

Emotional Contagion and Foreign Policy Preferences: The Case of National Humiliation on Chinese Social Media

Michael Masterson^{*}

21 September, 2020

Abstract

Researchers have linked a variety of emotions including anger, hatred, fear, and humiliation to inter-group conflict attitudes. National humiliation's effect on conflict attitudes has particularly come under scrutiny because of its influence on China's and Russia's foreign policies. How do group-relevant emotions spread socially to become politically relevant? Do emotional expressions of humiliation cause public advocacy of hostile foreign policy positions and shape foreign policy debates? To answer these questions, I use supervised machine learning to measure the prevalence of narratives of national humiliation as well as the foreign policy positions taken in a large (more than 1.6 billion posts), representative data set of Chinese social media posts. I find that national humiliation spreads through emotional contagion both when group members express national humiliation and when they express support for policies associated with national humiliation. I further find that Weibo posts invoking narratives of na-

^{*}This paper is a work in progress and results are subject to change. The author would like to thank Jialei Zhu for her excellent research assistance hand labeling the Weibo posts.

tional humiliation are more likely to support using military force, maintaining disputed territorial claims, and raising trade barriers.

The ongoing power transition between the United States and China has led scholars and observers alike to question whether the two most powerful countries on the planet are “destined for war” (Allison 2017). Recently, US national security strategy has openly labeled China a “strategic competitor” for the first time. China’s construction of artificial islands and increased military deployments in disputed areas of the East and South China Seas have further increased tensions between the US and China (Shen et al. 2019).

Because the stakes of a potential conflict are immense, understanding the “latent risk factors” that could increase its probability is critical (Stein 2019, 57). Scholars have argued that national humiliation narratives are just such a factor because they increase the probability disputes will escalate (Callahan 2010, 96; Wang 2012). In the early 1990s, the Chinese government launched the Patriotic Education Campaign to spread the narrative that China had a history of greatness but was humiliated by foreign powers. This narrative of national humiliation has since become ubiquitous in China through a multimedia campaign that exploits the power of the state to control the news and education system (Wang 2012). While the Chinese government is likely more insulated from public opinion pressure than a democratic government, the regime’s reliance on nationalism for legitimacy means that it cannot ignore citizens’ passions (Zhao 2004, 264–65).

Moreover, the influence of humiliation on international hostility is not limited to China, which makes understanding it even more important. Otto von Bismarck provoked the Franco-Prussia war, which killed nearly 200,000 people, by editing a telegram to give the appearance that the Prussian king had been “snubbed” by the French ambassador (O’Neill 1999, 143). The “humiliation” this loss caused led to French territorial conquest in Africa (Barnhart 2020, 113–16). Nixon used humiliation to justify continuing to sink costs in the US’ unpopular war in Vietnam, saying, “I know that a peace of humiliation for the United States would lead to a bigger war or surrender later” (Steinberg 1996, 204). Further, scholars point to “humiliation”

as a motivation for Russia’s 2014 annexation of Crimea (Larson and Shevchenko 2014).

Neither is humiliation the only emotion researchers have linked to foreign policy positions. Researchers have shown that other emotions including anger (Zeitsoff 2014), hatred (Halperin et al. 2011), and fear influence conflict attitudes (Hatemi et al. 2013). While I focus on the case of national humiliation in China, the theories of emotional contagion examined here apply generally to emotional influences on inter-group conflict.

Two key questions about emotions and international conflict preferences remain. First, how do individually experienced emotions such as humiliation spread socially and influence foreign policy debates? Second, does how this influence function in practice?

I argue 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. Drawing on work from psychology, 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). These mechanisms create a feedback effect where expressions of national humiliation lead to further such expressions from other group members. Moreover, because the symbols and narratives surrounding group-based emotions cumulate over time into an “emotional culture” that alters the way group members interpret subsequent events (Barsade and Knight 2015, 26), exposure to fellow group members who express support for policy positions associated with national humiliation will also increase the probability that individuals experience and express national humiliation.

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 they spread (Halperin 2014, 74; Hall and Ross 2015, 848; Kertzer and Zeitsoff 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 most likely to make a difference during a crisis (Weiss 2019).

To examine both emotional contagion as well as the relationship between expressed national humiliation and support for hostile foreign policies, I use a representative data set of Weibo (a Twitter-like Chinese social media platform) posts, which contains over 1.6 billion posts created from 2009 to 2014.¹ I use supervised machine learning to label these posts based on whether they contain national humiliation narratives and whether they express support for hostile foreign policies.

I find both that national humiliation posts increase the number of national humiliation posts on subsequent days and that posts advocating hostile foreign policy options increase the number of national humiliation posts on subsequent days. This suggests that national humiliation does spread through emotional contagion and that policy options become associated with national humiliation, furthering its spread. Additionally, I find that posts containing narratives of national humiliation are more likely to express support for using military force, maintaining disputed territorial claims, and erecting trade barriers. Further, posts containing national humiliation narratives elevate the number of posts expressing support for these policies on subsequent days. These results suggest that surges of posts about national humiliation push social media users to express support for costly, hostile policies.

This paper is structured as follows. Section 1 details a theory of how policy relevant emotions diffuse through social media to influence foreign policy debates and derives hypotheses about expressions of national humiliation and support for hostile policies. Section 2 discusses data collection and coding. Section 3 descriptively analyzes the data, and section 4 formally tests the hypotheses. Section 5 addresses concerns about censorship and regime commentators. The final section offers concluding comments.

¹Section 5 addresses potential concerns about censorship and regime commentators.

1 National Humiliation and Foreign Policy Preferences

A growing body of work makes the case that emotions affect international decisions. Emotionally aroused actors make different decisions than dispassionate actors (Loewenstein 1996; McDermott 2004; Sayette et al. 2008). In particular, research has shown that emotions alter decisions about international bargaining and confrontation. For example, Renshon, Lee, and Tingley (2017) show that emotional arousal can inhibit strategic thinking, causing actors to make sub-optimal decisions in bargaining. Halperin et al. (2011) find that individuals who are high in hatred of the out-group are less likely to compromise in negotiations when angry. Zeitzoff (2014) finds that anger increases the propensity to punish among Israelis living in areas more exposed to rocket fire. Hatemi et al. (2013) find that individuals high in social fear have more negative opinions about out-groups.

In particular, investigators have focused on humiliation, defined as “the deep dysphoric feeling associated with being, or perceiving oneself as being, unjustly degraded, ridiculed, or put down—in particular, one’s identity has been demeaned or devalued” (Hartling and Luchetta 1999, 264). Scholars have offered various mechanisms through which humiliation might increase international hostility. Some theorize that humiliation motivates revenge (Lindner 2006; Löwenheim and Heimann 2008; Wang 2012). Others argue that humiliation frames the humiliator as an enemy (Leidner et al. 2010; Wang 2012; Leidner, Castano, and Ginges 2013). Additional explanations emphasize the willingness humiliated of actors to use violence in order to regain status, dignity, or pride (Saurette 2006; Fattah and Fierke 2009; Barnhart 2017, 2020). There is also a substantial literature that associates humiliation with violent extremism (Pyszczynski, Motyl, and Abdollahi 2009; McCauley 2016; Webber et al. 2018).

Much of this research has focused on national humiliation, which is humiliation that *individuals* feel on behalf of the nationality that they identify with. This is a type of group-based emotion, which are emotions that are dependent on an individual’s self-categorization as a member of an identity group (Goldenberg et al. 2016, 118–19; Bar-Tal, Halperin, and

De Rivera 2007, 442; Mackie and Smith 2017, 659). These emotions are triggered when an individual is exposed to experiences of other group members that are perceived as relevant for the group as a whole (Goldenberg et al. 2016, 121). China has often been a focus of this research because of the prevalence of narratives of nation humiliation in the media and education system (Callahan 2010, 96; Wang 2012).

Recent experiments have distinguished support for a specific mechanism through which humiliation influences international hostility at the individual level (Masterson 2020). They show that humiliation motivates the support of hostile policies by decreasing individual's sensitivity to the cost of conflict. While this helps clarify a mechanism through which humiliation may function, it leaves open two questions that are needed to understand humiliation's influence on foreign policy outcomes rather than just individual preferences. First, how do these individual-level emotions affect debates about foreign policy in a way that becomes salient to policy makers. Secondly, do the results found on convenience samples (American MTurkers and college students) shed light on foreign policy debates outside of the lab?

I answer these questions by, first, building a theory of how national humiliation diffuses through social media to influence public opinion about foreign policy that is closely watched by policy makers. Second, I derive specific hypotheses about the kinds of policy positions national humiliation will be associated with on Chinese social media.

1.1 Social Media and Emotional Contagion

Social media is a key avenue for diffusing political opinion, emotion, and social pressure. As King, Pan, and Roberts (2017) argue, “scholars and policymakers should focus considerably more effort on the Chinese Internet and its information environment, which is a contested virtual space, one that may well be more important than many contested physical spaces. The relationship between the government and the people is defined in this space, and thus the world has a great interest in what goes on there” (497). Johnston (2017) calls for further exploration of Chinese nationalism using time series measures that include “online

opinion” because, even if online opinion is “less representative” than survey opinion, it is “more immediately salient for political leaders” (42).

Research on social media can shed light on what Kertzer and Zeitzoff (2017) refer to as “mesofoundations,” the social network context, in which citizens are embedded, which are vital to understanding public opinion and foreign policy because these networks are sources of information to citizens and produce pressure for citizens to conform (546). Because the deliberation over how group-related events should be perceived takes place in the public sphere, these deliberations can have a significance influence on the emotions that group identifiers experience (Bar-Tal, Halperin, and De Rivera 2007, 447). The social sharing of emotions and the connection of emotions to a goal can allow them to spread and motivate action (Frijda 2007, 189, 192). Understanding how individual emotional reactions spread is an important part of bridging the gap between individual-level behavioral research on emotions and foreign policy preferences and foreign policy outcomes themselves (Hall and Ross 2015, 848).

I draw on group-based emotions theory, which finds that group members tend to converge emotionally over time (Barsade and Knight 2015, 30, 36), to explain contagion of policy-relevant emotions through social media. Much of the work on emotional contagion focuses 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). Despite this, there are reasons to expect that even when interaction is only text-based, emotional contagion may still take place. Theoretical mechanisms that could lead to text-based 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, where individuals take cues from others as to what emotions are appropriate to express (Festinger 1954), and exposure to emotional words triggering facial muscles related to emotional expression (Feroni and Semin 2009). This contagion is not simply faking emotions to fit in. Instead it causes individuals “to actually feel the emotion, effectively catching the emotion of the other

person” (Barsade and Knight 2015, 24).

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). Researchers have experimentally induced emotional contagion in a small-group setting via emotional emails sent by a group member (Cheshin, Rafaeli, and Bos 2011) and in a large-scale social media context by manipulating users Facebook feeds (Kramer, Guillory, and Hancock 2014).

With a few important exceptions, research on the *politics* of emotion contagion is rather limited. Hall and Ross (2015) theorize that emotional reactions relevant to international politics can be activated through “bottom-up” processes where preexisting identities are activated, “horizontal” processes where emotion spreads across individuals, and “top-down” forces where the pressure to conform to the emotional norm of the group leads to emotional convergence (857–58). They further write of the possibility of “affective waves” that involve feedback effects where individuals spontaneously express emotions and this results in social pressure for others to follow suit, creating a wave of emotional activity that reaches a zenith and then dissipates over time (859). In more recent work, they encourage scholars to attend to the construction of group-based emotions through political discourse, the way the salience of these constructions changes dynamically over time, and the effect of these constructions on encouraging particular political positions (Hall and Ross 2019, 1369).

I also draw on appraisal theory to explain how emotional frames encountered on social media influence emotional experience. According to appraisal theory, 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). If a group identifier receives information from a fellow group member that frames an event as nationally humiliating, then

they 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 the group in conflict, they tend to increase their experience of that emotion (Goldenberg et al. 2016, 123). This process is not necessarily conscious. 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).

My account of emotional contagion on social media begins with individual group members spontaneously sharing content laden with group-based emotional frames. 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. They then become more likely to share such content on social media themselves.

H1 Past expressions of national humiliation on social media will increase future expressions of national humiliation.

Moreover, particular policy positions will become coded over time as reflecting particular emotions, so that expressions of support for these policies will lead to further emotional contagion, even absent an explicit emotional frame. Repeated appraisal of events with certain themes as evoking a particular group-based emotion can cumulate into a lasting sentiment that increases the probability that future events with these themes will also be appraised as evoking this emotion (Halperin 2014, 69). Indeed, the CCP often explicitly links current foreign policy humiliations to China’s humiliating past (Wang 2012). Symbols, narratives, and group categories build up over time into a “social context” or “emotional culture” that shapes the way group members impart meaning to emotional events (Bar-Tal, Halperin, and De Rivera 2007, 444; Barsade and Knight 2015, 26; Koschut 2017). This

can lead individuals to be more likely to emotionally respond in particular ways to mass media cues (Bar-Tal, Halperin, and De Rivera 2007, 446). This could indicate that previous encounters with social media posts that frame disputes, such as China’s territorial disputes, as nationally humiliating could lead group identifiers to experience national humiliation when subsequent events occur in these disputes, even absent an explicit emotional frame. While theories of “feedback loops” in group-based emotions are quite developed in psychology, they are largely ignored in political science, and empirical investigation of these dynamics over time is limited (Barsade and Knight 2015, 38–39).

In addition to the explicit expression of national humiliation in the context of particular policy positions, individuals also look at the behavior and negotiation positions of other group members for information about these group members’ emotional states (Mackie and Smith 2017, 662; Cheshin, Rafaeli, and Bos 2011). For these reasons, Chinese citizens encountering social media posts advocating policies associated with national humiliation may attribute the emotional experience of national humiliation to the creator of the post, even if the post itself lacks any mention of national humiliation. If this is the case, then posts advocating hostile foreign policy positions should themselves increase the frequency of national humiliation posts for the same reasons (pressure to conform, emotional norms, and social cues) that national humiliation posts feedback on themselves.

H2 Expressions of support for policies linked to narratives of national humiliation will increase future expressions of national humiliation.

1.2 National Humiliation and Policy Positions

Previous research has found that humiliation leads individuals to support hostile foreign policy actions by decreasing their sensitivity to the cost of these actions (Masterson 2020). This section derives some policies that individuals might be more likely to support. Humiliated individuals should be more likely to support using military force. Using military force is always a hostile action, and generally, it is considered costly (Fearon 1995). In the Chinese

case, 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).²

H3 Posts expressing national humiliation will be more likely advocate using or threatening to use military force.

Another costly, hostile foreign policy is maintaining territorial disputes. These disputes are costly both because of the risk that they could escalate to war (Vasquez and Henehan 2001), 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. The Patriotic Education Campaign “presented the communist state as the defender of China’s national interests, pride, and *territorial integrity* [my emphasis]” (Zhao 2004, 231). Chinese students study maps of China’s territorial claims, and the curriculum stresses the necessity of a strong military to regain these territories (Hughes 2017, 64). A best selling nationalist book, *The China That Can Say No* argued that China should “take back Taiwan by force, no matter what the cost, and confront the United States and Japan without consideration of the consequences” (Zhao 2004, 161–62). While this book was later suppressed by the government, it is anecdotal evidence that suggests the power of national humiliation to increase support for territorial aggrandizement in a way that is insensitive to the cost of this policy.

H4 Posts expressing national humiliation will be more likely to support China’s disputed territorial claims.

The final foreign policy position that I will investigate is in the economic rather than military realm. While lower on the hostility scale than military force, adopting economic barriers

²Although invoking the economic costs of war increased support for backing down, suggesting that citizens may evaluate economic and military costs differently.

against other countries can certainly be thought of as hostile. The current trade war between the US and China exemplifies this. Anti-trade measures are also costly to economic welfare (Alston, Kearn, and Vaughan 1992). Anecdotal evidence suggests that Chinese citizens connect China’s current trade war with the US 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). Could national humiliation lead individuals to be more supportive of trade barriers?

H5 Posts expressing national humiliation will be more likely to advocate adopting trade restrictions.

1.3 Does Public Opinion Matter in China?

Popular nationalism has the potential to constrain Chinese foreign policy. Nationalist criticisms of China’s policies as being too dovish towards the United States and Japan threaten Chinese Communist Party (CCP) legitimacy (Zhao 2004, 33). Zhao writes that “nationalism is a double-edged sword” because, while it allows the CCP to claim legitimacy, it could produce “serious backlash” if citizens demand more hawkish policies than the regime wishes to take (Zhao 2004, 264–65). This backlash could even threaten regime stability. While Zhao points out that the CCP has so far resolved international crises peacefully, despite its nationalist rhetoric, he also notes that with the rise of information exchange over the internet and leader accountability, it is not clear how long this can continue (272).

That China’s leaders get a daily briefing on online social media activity emphasizes the importance they ascribe to public opinion in this realm (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).³ 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

³In contrast, see (Quek and Johnston 2018, 10).

if the probability of losing office due to citizen dissatisfaction is lower for autocrats, they are sensitive to it because they face harsher consequences, such as death and imprisonment, than democratic leaders who lose office (Debs and Goemans 2010).

2 The Data

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).⁴ The Fudan NLP Group’s website appears to have been taken down, and the data set is no longer publicly available. The data was originally collected by randomly selecting 200 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 Chinese online opinion. In total, the data set has 1,679,006,899 posts from 2.07 million users.⁵ The earliest post is from August 13, 2009, and the latest post is from March 12, 2014. To my knowledge, this data set has not previously been used to study politics.

It is common in political science research, when dealing with data of this size, to select a small subset for analysis to ease computational challenges (Pan and Siegel 2020; Cairns and Carlson 2016). This is perfectly acceptable, but I want to make the case that the time and effort needed to analyze the full set of data, adds value. In addition to the obvious benefit of increasing statistical power, analyzing the full data set has three other advantages. First, labeling the full data set based on the content of the posts creates a valuable data contribution for others to use. Second, it reduces the problem of the “garden of forking paths” because there are nearly infinite subsets of the data that could be selected for analysis but only one set that contains all of the data (Gelman and Loken 2014). Subset analysis could be

⁴The 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.).

⁵This is after removing 203,341 posts that had no textual content. Presumably these posts only contained image or emojis content, which was not collected.

sensitive to subset selection and potentially allow the researcher to cherry-pick a hypothesis confirming subset. Lastly, using the full data set provides sufficient density for time series analysis at the day level. This fine-grained day unit of analysis is more in line with social media cycles that happen rapidly.

Possible concerns about censorship and regime commentators are discussed below, but an even more fundamental concern is whether the data set itself is real or fabricated. This problem often troubles research on politics in authoritarian countries. I follow best practices from King, Pan, and Roberts (2017). 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 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.

2.1 Coding

To code the data, I first assembled two lists of keywords. The first was designed to select posts that could contain narratives of national humiliation. The second was designed to select posts that could contain the policy variables of supporting the use of military force, China's territorial claims, or the adoption of trade barriers. Keyword selection is used because each variable is relatively rare in the overall data set (see the appendix section 6 for the full lists of keywords). It is not a good use of human coder time to code thousands of posts, none of which contain the variable being analyzed. Further, machine learning will perform poorly if some categories are too rare (Chang and Masterson 2020).

After constructing the initial lists of keywords, I supplement them with computer suggested keywords (King, Lam, and Roberts 2017). Next, I iteratively eliminate words that select a high proportion of documents that do not contain the sought after variables until

the frequency of sought after variables in the selected documents is high enough to be useful for human coding and machine labeling while remaining highly inclusive.

I then randomly selected 2,499 of the keyword-selected posts for national humiliation to be coded by a hypothesis-blind research assistant. The code book defines **national humiliation** posts as posts representing the nation of China as being humiliated by a foreign humiliator (see the appendix for the code book, which contains full definitions and examples for each category).

Another set of 3,499 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 military force against another country.⁶ Posts were coded as containing the **territory** variable if they mention 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 selected for hand coding and checked intercoder reliability. For each category, the research assistant’s coding agreed with my own at least 97% of the time.⁷

The human-labeled posts then became the training sets for separately trained Support Vector Machine (SVM) models, which were used to machine label the remaining keyword-selected posts for each variable.⁸ See the appendix for additional information on preprocessing and feature selection.

Posts that were not selected by keyword as possibly containing a variable were assumed to not contain that variable (meaning they were coded as 0 for that variable). Machine labeling

⁶Initially, 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.

⁷Krippendorff’s α for each category was as follows: 0.922 for national humiliation, 0.784 for military force, 0.889 for territory, and 1 for trade barrier.

⁸I follow guidance from Chang and Masterson (2020) for supervised learning with unbalanced data. These models were created using Python 3 and the machine learning package scikit-learn.

the 35,522,659 keyword-selected posts was computationally intensive. To accomplish it, I trained and deployed my SVM models on a Google AI Platform instance.

To evaluate the models, I use a 0.8/0.2 train/test split. The models are highly accurate, each model classified at least 96% percent of the posts in the withheld test set correctly (approximating the 97% intercoder agreement among the human coders). See the appendix for additional information on model performance.

I also include an `income inequality` variable both as a control for political issues in general 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 it is an issue ranked as the second biggest concern among Chinese survey respondents, so it should be salient (Chubb 2014).⁹ I use the same phrase to measure income inequality as was used in the survey: “rich-poor disparity” (贫富分化) (Chubb 2014, 22).

3 Descriptive Analysis

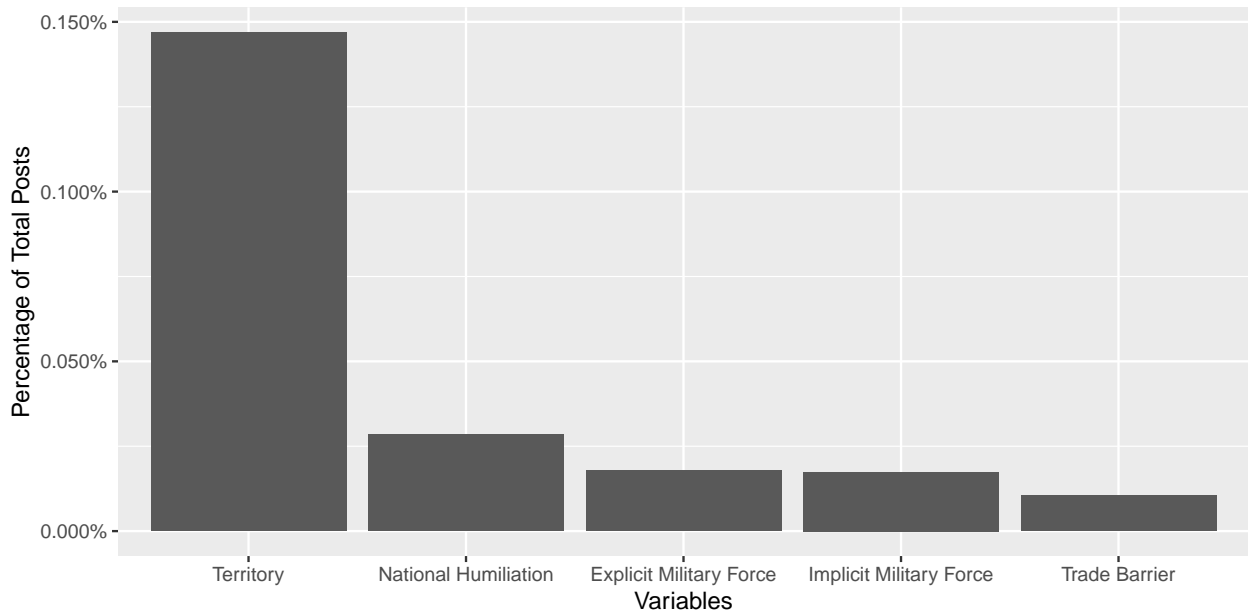
Because the prevalence of different foreign policy positions on Chinese social media is interesting in itself, this section descriptively analyzes the data. The following section formally assesses the hypotheses. Figure 1 shows the estimated percentage of total posts that fall into each category.¹⁰ These are consistent previous work, which has found that only a small minority of Weibo posts contain political content (Chang and Masterson 2020). However,

⁹I choose this instead of the number 1 issue of corruption because national humiliation posts often attribute China’s historic vulnerability to humiliation to historic corruption, for example, of the Qing dynasty. I needed a variable that measured a political issue for which there was no theoretical reason to expect it to covary with national humiliation.

¹⁰This estimate is calculated by taking the proportion of *hand-coded* posts in each category in the posts selected by keywords for that category and dividing by the proportion of keyword selected posts to total posts. Because the hand-coded posts were randomly selected from the posts selected by keyword to potentially contain content for each category, the proportion of the hand-coded sample should represent an unbiased estimate of the proportion in the keyword selected sample. Assuming that posts that were not selected by keyword do not belong in these categories, dividing by the proportion of keyword selected posts to total posts will give an unbiased estimate of the proportion of total posts from these categories.

the volume of posts in these categories is not trivial. The trade barrier category, which is the smallest category with only 0.0105% of posts, represents 176,787 posts in the overall sample. The territory category, which is the largest category with 0.1469% of posts, represents 2,467,334 posts in the overall sample. Further, the approximately 2 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

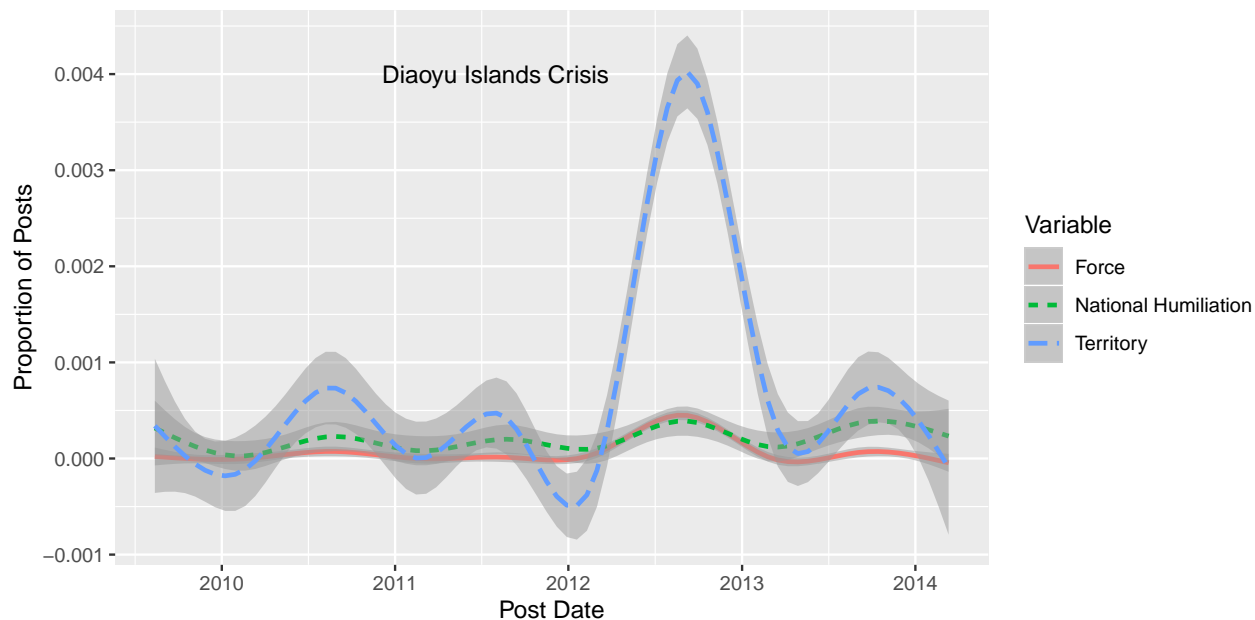


Posts that mention China’s territorial disputes are by far the most common. A distant second are posts that use national humiliation narratives. This means that most posts about territorial disputes do not explicitly invoke national humiliation. Posts advocating specific policies are the rarest, with explicit and implicit support for using military force both more common than those advocating trade barriers (implicit and explicit military force were coded separately in the hand coding but combined for machine labeling). This is somewhat surprising because trade barriers are a much less escalatory step, so one might expect they would receive wider support than more extreme measures like using force. It could be that if a citizen is emotionally aroused enough to post about foreign policy online,

then they are also excited enough to reach for more forceful positions.

Figure 2 shows how the proportion of posts in the overall sample that mention territorial disputes, contain national humiliation narratives, or advocate using force vary over time.¹¹ 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. This set off the largest anti-Japan protests in China since relations between China and Japan were normalized, suggesting that the machine-coded data are picking up real variation.

Figure 2: Posts Over Time
Variables smoothed using a generalized additive model.



4 Hypothesis Tests

First I examine the data at the post level. This is the most direct way to test whether the same posts that mention national humiliation are also more likely to mention territorial disputes, advocate using military force, or advocate trade barriers (H3–5). However, since

¹¹Here I use the machine-labeled posts to get a data set that is dense enough to show meaningful variation over time.

variables are present in the same post, the post-level analysis cannot determine the direction of the causal relationship between national humiliation and the other variables nor can it examine the temporal dynamics among the variables to test for feedback effects (H1–2). For this reason, the second part of this analysis aggregates the data to the day level and uses time series analysis to examine the causal order among the variables.

4.1 Post-level Analysis

4.1.1 Modeling Strategy

To analyze the post-level data, I use linear probability models.¹² I control for the number of characters in the social media post to control for the possibility that longer posts are more likely to contain both national humiliation narratives as well as each policy variable. I also control for the discussion of income inequality as a way to control for whether posts that discuss one political topic are more likely to discuss other political topics.

4.1.2 Results

Table 1 shows the results for each policy variable both with and without controls. In each case, the coefficient on national humiliation is positive and statistically significant, supporting H3–5. The coefficients can be interpreted as the increase in probability that a post containing a national humiliation narrative will contain the dependent variable compared with posts that do not contain national humiliation narratives. The substantive effect for humiliation on the territory variable is by far the largest, with national humiliation being associated with a 10.82 percentage point increase in the probability a post mentions a territorial dispute. National humiliation also increases the probability a post advocates using force and trade

¹²Because the post-level data contains 1,679,006,899 observations, the analysis cannot be conducted on a conventional computer. I conduct all of the post-level analysis using Google’s BigQuery ML (Machine Learning) platform. To save funds, I only calculate standard errors and evaluate significance for the treatment. This should not detract from the usefulness of the results because control variables do not have causal interpretations anyway (Keele, Stevenson, and Elwert 2020). I cannot use logistic regression because BigQuery logistic regression model objects do not currently return the information necessary to calculate standard errors.

barriers by between one and two percentage points each. The substantive importance of this effect becomes clear when it is compared to the baseline probabilities of posts advocating military force and trade barriers, which are increased by approximately 142 and 542 times respectively.

Table 1: Post-level Results

	Force	Territory	Trade Barrier	Force	Territory	Trade Barrier
Intercept	0.0001	0.0011	0.0000	-0.0001	0.0002	0.0000
National Humiliation	0.0143**	0.1089**	0.0119**	0.0142**	0.1082**	0.0119**
	(0.0000)	(0.0001)	(0.0000)	(0.0000)	(0.0001)	(0.0000)
Post Length				0.0002	0.0010	0.0000
Income Inequality				-0.0003	0.0007	0.0000

N = 1,679,006,899 Weibo posts. All results rounded to the fourth decimal place. Post Length is re-scaled to hundreds of characters. Standard errors and significance only shown for treatment. * indicates $p < 0.1$ ** indicates $p < 0.05$.

4.1.3 Placebo Test

A potential concern is that discussion of one political issue may correlate 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 run the analysis with income inequality, which is a political issue that should not theoretically be related to national humiliation, as the dependent variable. Table 2 shows the results.

Table 2: Post-level Placebo Test

	Income Inequality	Income Inequality
Intercept	0.00001	-0.00001
National Humiliation	0.00000	-0.00002**
	(0.00001)	(0.00001)
Post Length		0.00002

N = 1,679,006,899 Weibo posts. All results rounded to the fifth decimal place. Post Length is measured in hundreds of characters. Standard errors and significance only show for treatment. * indicates $p < 0.1$ ** indicates $p < 0.05$.

The relationship between national humiliation and income inequality is not significant in the bi-variate model. In fact, the coefficient rounds to 0 at the fifth decimal place. In the

model that accounts for post length, the relationship just barely reaches significance, but it is in the opposite of the direction expected if political discussion confounded the results. Further, the substantive effect is negligible. According to the model, posts that discuss national humiliation are -0.002 percentage points less likely to mention income inequality. Overall, the placebo tests increase confidence that the results in the previous section are not spurious to general political discussion.

4.2 Causal Order

Because the post-level analysis examines whether national humiliation and policy positions appear in the same post, it is unable to establish the temporal relationships among the variables. If national humiliation is causing the other variables, then, at a minimum, it should temporally precede them. Further, the theory makes predictions about the temporal dynamics among the variables. Specifically, it predicts that previous posts about national humiliation will increase subsequent humiliation posts (H1) and that posts about policies tied to national humiliation will increase future national humiliation posts (H2). This section aggregates the data to the day-level of analysis to examine these temporal relationships.

4.2.1 Modeling Strategy

Here, I use vector autoregression (VAR), which allows the inclusion of multiple endogenous variables in order to assess how shocks in one variable affect the others (Freeman, Williams, and Lin 1989). The endogenous variables I include are national humiliation, military force, territory, trade barrier, income inequality, total posts, and total length. Each of these variables (aside from total length, which is measured in characters) is measured as a count of posts per day. I select the lag order of the VAR model based on the Hannan–Quinn information criterion, which selects 9 lags as optimal (Hannan and Quinn 1979; Box-Steffensmeier et al. 2014 Ch. 4)¹³

¹³Specifically, I iteratively increase the maximum lag order examined until the information criterion selects a model with an order lower than the maximum lag order. Because non-stationary variables can pose prob-

First, I evaluate Granger causality for each endogenous variable on every other endogenous variable. For VAR, a variable Granger causes another if at least one of the lags of the first variable is significant in the outcome equation for the second variable (Floyd 2005, 31).¹⁴ Second, I examine the orthogonalized impulse response functions (OIRF), which show how each variable responds to a one standard deviation increase in another variable over time to learn about the magnitude and direction of these effects (Freeman, Williams, and Lin 1989, 847).¹⁵

In the results shown below, I do not include exogenous controls. However, in the appendix I show the Granger causality results when exogenous dummies are included for days of the week, months, and years.¹⁶ The robustness of the results to the inclusion of these fixed effects should increase confidence that the results are not spurious to unmeasured events that elevate discussion of both national humiliation and the policy variables although it is impossible to completely eliminate this concern.

4.2.2 Results

The results in Table 3 show whether each variable Granger causes the others. National humiliation both Granger causes (supporting H3–5) and is Granger caused by military force, territory, and trade barrier (supporting H2). Further, national humiliation Granger causes itself, supporting H1. However, national humiliation does not Granger cause the placebo variable of income inequality. This indicates that, as predicted, national humiliation is in a dynamic relationship with military force, territory, and trade barrier, in which past days’

lems for VAR (Box-Steffensmeier et al. 2014 Ch. 4), before selecting the model’s lag length, I examine each variable for stationarity. Kwiatkowski–Phillips–Schmidt–Shin (KPSS) tests fail to reject the null hypothesis that each variable is non-stationary (Kwiatkowski et al. 1992), so I difference each variable once. Both KPSS tests and augmented Dickey–Fuller (ADF) tests of the differenced variables indicate they are stationary.

¹⁴This is distinct from a block exogeneity test, which is a joint test of whether a variable Granger causes any of the other endogenous variables in the VAR. Throughout the paper I will use Granger causality to refer to whether past values of one variable predict future values of another rather than in the block exogeneity sense.

¹⁵I use bootstrapping to calculate the OIRF, so they are not sensitive to the order of the variables.

¹⁶Because the impulse response plots only depend on the endogenous variables, they are unaffected by whether exogenous variables are used.

counts of posts on each of these variables affect subsequent days' counts of the others. To understand the direction and magnitude of these effects, I turn to the orthogonalized impulse response plots shown in Figure 3.

Table 3: Does The Variable in the Row Granger Cause the Variable in the Column?

	Humiliation	Force	Territory	Trade	Inequality	Total Posts	Total Length
Humiliation	Yes	Yes	Yes	Yes	No	No	No
Force	Yes	Yes	Yes	Yes	No	No	No
Territory	Yes	Yes	Yes	Yes	No	No	No
Trade Barrier	Yes	Yes	Yes	Yes	No	No	No
Income Inequality	Yes	No	No	No	Yes	Yes	Yes
Total Posts	Yes	Yes	Yes	Yes	No	Yes	Yes
Total Length	Yes	Yes	Yes	Yes	Yes	No	Yes

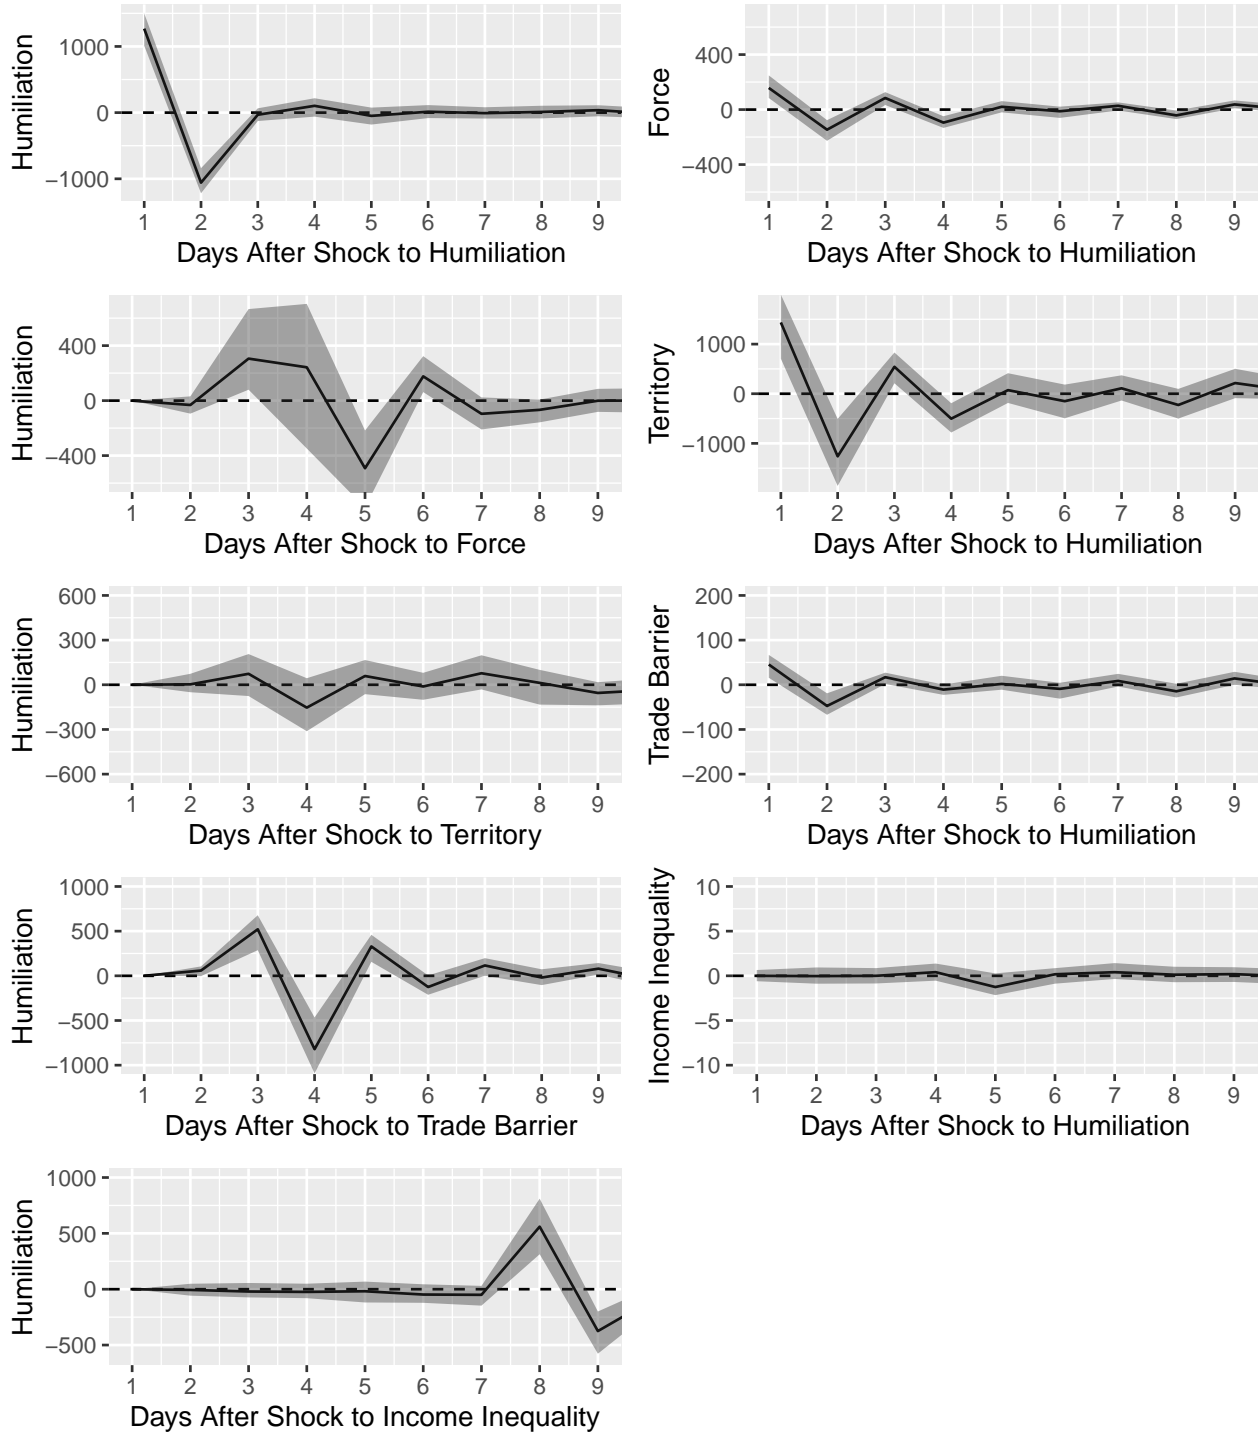
Yes indicates that the Granger causality test was significant at the 0.05 point level.

The OIRF plots show that a one standard deviation increase in posts about national humiliation results in an increase of approximately 1,250 posts about national humiliation, 200 posts about military force, 1,500 posts about territory, and 50 posts about trade barriers on the *following* day. On the second day after a shock to humiliation posts, posts about humiliation, force, territory, and trade barriers decline. Finally, over the next few days, posts about these variables stabilize to their normal levels. Lending confidence that this effect is not a result of general cycles of political discussion, the effect of a shock to humiliation posts on the placebo variable of income inequality is essentially 0 across all 9 days. The increase in posts about national humiliation on the day following a shock to national humiliation posts provides evidence for H1.

Shocks in the policy variables also affect future national humiliation posts. A shock to force increases national humiliation posts by about 350 posts three days later. This elevation continues on the fourth day before declining and ultimately stabilizing. The confidence interval for the territory variable's affect on national humiliation in the 9 days following a shock to territory always covers 0. This suggests that despite being a Granger cause of national humiliation, the substantive effect of posts about territory on posts about national humiliation is negligible.

Figure 3: Orthogonalised Impulse Responses

Orthogonalised impulse response to a 1 standard deviation shock with bootstrapped 95% confidence intervals shown. Results computed based on a single vector autoregression model with 9 lags. The impulse responses for total posts and post length as well as impulse response pairs not involving national humiliation are not shown.



The OIRF showing the effect of a shock to trade barrier on national humiliation shows a similar pattern to the OIRF for shocks to force on humiliation. On the third day after the shock, posts about national humiliation increase by about 500 before declining and then stabilizing. Overall, these results provide support for H2, with advocacy of two of the three hostile policy responses resulting in a substantial increase in national humiliation posts on subsequent days.

The OIRF showing the effect of shocks to income inequality on national humiliation is puzzling. The response of the humiliation variable is essentially zero for the first seven days. On the eighth day, humiliation increases by 500 posts before declining and then stabilizing. It is tempting to dismiss this change on the eight day, but I cannot rule out that there could be an effect after an 8 day lag or that income inequality could be proxying for some other variable that has such an effect.

Overall, these results suggest that there is an effect of national humiliation posts on subsequent posts that advocate each kind of costly, hostile foreign policy, even accounting for potential endogenous feedback effects between posts advocating these policies and posts containing national humiliation narratives. This provides further evidence for H3–5. Moreover, the fact that national humiliation narratives and these hawkish policy positions Granger cause each other combined with the fact that (with the exception of the effect of posts about territory disputes on humiliation posts) the OIRF analysis shows these effects to be substantial indicates that there is a feedback effect among these variables. In other words, posts about national humiliation encourage posts advocating hostile foreign policies, which, in turn, encourage more posts about national humiliation. This supports H1–2.

5 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 normal 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. 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.

The posts were collected the same day they were created. This decreases influence of human censors who remove posts after they are posted on the data. While this does not allow collection of posts that were blocked by automated keyword censorship, this kind of censorship is “not sophisticated or very successful, and therefore much of content filtering is done by hand” (Roberts 2018, 154). Further, 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.

Work on censorship has found 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). Online expression about international relations receives “relatively minimal interference from censors” compared to other political issues (Chubb 2014, 56). This makes it less likely posts containing the variables analyzed here would be specifically targeted.

Self censorship is also a concern. Citizens may be unwilling to express their true feelings because they fear state punishment. While there is no doubt this is a particularly prevalent concern for some groups in China, such as Uyghurs who face mass arbitrary detention and public opinion leaders who face extra scrutiny, for most Chinese the threat of punishment for online comments is not very credible and online criticism of the regime is common (Roberts 2018, 13). When Roberts examined the impact of censorship on Weibo users, contrary to the idea of prevalent self censorship, she found that being censored actually increased the likelihood the user would post about sensitive topics in the future (Roberts 2018, 117). Similarly, Pan and Siegel (2020) find that imprisoning dissidents in Saudi Arabia based on their social media posts tended to increase anti-regime social media content overall.

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). Regime commentators avoid taking stances for or against certain policies. 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 trying to defend its policies (Roberts 2018, 190–91).

Perhaps there is an exception to this for national humiliation posts? If regime posts were particularly likely to contain narratives of national humiliation, then we would expect bursts of regime commentator activity around dates like September 18, the anniversary of the Manchurian Incident, which is unofficially referred to as “National Humiliation Day.” However, King, Pan, and Roberts (2017) show regime commentators stay flat at near zero during this period and instead spike around periods when the regime fears that the potential for collective action is high (488, 496). This would lead to the expectation that posts made by the regime would simply add noise to the data by increasing the volume of irrelevant posts.

However, it is possible that “cheerleading” regime commentator posts may mention na-

tional 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). This *biases against* my hypotheses because I am predicting that posts containing national humiliation narratives will be more likely to advocate hostile policies. If instead of avoiding policy discussion, regime commentators took positions consistent with those taken in state-run media, then these posts would espouse positions less hawkish than typical online opinion, which would also bias against my hypotheses (Chubb 2014, 42–43).

Even in the absolute worst case where censorship and/or regime commentators not only bias the results but bias them in favor of my hypotheses, this analysis would still be illuminating. The consequence of this kind of bias would be that the interpretation of the results would change. Instead of concluding that national humiliation drives people to take particular kinds of foreign policy positions, the results would suggest that the Chinese government believes that national humiliation narratives on social media can lead people to support particular foreign policy positions and acts on this belief either by having its commentators create such posts or by censoring alternative views.

6 Conclusion

This article has examined the conditions under which emotions spread socially and affect policy debates. Specifically, I examined whether Weibo posts that contained narratives of national humiliation increased the probability of future posts about national humiliation. I find that they do, which suggests that the tendency for group identifiers to conform to emotional expressions from fellow group members allows national humiliation to spread. An important implication of this is that China’s narratives of national humiliation may be self-sustaining and even self-escalating. I further find that posts advocating policies positions associated with national humiliation promote future posts about national humiliation, even controlling

for whether or not these posts explicitly mention national humiliation. This suggests that Chinese citizens tend to attribute the emotion of national humiliation to individuals advocating these positions, even when these individuals do not explicitly convey this emotion. Further, this leads others to conform, expressing the emotion that they believe their fellow group members are expressing. This suggests that separating debates about Chinese foreign policy from national humiliation and its hostility promoting effects may be difficult because policy positions have become coded as reflecting national humiliation, which can trigger listeners to experience this emotion.

This article has also provided evidence that national humiliation affects policy debates on social media, which both provides external validity to experimental work on national humiliation and hostility as well as expands understanding of its practical effects. Specifically, Weibo posts that contained narratives of national humiliation were more likely to express support for each of the following costly, hostile policies: using military force, maintaining disputed territorial claims, and raising trade barriers. Further, posts about national humiliation increase the amount of posts advocating each of those policies on subsequent days. The absence of an increase in the placebo variable of income inequality due to national humiliation in both designs helps allay concerns that the results could be confounded by users' underlying propensities to discuss political topics and/or cycles of political discussion.

These findings have important implications for research on foreign policy preferences and political psychology. They show how emotional narratives can influence policy preferences. Further, they establish that these narratives can spread through social media to influence the policy options citizens consider in the future. This is particularly important in the case of China where the state 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 a feedback effect between posts containing national humiliation narratives and posts advocating hostile

foreign policies. This raises the specter of a vicious cycle that could magnify the effect of state media and nationalist groups pushing national humiliation narratives. The revelation that emotional narratives diffuse through social media networks also increases understanding of how emotions spread to become politically important, which is an under-studied aspect of political behavior (Hall and Ross [2015](#)).

Future research should examine the effect of social media discussions on policy itself. Examining this link is a challenge, partly because foreign policy change is rare, making data on foreign policy coarse and sparse in comparison with social media data. A convincing study of this kind might need to examine the internal deliberations of policy makers, to more fully capture variation in their positions and whether they mention popular pressure from social media as a policy justification. A second avenue for future research is examining how other kinds of emotional narratives diffuse through social media to influence policy positions.

7 References

- Allison, Graham T. 2017. *Destined for War: Can America and China Escape Thucydides's Trap?* Boston: Houghton Mifflin Harcourt.
- 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. 2017. "Humiliation and Third-Party Aggression." *World Politics* 69 (3): 532–68.
- . 2020. *The Consequences of Humiliation: Anger and Status in World Politics*. <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.
- Bar-Tal, Daniel, Eran Halperin, and Joseph De Rivera. 2007. "Collective Emotions in Conflict Situations: Societal Implications." *Journal of Social Issues* 63 (2): 441–60.
- Box-Steffensmeier, Janet M, John R Freeman, Matthew P Hitt, and Jon CW Pevehouse. 2014. *Time Series Analysis for the Social Sciences*. Cambridge University Press.
- 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>.
- Cheshin, Arik, Anat Rafaeli, and Nathan Bos. 2011. "Anger and Happiness in Virtual Teams: Emotional Influences of Text and Behavior on Others' Affect in the Absence of Non-Verbal Cues." *Organizational Behavior and Human Decision Processes* 116 (1): 2–16.
- Chubb, Andrew. 2014. *Maritime Awareness: Popular Attitudes Towards China's Maritime Disputes in the East and South China Seas*.
- 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>.
- Fattah, K., and K. M. Fierke. 2009. "A Clash of Emotions: The Politics of Humiliation and Political Violence in the Middle East." *European Journal of International Relations* 15 (1): 67–93. <https://doi.org/10.1177/1354066108100053>.
- 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.
- Floyd, John E. 2005. "Vector Autoregression Analysis: Estimation and Interpretation." *University of Toronto*. <https://conservancy.umn.edu/bitstream/handle/11299/54931/1977-91.pdf?sequence=1&isAllowed=y>.

- 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.
- Freeman, John R, John T Williams, and Tse-min Lin. 1989. "Vector Autoregression and the Study of Politics." *American Journal of Political Science*, 842–77.
- Frijda, Nico H. 2007. *The Laws of Emotion*. Mahwah, N.J: Lawrence Erlbaum Associates.
- Fudan NLP Group. n.d. "Shùjù Jí Xìnxī." Accessed November 22, 2019. <http://sma.fudan.edu.cn/datainfo/weibo.html>.
- Gelman, Andrew, and Eric Loken. 2014. "The Statistical Crisis in Science: Data-Dependent Analysis—a "Garden of Forking Paths"—Explains Why Many Statistically Significant Comparisons Don't Hold up." *American Scientist* 102 (6): 460–66.
- 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.
- 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>.
- Hall, Todd H, and Andrew AG Ross. 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.
- 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.
- Hannan, Edward J, and Barry G Quinn. 1979. "The Determination of the Order of an Autoregression." *Journal of the Royal Statistical Society: Series B (Methodological)* 41 (2): 190–95.
- Hartling, Linda M., and Tracy Luchetta. 1999. "Humiliation: Assessing the Impact of Derision, Degradation, and Debasement." *Journal of Primary Prevention* 19 (4): 259–78. <http://link.springer.com/article/10.1023/A:1022622422521>.
- 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.
- Hughes, Christopher R. 2017. "Militarism and the China Model: The Case of National Defense Education." *Journal of Contemporary China* 26 (103): 54–67.
- Incitez China. 2015. "Weibo MAUs 198 Mln in Q1 2015, up 38% YoY." China Internet Watch. <https://www.chinainternetwatch.com/13364/weibo-q1-2015/>.
- Johnston, Alastair Iain. 2017. "Is Chinese Nationalism Rising? Evidence from Beijing." *International Security* 41 (3): 7–43.
- Keele, Luke, Randolph T Stevenson, and Felix Elwert. 2020. "The Causal Interpretation of Estimated Associations in Regression Models." *Political Science Research and Methods* 8 (1): 1–13.

- 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.
- King, Gary, Patrick Lam, and Margaret E Roberts. 2017. "Computer-Assisted Keyword and Document Set Discovery from Unstructured Text." *American Journal of Political Science* 61 (4): 971–88.
- 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.
- Koschut, Simon. 2017. "The Structure of Feeling–Emotion Culture and National Self-Sacrifice in World Politics." *Millennium* 45 (2): 174–92.
- Kramer, Adam DI, Jamie E Guillory, and Jeffrey T Hancock. 2014. "Experimental Evidence of Massive-Scale Emotional Contagion Through Social Networks." *Proceedings of the National Academy of Sciences* 111 (24): 8788–90.
- Kwiatkowski, Denis, Peter CB Phillips, Peter Schmidt, Yongcheol Shin, and others. 1992. "Testing the Null Hypothesis of Stationarity Against the Alternative of a Unit Root." *Journal of Econometrics* 54 (1): 159–78.
- 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.
- Leidner, Bernhard, Emanuele Castano, and Jeremy Ginges. 2013. "Dehumanization, Retributive and Restorative Justice, and Aggressive Versus Diplomatic Intergroup Conflict Resolution Strategies." *Personality and Social Psychology Bulletin* 39 (2): 181–92.
- Leidner, Bernhard, Emanuele Castano, Erica Zaiser, and Roger Giner-Sorolla. 2010. "In-group Glorification, Moral Disengagement, and Justice in the Context of Collective Violence." *Personality and Social Psychology Bulletin* 36 (8): 1115–29.
- Lindner, Evelin. 2006. *Making Enemies: Humiliation and International Conflict*. Contemporary Psychology. Westport, Conn: Praeger Security International.
- Loewenstein, George. 1996. "Out of Control: Visceral Influences on Behavior." *Organizational Behavior and Human Decision Processes* 65 (3): 272–92.
- Löwenheim, Oded, and Gadi Heimann. 2008. "Revenge in International Politics." *Security Studies* 17 (4): 685–724. <https://doi.org/10.1080/09636410802508055>.
- 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.
- Masterson, Michael. 2020. "Humiliation and International Conflict Preferences." *R&R at the Journal of Politics*. https://michaelrmasterson.com/static/docs/conflict_humiliation_masterson.pdf.
- McCauley, Clark. 2016. "What Comes After ISIS? A Peace Proposal." *Perspectives on Terrorism* 10 (4): 62–66.
- 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.
- 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>.
- 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.
- O’Neill, Barry. 1999. “Honor, Symbols and War.” *Ann Arbor, Michigan University Press*.
- Pan, Jennifer, and Alexandra A Siegel. 2020. “How Saudi Crackdowns Fail to Silence Online Dissent.” *American Political Science Review* 114 (1): 109–25.
- Pyszczynski, Tom, Matt Motyl, and Abdolhossein Abdollahi. 2009. “Righteous Violence: Killing for God, Country, Freedom and Justice.” *Behavioral Sciences of Terrorism and Political Aggression* 1 (1): 12–39.
- 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.
- Renshon, Jonathan, Julia J Lee, and Dustin Tingley. 2017. “Emotions and the Micro-Foundations of Commitment Problems.” *International Organization* 71 (S1): S189–S218.
- Roberts, Margaret E. 2018. *Censored: Distraction and Diversion Inside China’s Great Firewall*. Princeton, New Jersey: Princeton University Press.
- Saurette, Paul. 2006. “You Dissin Me? Humiliation and Post 9/11 Global Politics.” *Review of International Studies* 32 (3): 495–522.
- Sayette, Michael A, George Loewenstein, Kasey M Griffin, and Jessica J Black. 2008. “Exploring the Cold-to-Hot Empathy Gap in Smokers.” *Psychological Science* 19 (9): 926–32.
- Shen, Dingli, Elizabeth Economy, Richard Haas, Joshua Kurlantzick, Sheila A. Smith, and Simon Tay. 2019. “China’s Maritime Disputes.” Council on Foreign Relations. https://www.cfr.org/interactives/chinas-maritime-disputes?cid=otr-marketing_use-china_sea_InfoGuide#!/chinas-maritime-disputes?cid=otr-marketing_use-china_sea_InfoGuide.
- Simmons, Beth A. 2002. “Capacity, Commitment, and Compliance: International Institutions and Territorial Disputes.” *Journal of Conflict Resolution* 46 (6): 829–56.
- Stein, Rachel Michelle. 2019. *Vengeful Citizens, Violent States: A Theory of War and Revenge*. 1 [edition]. New York: Cambridge University Press.
- 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>.
- Vasquez, John, and Marie T Henahan. 2001. “Territorial Disputes and the Probability of War, 1816-1992.” *Journal of Peace Research* 38 (2): 123–38.
- 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.
- Webber, David, Maxim Babush, Noa Schori-Eyal, Anna Vazeou-Nieuwenhuis, Malkanthi Hettiarachchi, Jocelyn J Bélanger, Manuel Moyano, et al. 2018. “The Road to Extremism: Field and Experimental Evidence That Significance Loss-Induced Need for Closure Fosters Radicalization.” *Journal of Personality and Social Psychology* 114 (2): 270.
- 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.
- Zeitsoff, 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.