

Putting Your Money Where Your Mouth Is: Do Expressions of Nationalist Sentiments Signal Loyalty or Genuine Nationalist Feelings?

Charles Chang and Michael Masterson*

22 April, 2022

Abstract

Do nationalist expressions reflect genuine sentiments or simply what citizens believe their government or compatriots want to hear? This is particularly important in context such as economic nationalism in which behaving in ways consistent with such expressions carries costs. Yet, past research has either relied on self reports that cannot rule out cheap talk or outcome measures at higher levels of analysis, such as international trade flows, that do not directly link individual-level nationalist sentiment to behavior. In contrast, we measure both nationalist sentiments and costly behavior on an individual level using supervised machine learning to label nationalist sentiment on Chinese social media and data on the nationality of the brand of device, for example type of cellphone, used to produce social media posts. Further, we recover the causal effect that nationalist sentiment makes citizens less likely to use US and Japanese brands, using an instrumental variable approach.

*The authors contributed equally and would like to thank Loe Bi, Elijah Cao, Tianshu Feng, Ganqi Li, Catherine Lin, Sylvia Ma, Yantao Mei, Chenghan Xue, Daniel Zhu, and Jialei Zhu for their excellent research assistance hand labeling the Weibo posts and device strings.

In a video uploaded to the Chinese website Meipai (美拍) on July 12, 2016, the same day the Permanent Court of Arbitration in the Hague ruled against China’s claims in the South China Sea (Lu and Westcott 2016), a young Chinese man going by going a username meaning ‘plan for a rainy day’ (未雨绸缪) shows a close up of his iPhone, so that the viewer can verify that it is real. Afterwards, he tells the viewer to “watch” (看着) and smashes it with a hammer. Despite his Meipai account only having 68 followers, this video has been viewed over 550,000 times.¹ “Tomorrow, I will buy a Chinese one,” he says. In the video description, he makes clear that his motivation comes from what he views as “America taking the lead to violate/attack Chinese territory” (美国带头侵犯我国领土). He calls on his “brothers and sisters to follow [his example] and smash [their iPhones]” (兄弟姐妹们跟起砸) and concludes his video description with 3 angry face emojis. Throughout the video, he wears a hat with a piece of paper stuck in it that says “those who violate/attack China, even if far away, must be punished/executed” (犯我中华者，虽远必诛).

This was not an isolated incident. Other Chinese citizens also posted videos smashing their iPhones to express their anger over the ruling against China’s claims. Observers attributed this to “nationalist sentiments,” and some of these videos featured the claim that “if you don’t smash it, you’re not Chinese!” (French 2016). Chinese state media eventually condemned the protests of the tribunal’s ruling, which targeted American companies, including KFC, Apple, and McDonald’s (Lu and Westcott 2016). However, this is hardly the first dispute linking Chinese nationalist sentiments with opposition to foreign products. Neither has the Chinese government consistently opposed economic nationalism. For example the Chinese government has adopted restrictions on Australian wine in 2018, tourism to South Korea in 2017, and Salmon from Norway in 2010 amid foreign policy spats with these states (Taylor 2018).

Protesters’ willingness to pay the cost of destroying their phone gives credibility to the sincerity of their nationalist sentiments. In other words, by paying this cost, they put their

¹As of 4/4/2022, the video can still be viewed [here](#).

money where their mouth is in terms of nationalist opposition to the United States. However, relatively few Chinese went as far as smashing their iPhones. Are more typical citizens who may express support for economic nationalism without participating in phone smashing or offline protests expressing genuine nationalist sentiment, or are they simply sending the message that they believe the government or their fellow citizens want them to send?

This also has important theoretical implications for economic nationalism. In particular, can economic nationalists sentiments actually affect individuals' preferences and behavior or are what observers often take to be sentiments that affect motivation merely cheap talk? By economic nationalism, we mean the pursuit of relative economic gains for one's own nation in comparison with other nations. This can manifest in raising trade barriers to the products of other nations or boycotts of these products. Because raising trade barriers and boycotting foreign products pose costs to citizens ([Alston, Kearl, and Vaughan 1992](#)), the genuineness of nationalist sentiments matters for whether citizens are willing to pay these costs. If citizens genuinely feel these sentiments, then they will be willing to pay the cost of missing out on foreign goods. Otherwise, they will maintain nationalist expressions, but these expressions will be cheap talk. In their private economic behavior, they will deviate from economic nationalism.

This also has implications for the broader influence of nationalism on foreign policy and government legitimacy. To the extent that nationalist expressions reflect genuine sentiments, this is consistent with explanations that highlight nationalism's role in regime legitimacy. Further, it would lend support to arguments that these sentiments shape citizens' policy demands on other issues, including territorial disputes ([Fang and Li 2020](#)). To the extent these expressions do not reflect genuine sentiments, it would indicate that the Chinese government's ability to rule rests more on rewarding loyalty and punishing deviance than nationalist legitimacy and that when Chinese citizens demand the government take action in the name of nationalism, such as standing firm on its territorial disputes, these utterances should be viewed more as expressions of loyalty rather than genuine demands that threaten

the government with a loss of support if it does not follow through (Quek and Johnston 2018).

Evaluating whether citizen expressions reflect genuine nationalist sentiments is challenging. In many cases the predictions of theories that argue nationalist expressions are genuine and theories that argue these expressions are motivated by citizen expectations about reward and punishment produce observationally equivalent predictions (Jiang and Yang 2016). Previous research has either relied on self-reports, which while informative cannot rule out that responses reflect cheap talk, or costly behavior measures at higher units of analysis, such as decreased trade volumes between countries after an event inflames nationalists, which cannot directly link individual-level nationalist sentiment to behavior. This raises the possibility that other factors could have driven these changes in trade and risks committing the ecological fallacy by concluding that because a country where nationalists reside decreased trade volumes, nationalists are the ones driving this change (King 1997).

We have adopted a unique research design that allows us to measure both nationalist sentiments and costly consumer behavior on an individual level. We use machine learning to label a large data set of Weibo posts based on whether they contain economic nationalist sentiments. We exploit data on the whether citizens use Chinese brand devices or foreign brands associated with countries towards which nationalists express animosity as a measure of costly economic behavior. If citizens' expressions of economic nationalism are genuine, then these citizens should be less likely to use brands from countries seen as harming China and more likely to use Chinese brands.

However, this analysis shares the limitation of previous observational work on economic nationalism of being unable to address unobservable confounds that vary over time, such as individuals' life experiences. To identify the *causal* effect of nationalist sentiment on individual behavior, we exploit the exogenous timing of the Japanese government's purchase of the disputed Diaoyu/Senkaku Islands as an instrumental variable for observed online nationalist sentiment. While descriptively we find no substantively meaningful relationship between

online nationalist sentiment and device choice, our instrumental variable approach recovers a large causal effect of nationalist sentiment. Citizens expressing nationalist sentiments are less likely to own Japanese or United State brand phones, indicating that nationalist sentiment does drive costly behavior.

This paper proceeds as follows. The first section explains our empirical contribution in the context of economic nationalism research. Section 2 discusses the our data and estimation strategies, including addressing the potential impact of censorship and regime commentators on the data and whether the instrument is truly exogenous. Section 3 presents the results first from a series of descriptive models and then from our instrumental variable models designed to identify the causal effect of economic nationalist sentiment. The final section offers concluding comments.

1 Economic Nationalism

While scholars originally considered economic nationalism only in terms of state policy, its meaning has now expanded to include a wide range of practices including protectionism, dumping of exports, subsidies, transfer of foreign property, discrimination in favor of citizen workers, and consumer antipathy towards foreign products (Baughn and Yaprak 1996, 761; Pickel 2003). The processes through which economic nationalism influences these outcomes share its psychological influence on individual’s cognitive frames, ideologies, and beliefs (Pickel 2003, 121). In particular, researchers theorize that nationalism could lead consumers to decrease their consumption of products from countries that these consumers see as opposed to their nation as a result of past events. This mechanism of economic nationalism’s influence on trade is referred to as “animosity” (Klein, Ettenson, and Morris 1998).

According to the animosity model, nationalist consumers avoid products from countries they see as having harmed their nation in the past because of nationalist sentiment rather than economic considerations, such as product cost or quality (Klein, Ettenson, and Morris

1998, 90). This animosity varies across individuals according to their attitudes and beliefs.

H1 Animosity: Individuals high in economic nationalist sentiment will be *less likely* to consume products with brands associated with foreign countries viewed as having harmed their nation.

In contrast to animosity, nationalists might also have an affinity towards national products either because they want to purchase from individuals that they see as similar and/or as in-group members or because they believe that such purchases will boost the national economy, which they care more about than the economies of other countries (Shimp and Sharma 1987; Mutz and Kim 2017).

H2 Affinity: Individuals high in economic nationalist sentiment will be *more likely* to consume products from domestic brands.

In particular, nationalist movements and protests in China have been associated with calls to boycott Japanese goods in the 1930s, 1985, early 2000s, and 2012 (Reilly 2014, 212). The 2012 protests included calls to buy Chinese rather than Japanese cars. Yet, indicating that Chinese citizens are aware of the costs, others argued that boycotting Japanese goods was too harmful to China's economy (Reilly 2014, 213).

Whether nationalist sentiments are genuine and able to drive costly action is closely linked with the issue of whether the Chinese government's propaganda is intended to shape citizens' beliefs or signal the strength of the regime and deter dissent rather than persuading citizens (Huang 2015, 2018). If propaganda is effective at persuading citizens, then citizens should both express agreement with the propaganda when these expressions are costless as well as engage in costly behavior consistent with these beliefs. If instead, propaganda is effective at signaling government strength rather than persuading citizens, then citizens should express agreement with propaganda when these expressions are costless but refrain from engaging in costly behavior motivated by a belief in the propaganda.

Scholars have provided two kinds of evidence for these animosity and affinity effects. The first set of evidence comes from surveys in which individuals self-report nationalist sentiment

as well as either their consumption practices or trade policy preferences. [Shimp and Sharma \(1987\)](#) find that American survey respondents who score higher on consumer ethnocentrism are more likely to prefer US-made to foreign products. [Klein, Ettenson, and Morris \(1998\)](#) finds that Chinese consumers who report higher animosity towards Japan report owning fewer Japanese products on surveys (96). Survey evidence also shows that individuals living in Chinese regions that have higher trade openness tend to express less nationalism ([Pan and Xu 2018](#)).

The other approach past research has taken is linking decreased trade volumes of foreign products to periods of tension between countries. For example, multiple studies find that the Diaoyu Islands dispute led to decreased purchases of Japanese brands within China ([Barwick et al. 2019](#); [Davis and Meunier 2011](#)). [Heilmann \(2016\)](#) finds similar results examining, in addition to China's boycott of Japan over the island dispute, Muslim boycotts of Danish products following the publication of cartoon depiction of Muhammad in a Danish newspaper, US boycotts of France over France's opposition to the Iraq War, and Turkey's boycott of Israeli goods over the 2014 Gaza War [Heilmann \(2016\)](#). Similarly [Hong et al. \(2011\)](#) find that following Chinese anger over the disruption of the Olympic torch relay in France, French vehicle sells declined in China.

While valuable, both of these approaches have important limitations that raise questions about whether they can establish a causal effect of economic nationalism on individual preferences and behavior. While survey questions can directly ask about consumption, these surveys still rely on self-reports rather than direct measures of costly-consumer behavior. Because survey responses are relatively costless, instrumentally motivated citizens may choose to express nationalism if they think there is even the slightest chance their responses will be able to be traced back to them. In the case of authoritarian countries like China this question is particularly important because the incentive to signal loyalty is much higher than in democratic contexts ([Jiang and Yang 2016](#)). Citizens could expect to be either rewarded for expressing views they perceive to be in line with the government narrative or punished

for not expressing them (Kuran 1997). This punishment could either come from the government or from other citizens who may shun those who fail to express loyalty (Jiang and Yang 2016, 623). Alternatively, citizens who have reported high levels of nationalist sentiment might report owning fewer foreign products in order to appear consistent (Schuman 1981, 27–28). Our point is not that these challenges prevent survey research from shedding light on nationalist sentiments. Instead, they highlight the importance of supplementing our understanding with measures linked to costly behavioral outcomes (Dickson 2011, 63).

However, studies measuring costly behavioral outcomes, such as trade flows, also face limitations. While these studies are important and help clarify the stakes of nationalist boycotts, they do not measure individual-level nationalist sentiments or behavior. This means these studies cannot be used to conclude that nationalist sentiment, which exists on and varies on the individual level, has caused these changes without risking committing the ecological fallacy (King 1997). Further, both types of studies lack an identification strategy to separate the effect of economic nationalism from unobservable confounds that could affect both economic nationalism and consumer behavior. Both nationalist expression and economic behavior are complex and can be affected by such a wide variety of events citizens experience and information they may be exposed to that the assumption of selection on observables, that all these factors have been measured and included in the model, is doubtful (Keele 2015, 321). We describe and carry out a research design below that addresses these issues.

2 Method

2.1 Weibo Data

[This subsection describing the providence of the data is reproduced from Masterson (2022a) for ASN readers' convenience.]

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).² 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 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 1,676,535,827 posts from 2.4 million users.³ The earliest post is from August 13, 2009, and the latest post is from March 12, 2014.

Possible concerns about censorship and regime commentators are discussed in section 2.5, but an even more fundamental concern is whether the data set itself is fabricated. We 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), we 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.

2.2 Operationalizing Nationalist Sentiments

We treat economic nationalist sentiments as a latent variable that cannot be directly observed. However, we can observe *expressions* of these sentiments. Expressions may not correspond exactly with true sentiment because individuals may feel incentive to express

²The 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.).

³This 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.

sentiments they do not feel or to suppress the expression of the sentiments that they do feel based on their beliefs about how the state and their peers will respond to these expressions. We include two separate measures of economic nationalist sentiment to decrease the chance that our results are sensitive to how economic sentiment is measured. The first measure, **trade barrier** includes any posts that advocate boycotting or raising trade barriers against foreign goods, including posts that advocate the substitution of domestic goods for foreign goods.

The second measure, **national humiliation**, indicates whether or not a post contains a national humiliation narrative. National humiliation is a particularly important and prominent nationalist sentiment in China (Wang 2012; Callahan 2010). National humiliation narratives represent a foreign humiliator as inflicting injustice on the Chinese nation. This has ties to economic nationalism as research has found that humiliation decreases sensitivity to the costs of hostile actions against non-nationals (such as more expensive goods as a result of raised tariffs) and that Weibo posts that contain national humiliation narratives are more likely to advocate raising trade barriers (Masterson 2022b, 2022a).

To measure whether a post invoked a national humiliation narrative or advocated raising trade barriers, we use Masterson (2022a)’s coding of the Fudan NLP data set. (For the conference version of this paper we reproduce relevant information about how these posts were coded from Masterson 2022a in the appendix including the full coding dictionary used for human coding as well as the accuracy of the supervised machine learning models trained based on this human-coded data and used to label the remaining posts.)

We choose two countries that modern Chinese nationalists are most likely to view as having harmed the Chinese nation to operationalize brands from countries towards which these nationalists have animosity, the United States and Japan. The United States is commonly viewed by Chinese nationalists as behind harmful events in world affairs. This can be seen everywhere from Chinese nationalists blaming the United States for instigating the war between Russia and Ukraine to the American companies, such as Apple and KFC, bearing the

brunt of Chinese nationalist ire for an international jurisdictional ruling on a case brought by the Philippines to a judicial body located in Europe (McCarthy 2022). Chinese nationalists have a wide range of grievances against the United States, some particularly emotional ones include Truman’s decision to send the US navy to deter a PRC invasion of Taiwan in 1950 and US bombs striking the Chinese embassy in Belgrade in 1999 (Spence 1990, 529; Wang 2012, 170).

Despite the United State’s prominence as a potential great power rival to China, the country that most provokes Chinese nationalist anger is Japan. This is a legacy of both the brutality of the Japanese military in the Second Sino-Japanese War as well as the Chinese government’s decision to emphasize this brutality in the media and education system (Gries 2005; Wang 2012).

Operationalizing domestic brands as brands associated with China is simpler than choosing the brands to examine animosity. The one key decision is whether or not devices associated with Taiwan should be considered domestic. Although Chinese nationalists would certainly claim that Taiwan is a part of China, they also regard the current political authorities in Taiwan with hostility, so it is possible that they do not feel the same affinity towards technology companies prospering under these authorities as they do towards mainland Chinese brands. For this reason, we do not label brands associated with Taiwan rather than the mainland as “Chinese” in our data.

2.3 Coding

Each post has a device string that provides information about the device used to create the post. For example, a post created by an iPhone could have a device string of ‘iPhone 客户端,’ a post created by a Samsung Galaxy phone could have a string of ‘三星 Galaxy,’ and a post created by an iPad might have a string of “iPad 客户端.” We use the device string to code both the nationality of the brand of the device and the price of the device. This is not possible for all device strings. For example, the most common device string is simply

Sina Weibo (新浪微博), which provides no actual information about the device, and we omit these posts (see the Appendix for a discussion of the selection of posts into informative and uninformative device strings). However, we are able to code brand nationality and price for device strings that correspond to 473,781,480 posts. Next to 新浪微博, 'iPhone 客户端' is the second most common device string, making up about 17.5% of posts in the dataset.

Nationality is coded based on the recognizable nationality of the brand of the device rather than the country where the device is manufactured. For example, iPhone is coded as a United State brand even though many iPhones are manufactured in China because iPhone is recognized in China as a US brand. This is consistent with previous work on economic nationalism that has emphasized the importance of a brand that is recognizable to consumers as foreign rather than the actual location of manufacture, which is less salient to consumers ([Barwick et al. 2019](#)).

Device prices are coded based on the price the device was listed at for sell in China in 2014. If we could not locate price information from 2014, we used prices from previous years and calculated the price using a depreciation factor that we constructed by collecting multiple prices for the same devices over different years. More information about how nationality and price were coded is available in the appendix.

2.4 Empirical Strategy

We conduct two sets of analysis. The first is designed to produce descriptive results and uses all of the posts for which the nationality and price of the device are coded. The second estimates the causal effect of nationalist sentiment on the choice of device nationality, using an indicator for days on and after September 11, 2012 when the Japanese government purchased the Diaoyu Islands from their private owners as an instrument for nationalist sentiment ([Mainichi 2012](#)). This second set of analysis only examines posts made one month before and after this date to decrease the chance that other events that took place in the time period analyzed after September 11, 2012 could confound the instrument while also providing

a large enough time span that some individuals will have made decisions to purchase new devices.

In each set of analysis, we show results both with and without the control variables of device price, post length (in characters), and income inequality. We follow [Masterson \(2022a\)](#) by including an income inequality variable that measures whether a post discussed income inequality as a control for non-foreign-policy-related political discussion as well as by controlling for the post length. Both sets of analysis use linear probability models with whether or not the device is a particular nationality as the dependent variable. However, in contrast with the descriptive models, the instrumental variable models use two-stage least-squares.

For the descriptive analysis, we show the results for different combinations of fixed effects because it is interesting to know both whether different people who have different values of nationalist sentiment make different brand choices, which is best examined without using user fixed effects, and whether the same user is more likely avoid American and Japanese phones when making posts that contain nationalist sentiment than when making posts that do not, which is best examined when including user fixed effects. For the causal analysis with the instrumental variable models, we include user fixed effects because we want to control for potential confounds that are constant across users. It is not possible to include day fixed effects in the instrumental variable models because the instrument is a day indicator, so it is perfectly colinear with day fixed effects.

Because each data set is quite large, even the one month window around September 11, 2012 contains 88,906,833 posts, we face the possibility that even trivially small effects could show up as statistically significant. To distinguish between these and effects that are substantively meaningful, we adopt [Hartman and Hidalgo \(2018\)](#)'s framework to select an effect size under which effects will not be considered substantively meaningful and test whether or not we can reject the null hypothesis that the effect we find is equivalent to this negligible effect. We adopt the “fairly conservative” substantively meaningful effect

threshold from [Hartman and Hidalgo \(2018\)](#) of 0.36 standard deviations of the dependent variable ($\pm 0.36\sigma$) (1011). In terms of percentage change in the probability a user’s device has a particular nationality, this means that an estimate must have a 95% confidence interval that does not overlap the following percentage point ranges to be considered a substantively significant effect: $\pm 9.8\%$ for Chinese, $\pm 3.3\%$ for Japanese phones, and $\pm 14\%$ for US phones.⁴

2.4.1 Instrument Justification

To be good instrumental variable, our indicator for days on and after the day the Japanese government officially purchased the Diaoyu Islands must meet a few conditions. These conditions are depicted graphically in the directed acyclic graph shown in Figure 1. Namely, there must be causal link between the instrument and nationalist sentiment, there must be no causal link between the confounds and the instrument, and there must be no causal link from the instrument to the outcome of phone choice except through nationalist sentiment. Each of these conditions is discussed below.

The release of the Japanese government’s arrangements to purchase the islands set off nationalist protests in over 200 cities in China, the largest nationalist protests in recent memory ([Wallace and Weiss 2015](#), 404–5, 413), so it seems likely this information had a large effect on nationalist sentiment. It should affect animosity not only towards Japan, but also towards the United States, which has had a policy that the US-Japan Security Treaty includes these islands since 1972 ([CRS 2021](#)). Figure 2 shows posts containing the trade barrier and national humiliation variables overtime. Supporting the strong relationship between the date of the purchase and nationalist sentiment, both variables have their largest spike immediately after the announcement.

Of course, the overall Diaoyu Islands dispute between China and Japan, which has been going on since the Chinese government officially made claims to the islands in the 1970s, is endogenous to the dynamics of nationalism in China and Japan, so how can this indicator

⁴We structure the equivalence tests this way instead of using the two one-sided test (TOST) because TOST is designed to test differences of means rather than regression results.

Figure 1: Identification Strategy

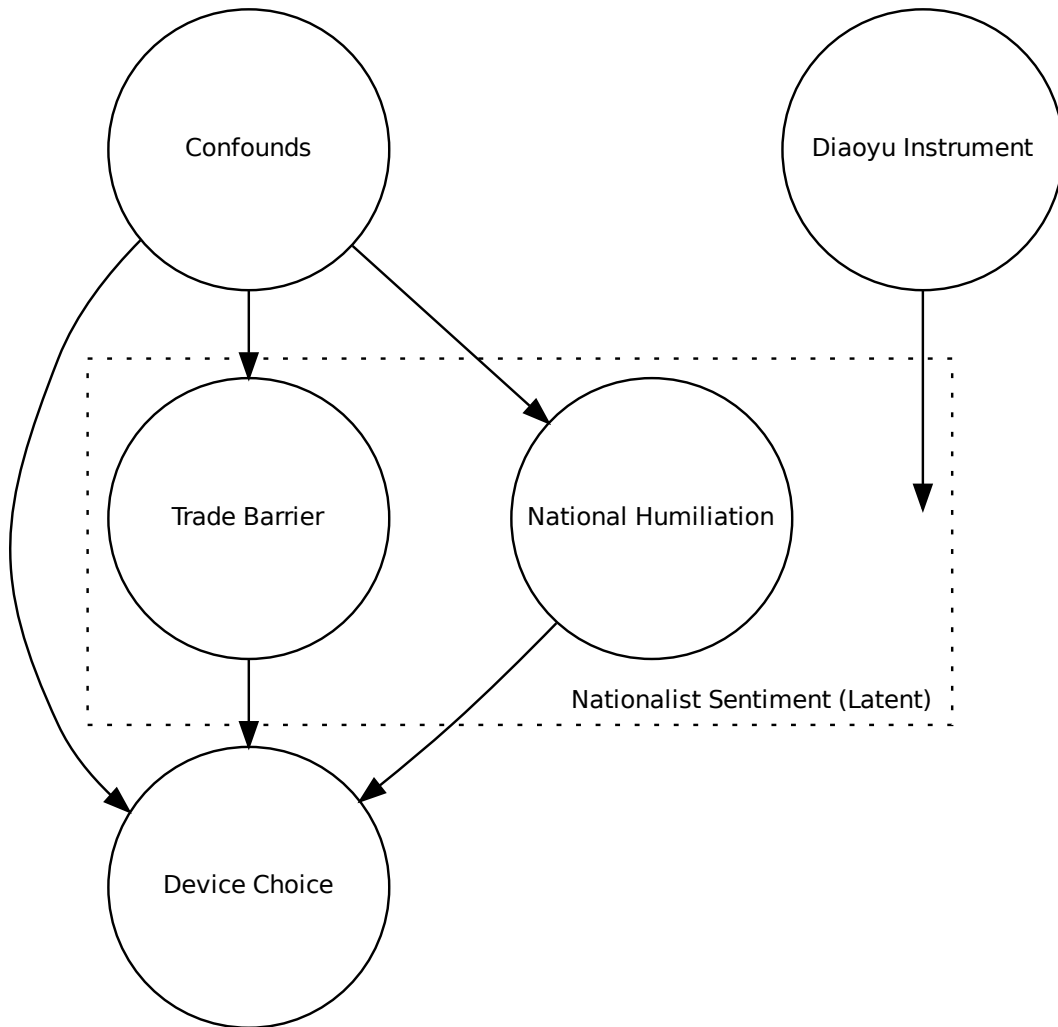



Figure 2: Posts Over Time

Variables smoothed using locally estimated scatterplot smoothing (LOESS).



be exogenous to confounds that affect Chinese nationalist sentiment? Our key identification assumption is that the *timing* of Japan’s decision to purchase the islands is exogenous from nationalist dynamics within China. This decision to purchase the islands was triggered by Japan’s domestic politics as the nationalist former governor of Tokyo, Shintaro Ishihara, planned to purchase the islands on his own initiative. Fearing the consequences, the Japanese central government made its own arrangements to purchase the islands, which it planned to do after diplomatically preparing the ground with China by explaining that this action was intended to prevent Ishihara’s purchase, which the Japanese government expected would be even more inflammatory. However, this plan was derailed and the Japanese government’s hand was forced when the news that Japan was going to purchase the islands leaked to the Japanese newspaper Asahi Shimbun, which revealed the information (Vogel 2019, 391). The exogeneity of the timing of this event from online nationalist dynamics within China is plausible because Chinese citizens posting online would have no ability to influence the leak that came from the Japanese government to a Japanese newspaper about the Japanese government’s plans that determined the timing of the event.

The Japanese government’s purchase of the islands should not affect phone choice through any means other than nationalist sentiment. The islands are uninhabited, so who owns the islands is unlikely to influence device production or consumers’ ability to access different devices in stores (CRS 2021). Further, international trade of electronic devices and the components needed to produce them remained open during this period (the possible temporary halt of rare earths exports from China to Japan occurred in 2010 not 2012, so it is not within the window analyzed here (Klinger 2018, 138)). Further, the purchase of these islands does not reveal any new information about the quality of electronic devices that could affect consumer purchase decisions. Because we only include a one month window before and after the purchase on September 11, 2012, there is less opportunity for events that occurred after Sept. 11, which could affect phone choice, to confound the instrument.

One limitation of this instrument is that it affects both operationalizations of nationalist sentiment, trade barrier and national humiliation. This means that the analysis cannot distinguish which of these two variables does the causal work. However, we conceptualize these variables as observable manifestations of the latent variable of economic nationalist sentiment within China, which is our true independent variable of interest. Including analysis of both is more of a way to ensure that our results do not depend on how this latent concept is measured rather than an attempt to make claims that particular kinds of nationalist sentiment are critical.

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.⁵ If, for example, censors were more likely to remove posts that came from US brand devices, this would lead to underestimating the amount of US brand devices, but

⁵This subsection largely comes from Masterson (2022a).

it would not bias the estimate of the relationship between nationalist sentiment and device choice unless censors were systematically more/less likely to remove posts from US devices that contained nationalist sentiment as opposed to other posts from US devices. Similarly, if regime commentators are more likely to post about nationalism than ordinary citizens, this would increase the amount of nationalist posts as a proportion of total posts, but it would not bias the hypothesis tests unless these posts are also more/less likely to use a particular device type. We 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, We do not discuss the increased restrictions on free discussion imposed after this period.

While it is impossible to directly measure the influence 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, we cross-reference the users from our data with the Weiboscope data set (Fu, Chan, and Chau 2013). The Weiboscope data contains posts that were collected in real time and automatically checked later to examine whether they were removed. The authors were able to exploit 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). Our dataset contains 126,574 users that overlap with the Weiboscope data, allowing us to view which of the 490,277 posts these users made in 2012 were censored. As Roberts (2018) writes, “The Weiboscope data provides an almost ideal dataset to test netizen reactions to censorship because many users

were followed over a relatively long time (125).” 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 our data set (Fu, Chan, and Chau 2013, 43).

In the appendix, we use regression analysis to examine both whether any particular kinds of posts predict the likelihood of being censored and whether nationalist sentiment posts are more or less likely to be censored based on the nationality of the device’s brand. Using Fu, Chan, and Chau (2013)’s measure of censorship, neither device nationality or nationalist sentiment is related to whether or not a post is censored. Similarly, no interaction between the nationalist sentiment variables and any of the phone nationalities is significantly related to censorship.

These findings are consistent with previous work on 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.

Self-censorship is also a concern. Citizens may be unwilling to express their true feelings because they fear 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. 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). While we cannot rule out self-censorship, individual fixed effects should help control for different levels of self censorship across individuals. This would only fail if individuals changed their level of self-censorship of nationalist posts at the same time that they changed the nationality of their device brand and this change was not driven by a change in their true nationalist sentiment.

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 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,” so the trade barrier measure of nationalist sentiment should be relatively resilient to regime commentator influence (King, Pan, and Roberts 2017, 499).


3 Results

3.1 Descriptive Results

Figure 3 shows the descriptive relationship between the trade barrier variable and user phone choice in terms of σ for Chinese, Japanese, and US brand phones. Generally, the results

are relatively null, with no model showing a relationship that is outside of the equivalence range. Further, most of the results have point estimates right around 0 and those that appear statically distinguishable from 0 often change directions when controls are added or removed (for example, the US and Japan results with no fixed effects) and attenuate towards 0 when fixed effects are added.

Figure 3: Descriptive Trade Results




$N = 473,781,480$ Weibo posts. 95% confidence intervals shown. Standard errors are clustered on the user. The control variables include device price, the length of the post in characters, and income inequality. The gray shaded region is the inverted equivalence range defined by $\pm 0.36\sigma$.

Figure 4 shows the descriptive relationship between national humiliation and device choice in terms of σ for Chinese, Japanese, and US brands. As with the trade barrier models, the results show negligible to no relationship with no model finding a relationship outside of the equivalence range. Those brands that show an effect statistically distinguishable from zero attenuate towards zero as fixed effects are added. However, unlike the trade barrier models, the results that most appear to show a relationship (those for US and Chinese phones with no fixed effects) do not flip signs when controls are added and are in the

expected direction, which implies there might be some basis to suggest a descriptive relationship exists across users. However, even in the models with no fixed effects, once controls are included, national humiliation in each case is associated with a less than 2 percentage point change in the probability that a user’s phone is either a Chinese or a US brand, which further suggests that to the extent any descriptive relationship exists, it is relatively minimal.

Figure 4: Descriptive Humiliation Results



$N= 473,781,480$ Weibo posts. 95% confidence intervals shown. Standard errors are clustered on the user. The control variables include device price, the length of the post in characters, and income inequality. The gray shaded region is the inverted equivalence range defined by $\pm 0.36\sigma$.

Overall, it seems that knowing whether a user expresses nationalist sentiment online, either in form of advocating raising trade barriers or national humiliation narratives, is not very informative about the nationality of their device. This implies that, descriptively speaking, users are not putting their money where their mouth is. However, it does not allow us to draw conclusions about the *causal* effect of nationalist sentiment on consumer choice. For that we turn to the instrumental variable models below.

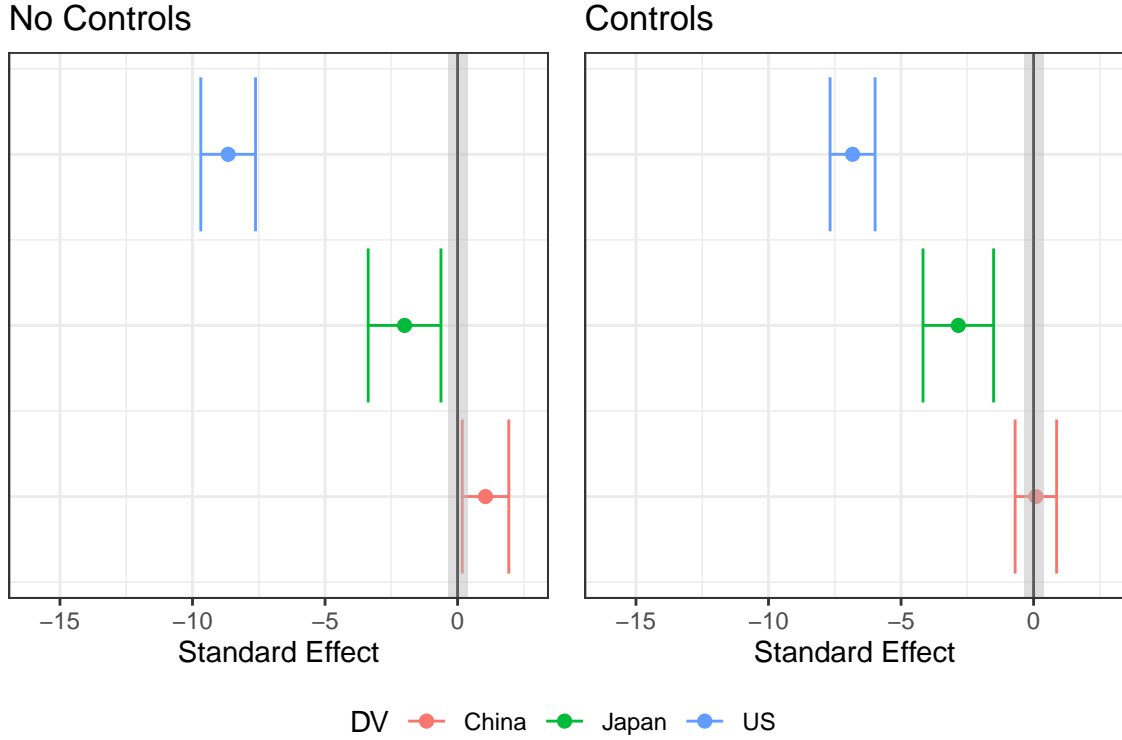
3.2 Diaoyu Instrument

Figure 5 shows the effect of trade barrier on phone choice as estimated in the instrumental variable models. The animosity hypothesis is strongly supported for both the United States and Japan, regardless of whether controls are included in the model. In all cases the effects are statically significant and the confidence intervals fall entirely outside of the equivalence range, allowing us to reject the null hypothesis that the effect is substantively unimportant. The models with controls indicate trade barrier is associated with about a 7σ drop in the probability a user is using a US branded phone and an approximately 3σ decline in the probability a user is using a Japanese phone. In other words, the models predict that users making trade barrier posts who have changed the nationality of their device after September 11, 2012 are approximately 100 percentage points less likely to choose a US and about 26 percentage points less likely to choose a Japanese brand.

In contrast, the affinity hypothesis (H2) is not supported. While the model with no controls finds a statically significant effect in the hypothesized direction, the confidence interval overlaps the equivalence range, implying we cannot reject the null hypotheses that the effect of trade barrier on whether or not individuals use Chinese brands is not substantively significant. Further casting doubt on H2, once controls are added, the point estimate of the effect is approximately zero and the effect is no longer statistically significant.

Figure 6 shows the effect of national humiliation on phone choice as estimated in the instrumental variable models. As with the trade barrier results, the national humiliation results provide strong evidence for the animosity hypothesis (H1) for both US and Japanese brand devices. For both brands, regardless of whether controls are included, national humiliation has a negative effect that is both substantively and statically significant. The models with controls imply national humiliation is associated with an approximately 11σ decrease in the probability of using a US brand device and about a 4σ decline in the probability of using a Japanese brand device. This equates to a decrease in the chance of choosing a US brand of approximately 100 percentage points and a Japanese brand of about 41 percentage

Figure 5: Instrument for Trade Barrier Results




$N = 88,906,833$ Weibo posts for the models without controls and $57,054,341$ for the models with controls. 95% confidence intervals shown. All models include user fixed effects and cluster standard errors on the user. The control variables include device price, the length of the post in characters, and income inequality. The gray shaded region is the inverted equivalence range defined by $\pm 0.36\sigma$.

points. Further consistent with the trade barrier results, these models do not support H2 as in no case can the null hypothesis that the affinity effect is substantively insignificant be rejected, and the effect of national humiliation on Chinese brands becomes statistically insignificant with a point estimate of close to 0 once controls are added.

Overall these results strongly support H1 but do not provide evidence of H2. This implies that when individuals experience nationalist sentiment they are likely to switch away from US and Japanese brands that they have animosity towards but are not necessarily more likely to choose a Chinese brand. Instead, they may be substituting other foreign brands towards which they have less animosity.

Figure 6: Instrument for National Humiliation Results



$N = 88,906,833$ Weibo posts for the models without controls and $57,054,341$ for the models with controls. 95% confidence intervals shown. All models include user fixed effects and cluster standard errors on the user. The control variables include device price, the length of the post in characters, and income inequality. The gray shaded region is the inverted equivalence range defined by $\pm 0.36\sigma$.

4 Conclusion

Our results show that while descriptively Chinese nationalists do not seem to be putting their money where their mouth is in terms of avoiding electronic devices with US and Japanese brands, the effect of economic nationalist sentiment on device choice is actually quite large once a research design able to control for unobserved confounds is applied. The same cannot be said for nationalist affinity for domestic products of which we find little evidence.

This has important implications both for the study of economic nationalism broadly and Chinese nationalism specifically. Our study provides strong evidence that economic nationalist sentiments are causally associated with the outcome of animosity towards foreign brands that are associated with countries viewed as having harmed the nation. This study provides several missing empirical links. First, it is, to our knowledge, the first study that combines

a measure of economic nationalist sentiment on the individual level (avoiding the ecological fallacy) with a costly behavioral measure of an individual’s economic choices (avoiding the tendency of individuals professing nationalism to self-report nationalist consumption on surveys due to consistency effects (Schuman 1981, 27–28)). Second, our instrumental variable approach allows us to be much more confident that we have recovered an accurate estimate of the causal effect of economic nationalist sentiment on individual behavior than existing observational studies, which rely on statistical controls of observable variables to adjust for confounds. The importance of this is dramatized by our descriptive results that show that without such an identification strategy, we would have been unable to recover a substantively significant effect of nationalist sentiment.

More concretely, our results show that the frequent examples of individuals who profess to be nationalists yet consume brands associated with countries they express animosity towards should *not* dissuade us of the power of economic nationalism to shape individual preferences and behavior. In the Chinese context, a common example is the open secret that top Chinese leaders frequently send their children to elite private universities in the United States despite their “anti-American rhetoric” (Higgins and Fan 2012). This should not convince us that their nationalist expressions do not reflect genuine nationalist sentiments that have real effects on their preferences and behavior. We cannot draw this conclusion because we do not observe the counterfactual rate at which these leaders, holding unobserved confounds constant, would send their children to the United States absent these sentiments. Of course other factors such as an individual’s socioeconomic status and position (or lack thereof) in the party matter, but our results provide evidence that economic nationalist sentiments have a large impact on individual economic preferences and behavior and that this influence should not be underestimated as a result of confounds that may obscure it in many observational contexts.

5 References

- 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.
- Barwick, Panle Jia, Shanjun Li, Jeremy Wallace, and Jessica Chen Weiss. 2019. "Commercial Casualties: Political Boycotts and International Disputes." *Available at SSRN 3417194*.
- Baughn, C Christopher, and Attila Yaprak. 1996. "Economic Nationalism: Conceptual and Empirical Development." *Political Psychology*, 759–78.
- Callahan, William A. 2010. *China: The Pessimist Nation*. Oxford ; New York: Oxford University Press.
- Chubb, Andrew. 2014. *Maritime Awareness: Popular Attitudes Towards China's Maritime Disputes in the East and South China Seas*.
- CRS. 2021. "The Senkakus (Diaoyu/Diaoyutai) Dispute: U.S. Treaty Obligations." *Congressional Research Service*, no. R42761 (March). <https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=&ved=2ahUKEwj7z4LuraP3AhXGm-AKHQ8MC4gQFnoECACQAQ&url=https%3A%2F%2Fsgp.fas.org%2Fcrs%2Frow%2FR42761.pdf&usg=AOvVaw0qUD20yB6ANqK1Fjs7xEGL>.
- Davis, Christina L, and Sophie Meunier. 2011. "Business as Usual? Economic Responses to Political Tensions." *American Journal of Political Science* 55 (3): 628–46.
- Dickson, Eric S. 2011. "Economics Versus Psychology Experiments: Stylization, Incentives, and Deception." In *Cambridge Handbook of Experimental Political Science*, edited by James N. Druckman, Donald P Green, James H. Kuklinski, and Arthur Lupia. Cambridge University Press.
- Fang, Songying, and Xiaojun Li. 2020. "Historical Ownership and Territorial Disputes." *The Journal of Politics* 82 (1): 000–000.
- French, Erik. 2016. "SHOULD INVESTORS FEAR CHINA'S ANTI-AMERICAN PROTESTS?" *Global Risk Insights*. <https://globalriskinsights.com/2016/07/china-anti-american-protests/>.
- 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 Xīnxī." Accessed November 22, 2019. <http://sma.fudan.edu.cn/datainfo/weibo.html>.
- Gries, Peter Hays. 2005. "China's "New Thinking" on Japan." *The China Quarterly*, no. 184 (December): 831–50. <http://www.jstor.org/stable/20192541>.
- Hartman, Erin, and F Daniel Hidalgo. 2018. "An Equivalence Approach to Balance and Placebo Tests." *American Journal of Political Science* 62 (4): 1000–1013.
- Heilmann, Kilian. 2016. "Does Political Conflict Hurt Trade? Evidence from Consumer Boycotts." *Journal of International Economics* 99: 179–91.
- Higgins, Andrew, and Maureen Fan. 2012. "Chinese Communist Leaders Denounce U.S. Values but Send Children to U.S. Colleges." *The Washington Post*, May. https://www.washingtonpost.com/world/asia_pacific/chinese-communist-leaders-denounce-us-

- [values-but-send-children-to-us-colleges/2012/05/18/gIQAiEidZU_story.html](https://www.bepress.com/jeap/vol11/iss1/art1/).
- Hong, Canhui, Wei-Min Hu, James E Prieger, and Dongming Zhu. 2011. “French Automobiles and the Chinese Boycotts of 2008: Politics Really Does Affect Commerce.” *The BE Journal of Economic Analysis & Policy* 11 (1).
- Huang, Haifeng. 2015. “Propaganda as Signaling.” *Comparative Politics* 47 (4): 419–44.
- . 2018. “The Pathology of Hard Propaganda.” *The Journal of Politics* 80 (3): 1034–38.
- Jiang, Junyan, and Dali L Yang. 2016. “Lying or Believing? Measuring Preference Falsification from a Political Purge in China.” *Comparative Political Studies* 49 (5): 600–634.
- Keele, Luke. 2015. “The Statistics of Causal Inference: A View from Political Methodology.” *Political Analysis* 23 (3): 313–35.
- King, Gary. 1997. *A Solution to the Ecological Inference Problem : Reconstructing Individual Behavior from Aggregate Data*. Course Book. Princeton, New Jersey: Princeton University Press.
- 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.
- Klein, Jill Gabrielle, Richard Ettenson, and Marlene D Morris. 1998. “The Animosity Model of Foreign Product Purchase: An Empirical Test in the People’s Republic of China.” *Journal of Marketing* 62 (1): 89–100.
- Klinger, Julie Michelle. 2018. *Rare Earth Frontiers: From Terrestrial Subsoils to Lunar Landscapes*. Cornell University Press.
- Kuran, Timur. 1997. *Private Truths, Public Lies: The Social Consequences of Preference Falsification*. Harvard University Press.
- Lu, Shen, and Ben Westcott. 2016. “South China Sea: Beijing Calls KFC, Apple Protests ‘Irrational.’” CNN. <https://edition.cnn.com/2016/07/20/asia/china-nationalism-south-china-sea-ruling/index.html>.
- Mainichi, The. 2012. “Japan Buys Three of Senkaku Islands from Owner.” *Mainichi Shimbun [Daily Newspaper]*, September. <https://archive.ph/20121209130337/http://mainichi.jp/english/english/newsselect/news/20120911p2g00m0dm074000c.html#selection-811.0-811.33>.
- Masterson, Michael. 2022a. “Catching Fire: How National Humiliation Spreads Hostile Foreign Policy Preferences on Chinese Social Media.” https://michaelfmasterson.com/static/docs/emotional_contagion.pdf.
- . 2022b. “Humiliation and International Conflict Preferences.” *The Journal of Politics* 84 (4). <https://doi.org/10.1086/715591>.
- McCarthy, Simone. 2022. “As War Breaks Out in Europe, China Blames the US.” CNN, February. <https://www.cnn.com/2022/02/25/china/china-reaction-ukraine-russia-intl-hnk-mic/index.html>.
- Mutz, Diana C, and Eunji Kim. 2017. “The Impact of in-Group Favoritism on Trade Preferences.” *International Organization* 71 (4): 827–50.
- Pan, Jennifer, and Alexandra A Siegel. 2020. “How Saudi Crackdowns Fail to Silence Online Dissent.” *American Political Science Review* 114 (1): 109–25.
- Pan, Jennifer, and Yiqing Xu. 2018. “China’s Ideological Spectrum.” *The Journal of*

- Politics* 80 (1): 254–73.
- Pickel, Andreas. 2003. “Explaining, and Explaining with, Economic Nationalism.” *Nations and Nationalism* 9 (1): 105–27.
- 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. 2014. “A Wave to Worry About? Public Opinion, Foreign Policy and China’s Anti-Japan Protests.” *Journal of Contemporary China* 23 (86): 197–215.
- 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.
- Schuman, Howard. 1981. *Questions and Answers in Attitude Surveys : Experiments on Question Form, Wording, and Context*. Quantitative Studies in Social Relations. New York: Academic Press.
- Shen, Xiaoxiao, and Rory Truex. 2021. “In Search of Self-Censorship.” *British Journal of Political Science* 51 (4): 1672–84.
- Shimp, Terence A, and Subhash Sharma. 1987. “Consumer Ethnocentrism: Construction and Validation of the CETSCALE.” *Journal of Marketing Research* 24 (3): 280–89.
- Spence, Jonathan D. 1990. *The Search for Modern China*. 1st ed. New York, NY: W.W. Norton.
- Taylor, Rob. 2018. “Sour Grapes: Why Australian Wine Can’t Get Into China Australia Faces an Awkward Balancing Act Because It Relies Heavily on China for Its Economic Well-Being.” *The Wall Street Journal*, May. <https://www.wsj.com/articles/australian-wine-trade-corked-as-china-retaliates-against-meddling-claims-1527426000>.
- Vogel, Ezra F. 2019. *China and Japan: Facing History*. Cambridge, Massachusetts: The Belknap Press of Harvard University Press.
- 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>.
- 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.
- 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*.