is more grounded in the experimental work on humiliation.

China is the context in which national humiliation has been most researched (Callahan 2010, 96; Wang 2012; Cohen 2002). Further, using Chinese social media has the advantage of combining its massive user base with the fact that, unlike other forms of social media such as Twitter, the Sina Weibo user base is not very internationalized, making it highly likely that followers will share the same national identity as users, which is a necessary condition for group-based emotions to operate (Wan 2018). China's rising economic and military power also make understanding preferences changes that could affect its foreign policy vital. For

these reasons, I choose China as the country context to examine emotional contagion.

I draw on previous work that has found that humiliation leads individuals to support hostile foreign policy actions by decreasing their sensitivity to the cost of these actions to derive some specific policies that humiliated individuals should be more likely to support (Masterson 2022). The first is using military force. Using military force is always a hostile action, and generally, it is considered costly (Fearon 1995). From China's perspective, using military force might be particularly dangerous because it could escalate to draw in the United States. Yet, research has found that invoking US deterrence and "cost-imposition threats" actually decreases support among Chinese citizens for backing down during disputes (Quek and Johnston 2018, 10).

Another costly, hostile foreign policy is maintaining territorial disputes. These disputes are costly both because they risk escalating to war, and because they entail substantial economic costs in the form of lost trade and investment (Simmons 2002). Chinese propaganda has linked national humiliation and China's territorial disputes, which could make it particularly likely that Chinese citizens think of these disputes in terms of national humiliation (Zhao 2004, 231).

The final foreign policy position that I investigate is economic rather than military. While lower on the hostility scale than military force, adopting economic barriers against other countries can be considered hostile. The current trade war between the United States and China exemplifies this. Anti-trade measures are also costly to economic welfare (Alston, Kearl, and Vaughan 1992). Anecdotal evidence suggests that Chinese citizens connect China's current trade war with the United States to China's historic national humiliation (Su 2019). Further, during anti-Japan protests associated with national humiliation, protesters often call to boycott Japan (N. D. 2012).

Therefore, users' posts about national humiliation will increase the likelihood followers make posts on the next day that…

H2a: advocate using or threatening to use military force.

H2b: support China's disputed territorial claims.

H2c: advocate adopting trade restrictions.

1.2 Does Public Opinion Matter in China?

While the focus of this paper is to show the effect of emotional contagion on policy preferences rather than to examine policy change itself, the significance of emotional contagion on Chinese social media would of course be diminished if these opinions failed to influence policy. Researchers have argued that Chinese leaders live in the same "social milieu" as ordinary citizens, implying that these leaders should be subject to the same forces of emotional contagion (Cairns and Carlson 2016, 40). Further those who identify more strongly with the group tend to experience stronger emotions from group-related events, which could suggest that leaders are even more susceptible to emotional influence (Goldenberg et al. 2016, 122; Mackie and Smith 2017, 660; Barnhart 2020, 60).

However, political scientists often believe leaders are more rational or strategic than ordinary citizens, which might make them less susceptible to these influences. For this reason, this section draws on previous research to argue that Chinese public opinion matters, making it important to consider even for readers who are skeptical of the claim that leaders themselves become emotional.

Previous research argues public opinion is an important policy input in China. Nationalism both produces legitimacy for the CCP and exposes it to possible "serious backlash" if citizens perceive the regime as too dovish (Zhao 2004, 264–65). This backlash could even threaten regime stability. For example, Chinese officials fear citizens will brand them "traitors" if they advocate moderation on China's territorial disputes (Lynch 2015, 196).

Emotional nationalism may play an even greater role in authoritarian regimes than in democracies because, unlike other viewpoints, citizens are more likely to feel free to express nationalism publicly. Further, when authoritarian regimes rely on nationalism for legitimacy, they may be unable to repress nationalists without risking stability (Gries 2005, 46, 120).

I am not suggesting that the Chinese state is simply a conduit that translates public opinion into policy. Chinese leaders both attempt to shape public opinion and manage the opinions that can be expressed (Creemers 2017). Leaders may selectively permit the expression certain opinions while quashing others (Reilly 2012). However, the effort that leaders exert to manage public opinion also shows that leaders care about its content. That China's leaders get a daily briefing on online social media activity further shows they regard it as important (Weiss 2019, 683).

Even if domestic political opinion is not an absolute constraint on Chinese foreign policy, it may increase leaders' political costs of compromise, making compromise less likely (Weiss 2019, 679). Further, popular pressure may be most likely to be decisive during "major crises and conflicts," when foreign policy decisions are critically important (Weiss 2019, 694). Even state-driven emotional messaging may blowback and become constraining because leaders may struggle "to simply switch off" contagion as it spreads to more actors who demand political responses consistent with this messaging (Hall 2015, 52–53). In sum, while public opinion is not an entirely independent influence on Chinese foreign policy, it is not epiphenomenal either. In an authoritarian context, control of opinion may be greater than in a democratic context, but this control is still imperfect. Because autocrats face harsher consequences upon removal from office than democratic leaders, such as death and imprisonment, autocrats have reason to fear disapproval (Debs and Goemans 2010).

2 The Data

Weibo, unlike other forms of Chinese social media, such as WeChat, has the advantage that users post publicly rather than sending individual or group messages, which both increases the potential for emotional contagion and enables these posts to be observed by researchers. The Weibo posts analyzed here come from a data set originally collected by a group of

Chinese scholars studying natural language processing (NLP) (Zhang et al. 2015).[1] The Fudan NLP Group's website seems to have removed the webpage hosting this data set, and the data set is no longer publicly available. The data were originally collected by randomly selecting Weibo users and collecting the first 50 pages of posts from these users and their followers. Because the initial selection of users was random, the data should be relatively representative of opinion on Weibo. In total, the data set has 1,676,535,827 posts from 2.4 million users.[2] The earliest post is from August 13, 2009, and the latest is from March 12, 2014. To my knowledge, this data set has not previously been used to study politics.

Unfortunately, many datasets that provide windows into authoritarian politics involve some form of opacity. This is likely to get worse as Chinese websites make it harder to scrap data and the Chinese government increases censorship (Yang 2022). It is not just online research that is affected. Those still able to conduct political surveys in China often must keep secrete the name of the company involved in conducting the survey, which obscures information about methods and sample quality, which can vary by company. This places researchers in a position in which they must either draw conclusions using data around which there is some uncertainty or largely forego studying these countries.

However, there are actions researchers can take to increase confidence in their findings. I follow best practices from King, Pan, and Roberts (2017). King, Pan, and Roberts (2017) analyzed a dataset of regime commentator posts leaked by an anonymous blogger whose blog now appears to be defunct. Similarly to the data set they used, the massive size of this data set and its complexity suggest that fabrication would be extremely difficult and is therefore unlikely (495). Further, like King, Pan, and Roberts (2017), I verify external references in select posts as well as checking to see if these posts correspond to real posts

[1]The data set was updated after Zhang et al. (2015) was published. The latest information is based on the Fudan NLP Group's website's description of the data before it was taken down. Regarding data usage permissions, the website said, "The platform shares the data sets of social media, such as public comments and Weibo posts, for non-commercial, non-profit research, and the laboratory reserves the right to interpret" (Fudan NLP Group n.d.).

[2]This is after removing 2,674,413 posts that had no textual content or were missing a unique user or row identifier. Presumably the posts without content only contained image or video content, which was not collected.

that exist online. The fact that this data set was initially collected by computer science researchers who were unlikely to be thinking about the possibility it might be later used to study politics also decreases the likelihood it was manipulated for political purposes. Section 2.2 addresses concerns about censorship and regime commentators.

2.1 Coding

Before coding the data, I filter it to decrease variable rarity. It is inefficient and resource prohibitive to hand code randomly selected posts if less than 0.2% of them contain the variables being analyzed (see Figure 1). Further, machine learning will perform poorly if some categories are too rare (Chang and Masterson 2020, 406–7). I filter the data with two highly inclusive keyword lists. The first was designed to select posts that *could* contain narratives of national humiliation. The second was designed to select posts that *could* advocate: using military force, maintaining disputed territorial claims, or adopting trade barriers (see appendix section "Keyword Lists" for the full lists of keywords).

To be clear, this keyword selection is only to eliminate posts that have a very low probability of containing the key variables. That is why the lists are highly inclusive, meaning that they are designed to include nearly all posts that contain the key variables. Because the threshold for inclusion is intentionally set low, the selected sets still contain many irrelevant posts. Reflecting this, only about 5.4% of the posts selected by keyword to be coded for national humiliation were confirmed by human coders as actually representing this variable. The percentages for the international policy variables are similar.[3] In each case, while the sought after variables are still rare in the keyword-selected set, they are much more common than they are in the overall data set. Given the highly inclusive nature of the keyword-selected posts, I assume that all posts that contain the key variables fall in these sets, meaning posts not keyword selected are not subjected to machine labeling and are coded as zero for those variables.

[3] 1.8% for military force, 8.1% for territory, and 0.6% for trade barrier.

Importantly, the supervised machine learning algorithms *do not use the keyword lists to classify posts*. Instead, they are trained based on *human-coded* training sets of posts. To construct the training set for national humiliation posts, I randomly selected 5,000 of the keyword-selected posts for national humiliation to be coded by a research assistant. Because Weibo posts often criticize and reference each other, the coder was instructed to code posts based on the message the author of the post wrote and only use quoted text from other posts as context for interpreting the author's message. In cases in which posts deployed jokes or sarcasm, the coder (a native Chinese) was instructed to use her knowledge of the cultural context as well as information about the timing of the post to interpret its meaning (Hall 2017, 491). While the "Coding Dictionary" section of the appendix contains the full definitions used to code each category along with example posts, this section contains summaries of how each category was defined. First, `national humiliation` posts are posts that represent a foreign humiliator inflicting injustice on the Chinese nation.

Another set of 4,999 posts was randomly selected from the posts found by keyword to potentially contain one or more of the policy variables and coded, separately from the humiliation posts, by a research assistant for these variables. Posts were coded as supporting military `force` if they advocated using or threatening to use China's armed forces against another country.[4] Posts were coded as containing the `territory` variable if they raise China's claim to a disputed territory. Posts were coded as containing the `trade barrier` variable if they advocated boycotting or raising trade barriers against foreign goods. As a quality control on the human labels, I coded 200 of the posts from both sets and checked intercoder reliability. For each category, the research assistant's coding agreed with my own at least 95% of the time.[5]

While researchers have previously used other methods to classify texts based on psy-

[4]Initially, there were two military force variables. One for posts that implicitly advocated force and another for posts that explicitly advocated it. These were combined for machine labeling and analysis to decrease rarity and increase accuracy of machine labeling.

[5]Krippendorff's α for each category was as follows: 0.688 for national humiliation, 0.705 for military force, 0.878 for territory, and 1 for trade barrier.

chological concepts (Tausczik and Pennebaker 2010; Neuman and Cohen 2014; Xue et al. 2020), none of these models are designed to measure humiliation specifically, let alone national humiliation. As I now have "gold standard" human-coded training sets based on my coding dictionary and wish to label posts based on these categories rather than discover new categories, I use a supervised machine learning algorithm to label the posts (Wilkerson and Casas 2017, 4).

To select which algorithm to use, I follow Chang and Masterson (2020) on selecting supervised classification algorithms for unbalanced text data and try long short-term memory models, gradient boosting machines, and support vector machine (SVM) models with the loss function weighted to penalize misclassification of the rarer category more highly (406–07). To evaluate the models, I use a 0.8/0.2 train/test split. I find that SVM models perform best for my particular application, so I use these models.[6] Each model classified at least 96% percent of the posts in the withheld test set correctly. See the appendix section 1 for additional information on model performance, preprocessing, and feature selection.

I include two additional variables that measure post content, which, unlike the others, are coded based on keywords only. First, I create a `patriotism` variable, which takes the value of 1 if a post contains the Chinese word 'patriotism' (爱国) and 0 otherwise. This variable serves as a control for general nationalism that could make users both more likely to share national humiliation narratives and particular policy positions.

I further include an `income inequality` variable both as a control for non-foreign-policy-related political discussion and as a placebo test, since it is a political issue that national humiliation narratives should not theoretically be related to. Both of these steps are designed to deal with the concern that national humiliation narratives might be more common in political posts generally. I selected income inequality because Chinese survey respondents ranked it as their second biggest concern, so it should be salient (Chubb 2014).[7] I use the

[6]These models were created using Python 3 and the machine learning package scikit-learn.

[7]I choose this instead of the number 1 issue of corruption because national humiliation posts often attribute China's historic vulnerability to humiliation to corruption, for example, of the Qing dynasty.

same phrase to measure income inequality as was used in the survey: "rich-poor disparity" (贫富分化) (Chubb 2014, 22).

2.2 Censorship and Regime Commentators

When assessing the possible impact of both censorship and regime commentators, it is important to distinguish between conditions that would bias the descriptive estimates of each type of post as a proportion of total Weibo posts and conditions that would bias the hypothesis tests. If, for example, censors were more likely to remove posts that advocated military force, this would lead to underestimating the amount of calls for China to use military force on Weibo, but it would not bias the estimate of the relationship between humiliation and force unless censors were systematically more/less likely to remove calls for force that contained narratives of national humiliation as opposed to other calls for force. Similarly, if regime commentators are more likely to post about national humiliation than ordinary citizens, this would increase the amount of national humiliation posts as a proportion of total posts, but it would not bias the hypothesis tests unless these posts are also more/less likely to express support for a particular policy. Even if regime interference on the treatment and one or more of the outcomes is related within posts, this would not bias the analysis of the effects of users' posts on their followers' posts the following day so long as censors/commentators cannot condition their activity on followers' future responses. I discuss the likelihood and possible direction of biases driven by censorship and regime commentators below with a focus on biases that would confound the hypothesis tests. Because the data set ends in 2014, I do not discuss the increased restrictions on free discussion imposed after this period.

While it is impossible to directly measure the influence of automated keyword censorship, which prevents messages from being posted, this kind of censorship is "not sophisticated or very successful, and therefore much of content filtering is done by hand" (Roberts 2018, 154). If posts about foreign policy crises were keyword filtered during this period, then we would expect no posts about the Diaoyu Islands during the massive 2012 anti-Japan protests over

the dispute. Instead, posts about the Diaoyu Islands reach a peak during this period.

To investigate the influence of the more threatening kind of censorship in which humans working for the regime manually censor posts, I cross-reference the users from my data with the WeiboScope data set (Fu, Chan, and Chau 2013). The WeiboScope data contain posts that were collected in real time and shows whether they were later removed. The authors exploited a feature of a previous version of the Weibo API to distinguish between inaccessible posts that returned as "permission denied," indicating censorship, and posts that were simply deleted later, which could indicate that the user deleted their own post (Fu, Chan, and Chau 2013, 44). My data set contains 396,947 users that overlap with the WeiboScope data, allowing me to view which of the more than 3.4 million posts that these users made in 2012 were censored. Because the WeiboScope data focused on users that had many followers, these users are especially influential and, hence, if anything, more likely to experience censorship than the average user in my data set (Fu, Chan, and Chau 2013, 43).

In the appendix "Censorship" section, I examine both whether any particular kinds of posts predict the likelihood of being censored, which would bias the descriptive analysis but not necessarily the relationship between the treatment and the outcome, and whether humiliation posts are more or less likely to be censored when they also contain one of the outcome variables. I use both the preferred measure of censorship from Fu, Chan, and Chau (2013), which I label 'narrow' as well as a 'broad' measure of censorship that considers all posts deleted for any reason as censored for robustness.

For the narrow measure, neither the treatment nor any of the outcomes predict censorship (Table 6 in appendix). For the broader measure, humiliation does not predict censorship, but there is some evidence that posts about territory and trade barriers are more likely to be deleted. This implies the descriptive analysis may undercount these posts. Regardless of whether the narrow or broad measure is used, none of the interactions of humiliation and the outcomes predict censorship once user and day fixed effects are included (Table 7 in appendix), which increases confidence that the results are not driven by censorship.

These findings are consistent with previous work on Chinese censorship, which finds posts are removed not because they contain political content but because they contain calls to collective action. Censorship is costly to the regime because it risks backlash and decreases the information the regime can gain about public opinion, so political posts are not blanketly censored (Roberts 2018, 13). From 2009 to 2014, online expression about international relations received "relatively minimal interference from censors" compared to other political issues (Chubb 2014, 56). This makes it less likely censors specifically target posts containing these variables.

Another concern is self-censorship. While self-censorship is unique neither to authoritarian contexts nor social media (Das and Kramer 2013; Hayes, Scheufele, and Huge 2006), readers might reasonably fear it is particularly of concern for online expression in China. However, for most Chinese the threat of punishment for online comments is not very credible, and online criticism of the regime on is common (Roberts 2018, 13). Contrary to the idea of prevalent self-censorship, Chinese citizens experiencing censorship on Weibo may actually become more likely to post about sensitive topics in the future (Roberts 2018, 117). However, recent studies have found evidence of self-censorship in China on sensitive survey questions, such as how citizens evaluate the regime (Shen and Truex 2021; Robinson and Tannenberg 2019).

To the extent that self-censorship removes alternative narratives about foreign policy from public discourse in China, this may create an environment more conducive to emotional contagion by leading citizens to draw the conclusion that their fellow group members are having a more uniform emotional response to events than they really are, increasing the pressure to emotionally conform (Hall and Ross 2019, 1366). However, this is not the same as biasing the study in the sense of causing hypotheses tests to be supported when no emotional contagion has actually taken place. Instead, if self-censorship functions in this way, it would imply that authoritarian contexts in which the regime restricts online foreign policy discourse to a single emotional narrative are cases in which emotional contagion is

particularly likely to occur.

Regime commentators, who create approximately 1 in every 178 Chinese social media posts, are another concern (King, Pan, and Roberts 2017). However, these posts rarely contain political content and are instead intended to distract from political issues (King, Pan, and Roberts 2017, 485). For example, after a 2014 earthquake that led the regime to expect criticism for poor building construction, the regime tried to encourage online commentators to discuss a celebrity scandal rather than to defend its policies (Roberts 2018, 190–91).

However, it is possible that "cheerleading" regime commentator posts may mention national humiliation because these posts sometimes reference patriotism and martyrdom (King, Pan, and Roberts 2017, 489–90). Still, these cheerleading posts do not discuss "specific policies" (King, Pan, and Roberts 2017, 499). To get an empirical sense of whether regime commentator posts might bias the hypothesis tests, I examine the data set of 43,757 leaked regime commentator posts analyzed in King, Pan, and Roberts (2017). Of these posts, 61 contain the phrase, 'national humiliation' (国耻). None of these posts mention territorial disputes, advocate using military force, or advocate raising trade barriers. This increases confidence that the hypothesis tests will not be driven by regime commentator posts. See the "Regime Commentator" section of the appendix for further discussion and the full text of these posts.

3 Descriptive Analysis

Because the prevalence of different foreign policy positions on Chinese social media is interesting in itself, this section descriptively analyzes the data. The following section formally assesses the hypotheses. Figure 1 shows the estimated percentage of total posts that fall into each category.

These are consistent with previous work, which has found that only a small minority

of Weibo posts contain political content (Chang and Masterson 2020, 407). However, the volume of posts in these categories is not trivial. The trade barrier category, which is the smallest category with only 0.009% of posts, represents 150,641 posts in the overall sample. The territory category, which is the largest category with 0.1306% of posts, represents 2,189,668 posts in the overall sample. Further, the approximately 2.4 million users in the data set are only a small proportion of the nearly 200 million monthly active Weibo users at the beginning of 2015 (Incitez China 2015). Multiplying these totals by 100 to account for this amounts to a substantial discussion of foreign policy on Weibo.

Figure 1: Post Frequency

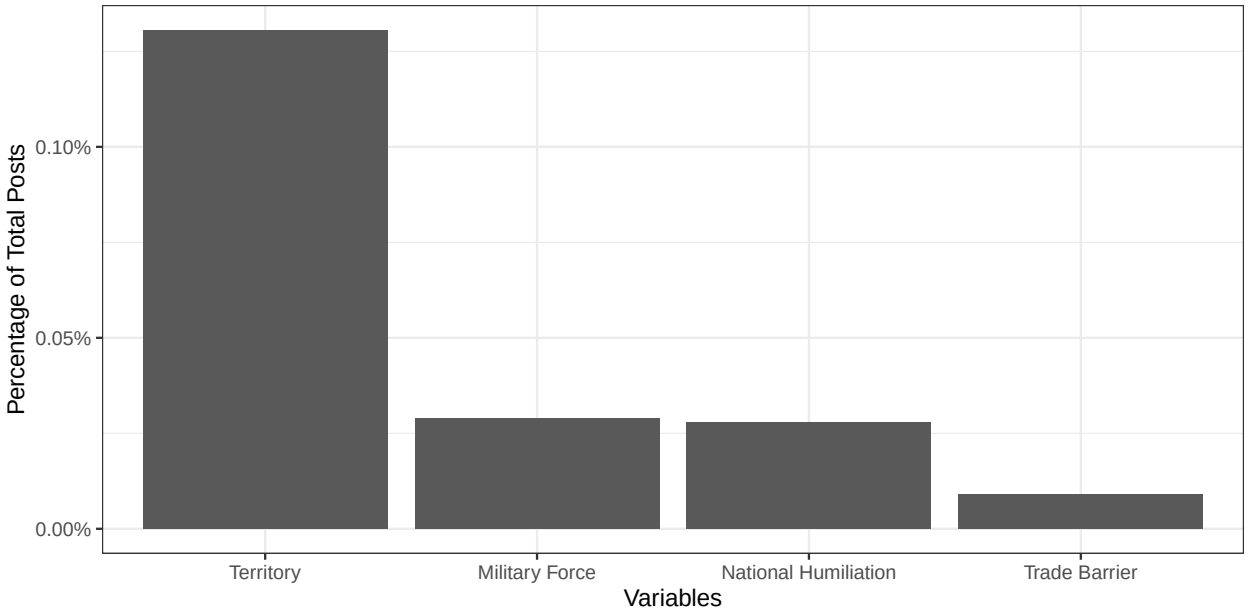

Figure 2 shows how the proportions of posts in the overall sample that mention territorial disputes, contain national humiliation narratives, or advocate using force vary over time. It appears that these posts covary. Further, posts about territory have their largest spike during the 2012 Diaoyu Islands crisis when the Japanese government moved to purchase the islands from their private owners, suggesting that the machine-coded data are picking up real variation.

Figure 2: Posts Over Time

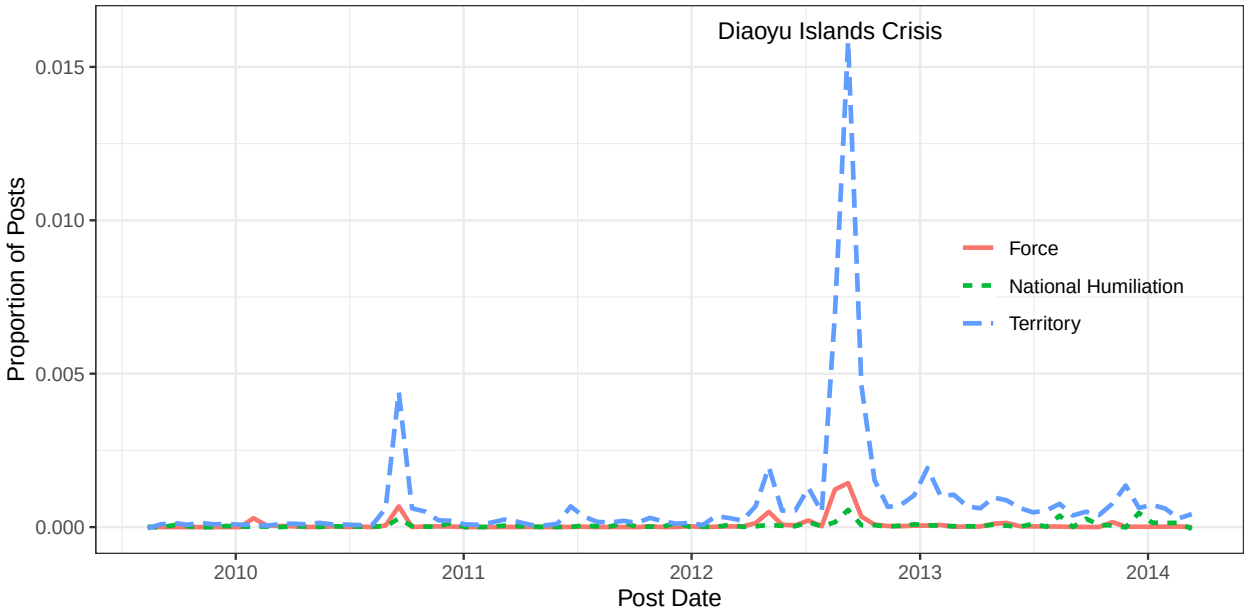

Variables smoothed using locally estimated scatterplot smoothing (LOESS).

4 Hypothesis Tests

Because the hypotheses are about the effect of users' posts on followers' posts, I transform the data set to the user-day level of analysis, resulting in 1,030,521,880 user-day observations. The user portion of the unit of analysis only includes the 628,367 users (out of the original 2.4 million) who have at least one follower in the data set who posts during the time period data was collected. The national humiliation, territory, force, trade barrier, income inequality, and patriotism variables become daily counts of a user's posts in those categories.

4.1 Modeling Strategy

I lag the user variables (and controls) one day to examine the effect of users' posts on their *followers' posts the following day.* All models are OLS linear regressions. In addition to the income inequality and patriotism variables discussed in the coding section, I control for all *People's Daily* articles in the *People's Daily* OriProbe Information Services data base that contain "national humiliation" (国耻) in the text. The *People's Daily* is the CCP's main

"mouthpiece" newspaper, and it sets the tone for other Chinese media, which sometimes must copy its reports verbatim (Wu 1994, 195). For this reason, even though it is only a single source, it should be a good control for discussion of national humiliation and the framing of events as nationally humiliating in Chinese media overall. This will help control for confounds that result from users and their followers having similar tendencies to follow media, which could trigger them both to post about national humiliation after such stories.

Another potential confound is that an international event could trigger both users and their followers to post about national humiliation and particular policy responses. To control for international conflict events, I include a dummy variable for posts that occurred during days when China had an ongoing militarized interstate dispute (MID) (Palmer et al. 2021). While this might seem a coarse measure, 186 MIDs involving China took place during the time period that overlaps with the data, making up about 25% of the days in the data set, so the variable contains substantial variation. Further, it does capture the period in September 2012 when nationalist posts surged because of the Diaoyu Islands crisis.

I also control for attributes of the user and their daily activity. These controls include: post length, which measures the length in characters of all of a user's posts on a particular day, each user's number of followers, and the number of posts each user made on each day. Finally, I include both user and day fixed effects to control for unobserved confounds that are constant for particular users or days. Examples of such confounds include: politically engaged users being more likely to have politically engaged followers, particular days when users were more likely to post about both foreign policy and national humiliation due to events not captured by the MIDs variable, and user demographic features.

4.2 Results

While Table 1 shows the effect of the number of posts a user makes about national humiliation on their followers' posts both with and without controls, all discussion of effect size will focus on the models that include controls. For every post a user writes about national humiliation,

their followers post 0.13 (0.62σ) more posts about national humiliation on the following day. This supports H1 that the emotion of national humiliation spreads from users to followers who read their posts. H2a–H2b are also supported as each post a user produces about national humiliation corresponds to increases of 0.106 (0.63σ) posts about using force and 2.314 (0.73σ) posts about maintaining territorial disputes among their followers on the next day. In practice, these effects would add up as followers follow multiple users who are posting about national humiliation.

In contrast, the effect of users' humiliation posts on followers' trade barrier posts (H2c), while statistically significant and in the hypothesized direction, has an effect size so small that it is difficult to argue it is substantively significant. A national humiliation post only produces an increase of 0.04 posts (0.26σ) about raising trade barriers among followers on the following day. This implies 25 posts about national humiliation are required, on average, to get a follower to post about trade barriers once.

Table 1: Contagion Results: Humiliation's Effect on Followers' Next-Day Posts

	Humiliation	Force	Territory	Trade Barrier	Humiliation	Force	Territory	Trade Barrier
National Humiliation	0.275**	0.174**	4.005**	0.093**	0.13**	0.106**	2.314**	0.04**
	(0.002)	(0.001)	(0.025)	(0.001)	(0.002)	(0.001)	(0.024)	(0.001)
Controls	No	No	No	No	Yes	Yes	Yes	Yes
User Fixed Effects	No	No	No	No	Yes	Yes	Yes	Yes
Day Fixed Effects	No	No	No	No	Yes	Yes	Yes	Yes

N = 1,030,521,880 user days. Controls include: *People's Daily* national humiliation articles, MID indicator, patriotism, income inequality, characters posted, number of followers, number of posts. All results rounded to the third decimal place. * indicates $p < 0.1$ ** indicates $p < 0.05$.

It may be that discussion of one political issue correlates with discussion of other political issues. In this case, the relationships between national humiliation and the policy variables might be spurious to general cycles of political discussion. To assess this concern, I conduct placebo tests with income inequality, which is a political issue that should not theoretically be related to national humiliation, as the dependent variable. I find no evidence of this kind of bias.[8]

[8]The results are available from the author upon request.

4.3 Limitations and Areas for Further Research

While examining observational data captures expressed preference change in an important real-world case, it also creates challenges to inference that would be absent in a laboratory context. This trade off between a real-world context and precise control cannot be resolved in any single study but can be addressed as studies of different types accumulate within the research program as a whole.

The first limitation is that observed emotional expressions may differ from felt emotions. Followers might echo the posts of users for reasons seemingly unrelated to emotional contagion. For example, followers might want to get on board with a lively trend, wish to validate the user, signal their values, or simply gain entertainment from using emotional language in the discussion.

While the possibility expressed emotions may not match felt emotions is indeed an important concern, if taken too far it would reduce researchers to "an old form of behaviorism in that they restrict themselves to observable billiard-ball interactions," effectively restricting the study of emotions to laboratory contexts (Jasper 2018, 11). Further, *understanding emotional expression is critical* to understand emotional contagion because both fellow identity group members and the regime are in the same boat as researchers in that they only observe expressed emotion. This allows these expressions to influence their assessment of how common a particular emotion is among fellow identifiers and potentially result in emotional contagion, whether or not the expressed emotion is genuinely felt as discussed in section 1.

Likewise, in terms of policy influence, what these posts signal to leaders about citizens' sentiments may be more significant than what citizens actually feel. Autocrats want to avoid collective action, such as protests, that emotionally motivated citizens might engage in, so they have incentives to try to placate these sentiments before they escalate to offline action (King, Pan, and Roberts 2013). This means autocrats have incentives to act before citizens incur the cost of protesting that would differentiate emotionally motivated citizens from those simply following along.

A second limitation is that the treatment of ingroup emotional expression is potentially bundled with other variables, such as information about events, which users' Weibo posts may also contain. While I have attempted to account for this by controlling for events (MIDs), reporting on events (*People's Daily* articles), and variables that measure the content of the Weibo posts and characteristics of the user, these steps are inevitably imperfect.

5 Conclusion

This article has examined the conditions under which emotions spread socially and affect preferences. I find both that national humiliation posts make followers more likely to share posts containing national humiliation narratives in the future and that these posts motivate followers to take more hostile foreign policy positions on using force and maintaining territorial disputes. These findings hold even controlling for national humiliation news, international conflict events, and user and post characteristics.

The results suggest emotions spread and influence preferences rapidly and broadly across identity groups. This implies emotional activation can transform public opinion in response to events that implicate political identity groups, such as nations, which makes it more likely that public pressure will be sufficient to influence policy makers. Further, it also means these triggers are more likely to reach policy makers because leaders need not have individuals experiencing these emotions in their inner circle to be exposed. This has applications for a broad range of theories about political identities, including national, ethnic, and religious identities, which often argue that these identities influence policy but may have a hard time articulating how an identity that is relatively constant in the short to medium term produces sudden political movements and demands for policy change.

These findings also have implications for international conflict and Chinese foreign policy. The Chinese government promotes narratives of national humiliation with regard to many of its current foreign policy disputes. If these narratives drive individuals towards more costly,

hostile policy options, their propagation may increase pressure on the state to adopt more hawkish policies in the future. Further, the evidence suggests that users' emotional narratives influence their followers. This implies that once social media users begin posting about national humiliation, this activity may take on a life of its own as users' posts influence their followers whose posts in turn influence their followers' followers, and so on. This raises the specter of a vicious cycle that could magnify the effect of crises, propaganda, and nationalist groups.

Showing emotions can spread broadly and rapidly is a necessary condition for citizens emotions to influence political outcomes, but it is not the same as showing that these pressures directly shape these outcomes. Future research should examine the effect of social media discussions on policy itself. A convincing study of this kind might need to examine the internal deliberations of policy makers, to capture variation in their positions and whether they mention popular pressure from social media as a policy justification. Future experimental research might also examine which kinds of non-face-to-face transmission mechanisms are needed for emotions to spread and retain their political influence. For example, researchers could independently manipulate social comparison, emotional norms, and the emotional words contained in messages.

6 References

Abrams, Dominic, Margaret Wetherell, Sandra Cochrane, Michael A Hogg, and John C Turner. 1990. "Knowing What to Think by Knowing Who You Are: Self-Categorization and the Nature of Norm Formation, Conformity and Group Polarization." *British Journal of Social Psychology* 29 (2): 97–119.

Alston, Richard M, James R Kearl, and Michael B Vaughan. 1992. "Is There a Consensus Among Economists in the 1990's?" *The American Economic Review* 82 (2): 203–9.

Barnhart, Joslyn. 2020. *The Consequences of Humiliation: Anger and Status in World Politics*. Ithaca: Cornell University Press. http://public.eblib.com/choice/PublicFullRecord.aspx?p=5964952.

Barsade, Sigal G, and Andrew P Knight. 2015. "Group Affect." *Annu. Rev. Organ. Psychol. Organ. Behav.* 2 (1): 21–46.

Cairns, Christopher, and Allen Carlson. 2016. "Real-World Islands in a Social Media Sea: Nationalism and Censorship on Weibo During the 2012 Diaoyu/Senkaku Crisis." *The China Quarterly* 225: 23–49.

Callahan, William A. 2010. *China: The Pessoptimist Nation*. Oxford ; New York: Oxford University Press.

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.

Chubb, Andrew. 2014. *Maritime Awareness: Popular Attitudes Towards China's Maritime Disputes in the East and South China Seas*.

Cohen, Paul A. 2002. "Remembering and Forgetting National Humiliation in Twentieth-Century China." *Twentieth-Century China* 27 (2): 1–39.

Creemers, Rogier. 2017. "Cyber China: Upgrading Propaganda, Public Opinion Work and Social Management for the Twenty-First Century." *Journal of Contemporary China* 26 (103): 85–100.

Das, Sauvik, and Adam Kramer. 2013. "Self-Censorship on Facebook." In *Proceedings of the International AAAI Conference on Web and Social Media*, 7:120–27.

Debs, Alexandre, and H. E. Goemans. 2010. "Regime Type, the Fate of Leaders, and War." *American Political Science Review* 104 (3): 430–45. https://doi.org/10.1017/S0003055410000195.

Fearon, James D. 1995. "Rationalist Explanations for War." *International Organization* 43 (3): 379–414.

Festinger, Leon. 1954. "A Theory of Social Comparison Processes." *Human Relations* 7 (2): 117–40.

Foroni, Francesco, and Gün R Semin. 2009. "Language That Puts You in Touch with Your Bodily Feelings: The Multimodal Responsiveness of Affective Expressions." *Psychological Science* 20 (8): 974–80.

Frijda, Nico H. 2007. *The Laws of Emotion*. Mahwah, N.J: Lawrence Erlbaum Associates.

Fu, King-wa, Chung-hong Chan, and Michael Chau. 2013. "Assessing Censorship on Microblogs in China: Discriminatory Keyword Analysis and the Real-Name Registration Policy." *IEEE Internet Computing* 17 (3): 42–50. https://doi.org/10.1109/MIC.2013

.28.

Fudan NLP Group. n.d. "Shùjù Jí Xìnxī." Accessed November 22, 2019. http://sma.fudan.edu.cn/datainfo/weibo.html.

Gao, Jimmy. 2012. "China's Anti-Japan Protests Attract Some Chinese Criticism." *The Atlantic*, September. https://www.theatlantic.com/international/archive/2012/09/chinas-anti-japan-protests-attract-some-chinese-criticism/262481/.

Goldenberg, Amit, Eran Halperin, Martijn van Zomeren, and James J Gross. 2016. "The Process Model of Group-Based Emotion: Integrating Intergroup Emotion and Emotion Regulation Perspectives." *Personality and Social Psychology Review* 20 (2): 118–41.

Gries, Peter Hays. 2005. "China's "New Thinking" on Japan." *The China Quarterly*, no. 184 (December): 831–50. http://www.jstor.org/stable/20192541.

Gries, Peter Hays, Derek Steiger, and Tao Wang. 2016. "Popular Nationalism and China's Japan Policy: The Diaoyu Islands Protests, 2012–2013." *Journal of Contemporary China* 25 (98): 264–76.

Hall, Todd H. 2015. *Emotional Diplomacy: Official Emotion on the International Stage*. Ithaca; London: Cornell University Press.

———. 2017. "Three Approaches to Emotion and Affect in the Aftermath of the Zhuhai Incident." *International Studies Review* 19 (3).

Hall, Todd H., and Andrew A. G. Ross. 2015. "Affective Politics After 9/11." *International Organization* 69 (4): 847–79. https://doi.org/10.1017/S0020818315000144.

———. 2019. "Rethinking Affective Experience and Popular Emotion: World War i and the Construction of Group Emotion in International Relations." *Political Psychology* 40 (6): 1357–72.

Hall, Todd H., and Keren Yarhi-Milo. 2012. "The Personal Touch: Leaders' Impressions, Costly Signaling, and Assessments of Sincerity in International Affairs." *International Studies Quarterly* 56 (3): 560–73.

Halperin, Eran. 2014. "Emotion, Emotion Regulation, and Conflict Resolution." *Emotion Review* 6 (1): 68–76.

Halperin, Eran, Alexandra G Russell, Carol S Dweck, and James J Gross. 2011. "Anger, Hatred, and the Quest for Peace: Anger Can Be Constructive in the Absence of Hatred." *Journal of Conflict Resolution* 55 (2): 274–91.

Hatemi, Peter K, Rose McDermott, Lindon J Eaves, Kenneth S Kendler, and Michael C Neale. 2013. "Fear as a Disposition and an Emotional State: A Genetic and Environmental Approach to Out-Group Political Preferences." *American Journal of Political Science* 57 (2): 279–93.

Hayes, Andrew F, Dietram A Scheufele, and Michael E Huge. 2006. "Nonparticipation as Self-Censorship: Publicly Observable Political Activity in a Polarized Opinion Climate." *Political Behavior* 28 (3): 259–83.

Hogg, Michael A, Dominic Abrams, and Marilynn B Brewer. 2017. "Social Identity: The Role of Self in Group Processes and Intergroup Relations." *Group Processes & Intergroup Relations* 20 (5): 570–81.

Hogg, Michael A, and Scott A Reid. 2006. "Social Identity, Self-Categorization, and the Communication of Group Norms." *Communication Theory* 16 (1): 7–30.

Holmes, Marcus. 2013. "The Force of Face-to-Face Diplomacy: Mirror Neurons and the

Problem of Intentions." *International Organization* 67 (4): 829–61.
Incitez China. 2015. "Weibo MAUs 198 Mln in Q1 2015, up 38% YoY." China Internet Watch. https://www.chinainternetwatch.com/13364/weibo-q1-2015/.
Jasper, James M. 2018. *The Emotions of Protest*. Chicago ; London: The University of Chicago Press.
Johnston, Alastair Iain. 2001. "Treating International Institutions as Social Environments." *International Studies Quarterly* 45 (4): 487–515.
Keltner, Dacher, and Jonathan Haidt. 1999. "Social Functions of Emotions at Four Levels of Analysis." *Cognition & Emotion* 13 (5): 505–21.
Kertzer, Joshua D, and Thomas Zeitzoff. 2017. "A Bottom-up Theory of Public Opinion about Foreign Policy." *American Journal of Political Science* 61 (3): 543–58.
Keys, Barbara, and Claire Yorke. 2019. "Personal and Political Emotions in the Mind of the Diplomat." *Political Psychology* 40 (6): 1235–49.
King, Gary, Jennifer Pan, and Margaret E Roberts. 2013. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review* 107 (2): 326–43.
———. 2017. "How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, Not Engaged Argument." *American Political Science Review* 111 (3): 484–501.
Lake, David A., and Robert Powell, eds. 1999. *Strategic Choice and International Relations*. Princeton, N.J: Princeton University Press.
Larson, Deborah Welch, and Alexei Shevchenko. 2014. "Russia Says No: Power, Status, and Emotions in Foreign Policy." *Communist and Post-Communist Studies* 47 (3): 269–79.
Lynch, Daniel C. 2015. *China's Futures: PRC Elites Debate Economics, Politics, and Foreign Policy*. Stanford, California: Stanford University Press.
Mackie, Diane M, Thierry Devos, and Eliot R Smith. 2000. "Intergroup Emotions: Explaining Offensive Action Tendencies in an Intergroup Context." *Journal of Personality and Social Psychology* 79 (4): 602.
Mackie, Diane M, and Eliot R Smith. 2017. "Group-Based Emotion in Group Processes and Intergroup Relations." *Group Processes & Intergroup Relations* 20 (5): 658–68.
Mackie, Diane M, Eliot R Smith, and Devin G Ray. 2008. "Intergroup Emotions and Intergroup Relations." *Social and Personality Psychology Compass* 2 (5): 1866–80.
Masterson, Michael. 2022. "Humiliation and International Conflict Preferences." *The Journal of Politics* 84 (4): 874–88. https://doi.org/10.1086/715591.
McDermott, Rose. 2004. "The Feeling of Rationality: The Meaning of Neuroscientific Advances for Political Science." *Perspectives on Politics* 2 (4): 691–706. http://journals.cambridge.org/abstract_S1537592704040459.
———. 2014. "The Body Doesn't Lie: A Somatic Approach to the Study of Emotions in World Politics." *International Theory* 6 (3): 557–62.
Mercer, Jonathan. 2014. "Feeling Like a State: Social Emotion and Identity." *International Theory* 6 (3): 515–35. https://doi.org/10.1017/S1752971914000244.
Moravcsik, Andrew. 1997. "Taking Preferences Seriously: A Liberal Theory of International Politics." *International Organization* 51 (4): 513–53.
N. D. 2012. "Anti-Japan Protests: Outrage, to a Point." *The Economist*, September.

https://www.economist.com/analects/2012/09/17/outrage-to-a-point.
Neuman, Yair, and Yochai Cohen. 2014. "A Vectorial Semantics Approach to Personality Assessment." *Scientific Reports* 4 (1): 1–6.
Nummenmaa, Lauri, Jussi Hirvonen, Riitta Parkkola, and Jari K Hietanen. 2008. "Is Emotional Contagion Special? An fMRI Study on Neural Systems for Affective and Cognitive Empathy." *Neuroimage* 43 (3): 571–80.
Palmer, Glenn, Roseanne W McManus, Vito D'Orazio, Michael R Kenwick, Mikaela Karstens, Chase Bloch, Nick Dietrich, Kayla Kahn, Kellan Ritter, and Michael J Soules. 2021. "The MID5 Dataset, 2011–2014: Procedures, Coding Rules, and Description." *Conflict Management and Peace Science*, 0738894221995743.
People's Daily Press. 2017. "People's Daily (1946–Present)." OriProbe Information Services. https://www.oriprobe.com/peoplesdaily.shtml.
Powell, Robert. 2017. "Research Bets and Behavioral IR." *International Organization* 71: S265–77.
Quek, Kai, and Alastair Iain Johnston. 2018. "Can China Back down? Crisis de-Escalation in the Shadow of Popular Opposition." *International Security* 42 (3): 7–36.
Reilly, James. 2012. *Strong Society, Smart State: The Rise of Public Opinion in China's Japan Policy*. Contemporary Asia in the World. New York: Columbia University Press.
Renshon, Jonathan, Julia J Lee, and Dustin Tingley. 2017. "Emotions and the Micro-Foundations of Commitment Problems." *International Organization* 71: S189–218.
Renshon, Jonathan, and Jennifer S Lerner. 2012. "The Role of Emotions in Foreign Policy Decision Making." *Encyclopedia of Peace Psychology*, 313–17.
Roberts, Margaret E. 2018. *Censored: Distraction and Diversion Inside China's Great Firewall*. Princeton, New Jersey: Princeton University Press.
Robinson, Darrel, and Marcus Tannenberg. 2019. "Self-Censorship of Regime Support in Authoritarian States: Evidence from List Experiments in China." *Research & Politics* 6 (3): 2053168019856449.
Sasley, Brent E. 2011. "Theorizing States' Emotions." *International Studies Review* 13 (3): 452–76.
Shen, Xiaoxiao, and Rory Truex. 2021. "In Search of Self-Censorship." *British Journal of Political Science* 51 (4): 1672–84.
Simmons, Beth A. 2002. "Capacity, Commitment, and Compliance: International Institutions and Territorial Disputes." *Journal of Conflict Resolution* 46 (6): 829–56.
Smith, Eliot R, Charles R Seger, and Diane M Mackie. 2007. "Can Emotions Be Truly Group Level? Evidence Regarding Four Conceptual Criteria." *Journal of Personality and Social Psychology* 93 (3): 431.
Steinberg, Blema S. 1996. *Shame and Humiliation: Presidential Decision Making on Vietnam*. Pittsburgh, Pa: University of Pittsburgh Press.
Su, Alice. 2019. "As Trade War Escalates, Chinese Remember 'National Humiliation'." *LA Times*, May. https://www.latimes.com/world/la-fg-china-trade-war-tariffs-colonialism-humiliation-20190513-story.html.
Tausczik, Yla R, and James W Pennebaker. 2010. "The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods." *Journal of Language and Social Psychology* 29 (1): 24–54.

Turner, John C. 1988. *Rediscovering the Social Group: A Self-Categorization Theory*. 1. publ. in paperback. Oxford: Blackwell.

Wallace, Jeremy L., and Jessica Chen Weiss. 2015. "The Political Geography of Nationalist Protest in China: Cities and the 2012 Anti-Japanese Protests." *The China Quarterly* 222: 403–29. https://doi.org/10.1017/S0305741015000417.

Waltz, Kenneth N. 1986. "Political Structures." In *Neorealism and Its Critics*, edited by Robert O. Keohane, 70–97. New York: Columbia University Press.

Wan, Vanessa. 2018. "The Ultimate Guide to Sina Weibo: The Largest Micro-Blogging Platform in China." *Dragon Social*, April. https://www.dragonsocial.net/blog/chinese-social-media-weibo-and-twitter-comparison/#comments.

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.

Weiss, Jessica Chen. 2014. *Powerful Patriots: Nationalist Protest in China's Foreign Relations*. New York, NY: Oxford University Press.

———. 2019. "How Hawkish Is the Chinese Public? Another Look at 'Rising Nationalism' and Chinese Foreign Policy." *Journal of Contemporary China* 28 (119): 679–95.

Wendt, Alexander. 1999. *Social Theory of International Politics*. 1st ed. Cambridge University Press. https://doi.org/10.1017/CBO9780511612183.

Wilkerson, John D, and Andreu Casas. 2017. "Large-Scale Computerized Text Analysis in Political Science: Opportunities and Challenges." *Annual Review of Political Science* 20 (1).

Wong, Seanon S. 2019. "Stoics and Hotheads: Leaders' Temperament, Anger, and the Expression of Resolve in Face-to-Face Diplomacy." *Journal of Global Security Studies* 4 (2): 190–208.

Wu, Guoguang. 1994. "Command Communication: The Politics of Editorial Formulation in the People's Daily." *The China Quarterly* 137: 194–211.

Xue, Jia, Junxiang Chen, Chen Chen, Chengda Zheng, Sijia Li, and Tingshao Zhu. 2020. "Public Discourse and Sentiment During the COVID 19 Pandemic: Using Latent Dirichlet Allocation for Topic Modeling on Twitter." *PloS One* 15 (9): e0239441.

Yang, Stephanie. 2022. "As China Shuts Out the World, Internet Access from Abroad Gets Harder Too." *LA Times*, June. https://www.latimes.com/world-nation/story/2022-06-23/china-great-firewall-foreign-domestic-virtual-censorship.

Zeitzoff, Thomas. 2014. "Anger, Exposure to Violence, and Intragroup Conflict: A 'Lab in the Field' Experiment in Southern Israel." *Political Psychology* 35 (3): 309–35.

Zhang, Qi, Yeyun Gong, Ya Guo, and Xuanjing Huang. 2015. "Retweet Behavior Prediction Using Hierarchical Dirichlet Process." In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.

Zhao, Suisheng. 2004. *A Nation-State by Construction: Dynamics of Modern Chinese Nationalism*. Stanford, California: Stanford University Press.