

---

# Bada Bing, Bada Boom

## Microsoft Bing's Chinese Political Censorship of Autosuggestions in North America

By Jeffrey Knockel and Lotus Ruan

**MAY 19, 2022**

**RESEARCH REPORT #156**

---

---

# Copyright

© Citizen Lab



Licensed under the Creative Commons BY-SA 4.0 (Attribution-ShareAlike licence). Electronic version first published in 2022 by the Citizen Lab. This work can be accessed through <https://citizenlab.ca/2022/05/bada-bing-bada-boom-microsoft-bings-chinese-political-censorship-autosuggestions-north-america/>.

Document Version: 1.0

The Creative Commons Attribution-ShareAlike 4.0 license under which this report is licensed lets you freely copy, distribute, remix, transform, and build on it, as long as you:

- give appropriate credit;
- indicate whether you made changes; and
- use and link to the same CC BY-SA 4.0 licence.

However, any rights in excerpts reproduced in this report remain with their respective authors; and any rights in brand and product names and associated logos remain with their respective owners. Uses of these that are protected by copyright or trademark rights require the rightsholder's prior written agreement.

---

## Suggested Citation

Jeffrey Knockel and Lotus Ruan. “Bada Bing, Bada Boom: Microsoft Bing’s Chinese Political Censorship of Autosuggestions in North America,” Citizen Lab Report No. 156, University of Toronto, May 2022.

---

## Acknowledgements

We would like to thank Jedidiah Crandall, Miles Kenyon, and Gabby Lim for helpful comments and peer review. Funding for this research was provided by foundations [listed on the Citizen Lab's website](#). Research for this project was supervised by Masashi Crete-Nishihata and Professor Ron Deibert.

---

## About the Citizen Lab, Munk School of Global Affairs & Public Policy, University of Toronto

**The Citizen Lab** is an interdisciplinary laboratory based at the Munk School of Global Affairs & Public Policy, University of Toronto, focusing on research, development, and high-level strategic policy and legal engagement at the intersection of information and communication technologies, human rights, and global security.

We use a “mixed methods” approach to research that combines methods from political science, law, computer science, and area studies. Our research includes investigating digital espionage against civil society, documenting Internet filtering and other technologies and practices that impact freedom of expression online, analyzing privacy, security, and information controls of popular applications, and examining transparency and accountability mechanisms relevant to the relationship between corporations and state agencies regarding personal data and other surveillance activities.

---

# Contents

<b>Key Findings</b>	1
<b>Introduction</b>	1
<b>Background</b>	3
Search engine autosuggestions	3
Communicating via autosuggestions	3
Microsoft Bing	4
Microsoft in China	4
<b>Methodology</b>	5
<b>Experimental setup</b>	8
<b>Findings</b>	8
Is there Chinese political censorship in North America?	12
<b>Content analysis</b>	18
Cross-region comparison	21
Glitches in the matrix	24
Jeff Widener	24
Hillary Clinton and Alex Jones	25
Other censorship findings	26
Censorship implementation	28
Affected services beyond Bing	28
Windows Start Menu	28
Microsoft Edge	29
DuckDuckGo	30
Autosuggestion censorship impact	30
Hillary Clinton and Alex Jones	31
Shutdown of autosuggestion functionality in China	33
<b>Limitations</b>	36
<b>Discussion</b>	37
Does Bing's development team follow Chinese principles and norms concerning political expression?	38
Is Microsoft censoring outside of China to appease the Chinese government?	38
Is Bing's censorship of users globally an unintentional result of their censorship of Chinese users?	39
Is Chinese users' self-censorship affecting autosuggestions globally?	39
Can Microsoft even permanently address this problem?	40
<b>Availability</b>	41
<b>Appendix: Bug in Keyword Research tool</b>	41
<b>Appendix: Comments on Microsoft's response</b>	42

# Key Findings

Read the report's key findings in [Simplified](#) and [Traditional Chinese](#)

- › We analyzed Microsoft Bing's autosuggestion system for censorship of the names of individuals, finding that, outside of names relating to eroticism, the second largest category of names censored from appearing in autosuggestions were those of Chinese party leaders, dissidents, and other persons considered politically sensitive in China.
- › We consistently found that Bing censors politically sensitive Chinese names over time, that their censorship spans multiple Chinese political topics, consists of at least two languages, English and Chinese, and applies to different world regions, including China, the United States, and Canada.
- › Using statistical techniques, we preclude politically sensitive Chinese names in the United States being censored purely through random chance. Rather, their censorship must be the result of a process disproportionately targeting names which are politically sensitive in China.
- › Bing's Chinese political autosuggestion censorship applies not only to their Web search but also to the search built into Microsoft Windows as well as DuckDuckGo, which uses Bing autosuggestion data.
- › Aside from Bing's Chinese political censorship, many names also suffer from collateral censorship, such as Dick Cheney or others named Dick.

## Introduction

Companies providing Internet services in China are [held accountable](#) for the content published on their products and are expected to invest in technology and human resources to censor content. However, as China's economy expands, more Chinese companies are [growing](#) into markets beyond China, and, likewise, the Chinese market itself has also [become](#) a significant portion of international companies' sources of profit. Companies operating Internet platforms with users inside and outside of China increasingly face the dilemma of appeasing Chinese regulators while providing content without politically motivated censorship for users outside of China. Such companies adopt different approaches to meeting the expectation of international users while following strict regulations in China.

Some companies such as [Facebook](#) and [Twitter](#) do not presently comply with Chinese regulations, and their platforms are blocked by China's national firewall. Other companies operate their platforms in China but fragment their user bases. For instance, Chinese tech giant ByteDance operates Douyin inside of China and TikTok outside of China, [subjecting Douyin users](#) to Chinese laws and regulations, while TikTok is [blocked](#) by the national firewall. Users of one fragment of the platform are not able to interact with users in the other. Finally, companies can combine user bases but only subject some communications to censorship and surveillance. Tencent's WeChat [implements](#) censorship policies only on accounts registered to mainland Chinese phone numbers, and, until 2013, Microsoft's Skype [partnered](#) with Hong Kong-based TOM Group to provide a version of Skype for the Chinese market that included [censorship and surveillance of text messages](#). Platforms with combined user bases often provide users with limited transparency over whether their communications have been subjected to censorship and surveillance due to Chinese regulations.

Previous research has demonstrated a growing number of companies that have either accidentally or intentionally enabled censorship and surveillance capacities designed for China-based services on users outside of China. Our analysis of Apple's filtering of product engravings, for instance, [shows](#) that Apple censors political content in mainland China and that this censorship is also present for users in Hong Kong and Taiwan despite there existing no written legal requirement for Apple to do so. While WeChat only implements censorship on mainland Chinese users, we found that communications made on the platform entirely among non-Chinese accounts were subject to [content surveillance](#) which was used to train and build up WeChat's political censorship system in China. TikTok has reportedly [censored](#) content posted by American users which was critical of the Chinese government. Zoom (an American-owned company based in California) [worked](#) with the Chinese government to terminate the accounts of US-based users and disrupt video calls about the 1989 Tiananmen Square Massacre.

In the remainder of this report, we analyze Microsoft Bing's autosuggestion system for censorship of people's names. We chose to test people's names since individuals can represent highly sensitive or controversial issues and because, unlike more abstract concepts, names can be easily enumerated into lists and tested. We begin by providing background on how search autosuggestions work and their significance. We then set out an experimental methodology for measuring the censorship of people's names in Bing's autosuggestions and explain our experimental setup. We then describe the results of this experiment, which were that the names Bing censors in autosuggestions were primarily related to eroticism or Chinese political sensitivity, including for users in North America. We then discuss the consequences of these findings as well as hypothesize why Bing subjects North American users to Chinese political censorship.

# Background

In this section, we provide background on search autosuggestions and their significance as well as discuss Microsoft Bing and Microsoft's history in China.

## Search engine autosuggestions

Search engines play an important role in distributing content and shaping how the public perceives certain issues. Previous studies have analyzed algorithmic biases and subtle censorship implemented by [Baidu in China](#) and [Yandex in Russia](#), each favoring pro-regime and pro-establishment results via source bias and reference bias.



Figure 1: Left, Baidu censoring mention of “Xi Jinping” in autosuggestions for “xi” followed by a space; right, Baidu censoring all autosuggestions for “xi jin”.

In addition to displaying search engine results, search engines also implement *autosuggestion* (sometimes called *autofill* or *autocomplete*) functionality. Autosuggestions are used to fix user typos and also guide and suggest search queries, and autosuggestions often contain answers to a user's question in themselves. Accordingly, autosuggestions play an important role in informing the user. For example, recent [reports](#) on COVID-19 misinformation found that online autosuggestion results influence how users are subject to medical misinformation. These studies collectively demonstrate that search engines can potentially be “[architecturally altered](#)” to serve a particular political, social, or commercial purpose by controlling not only what users are able to see in search results but also the search phrases users might enter in the first place.

## Communicating via autosuggestions

While autosuggestion systems can be thought of as a means for users to quickly obtain information, the communication in these systems is usually not one-way. Microsoft researchers have [previously noted](#) that “[a]utosuggestion systems are typically designed to predict the most likely intended queries given the user's partially typed query, where the

predictions are primarily based on frequent queries mined from the search engine’s query logs” and that, “[s]ince the suggestions are derived from search logs, they can, as a result, be directly influenced by the search activities of the search engine’s users.” Resultantly, Google has [repeatedly struggled](#) to keep its autosuggestions free of hate speech.

Autosuggestion features are known to be under censorship in China. For example, Baidu is known to [filter autosuggestions](#) relating to sensitive topics (see Figure 1). This practice is consistent with the general information control regime in China, which requires all Internet communications to be subject to political censorship. In our previous work, we have found a wide range of user content subject to censorship, including [messages](#), [group chat](#), [file contents](#), [usernames](#), [mood indications](#), and [user profile descriptions](#). As autosuggestions are based on users’ historical searches, they are the result of users’ input and thus are required to be moderated and censored for prohibited content in China.

## Microsoft Bing

In 1998, Microsoft launched Bing’s predecessor, MSN Search. After transitioning through multiple name changes, Microsoft rebranded the search engine as [Bing](#) in 2009 and finally as [Microsoft Bing](#) in 2020.

While Bing’s market share varies regionally, as of March 2022, Bing is used by [6.6 percent](#) of Web users in the United States, [5.4 percent](#) of users in Canada, and [6.7 percent](#) of users in China. While Google is the most popular search engine in North America, Baidu is Bing’s [primary competitor](#) in China.

In addition to Web usage, Bing also sees usage through its integration into multiple Microsoft products and through other search engines which use its data. Since Windows 8.1, Bing has been built into the [Windows start menu](#), providing autosuggestions and search results for queries searched using the Windows start menu search functionality. Bing is also the default search engine in [Microsoft Edge](#), Microsoft’s cross-platform, Chromium-based Web browser, providing both autosuggestions and search results for queries typed into the browser’s search bar. Finally, Bing provides autosuggestion and search result data for other search engines, including [DuckDuckGo](#) and [Yahoo](#).

## Microsoft in China

Entering the Chinese market in 1992, Microsoft [established](#) an early presence in China long preceding the Chinese debut of its Internet search engine. From computers’ operating systems to gaming consoles and communications platforms, the American company has invested in multiple technology sectors and is largely successful in China despite the country’s restricted regulatory environment. It is unclear exactly how much

the China market accounts for Microsoft’s [global revenues](#), as the company has [kept](#) it a secret for years, leaving the estimated figure to be as [low](#) as around one percent to as [high](#) as 10 percent. It is clear, however, that Microsoft continues to [expand](#) in China and its relations with Chinese regulators appear relatively stronger than many other American tech companies despite some [rough interactions](#).

One of the reasons for Microsoft’s continued success in China might be its implementation of censorship in response to China’s content regulations. Before Microsoft announced that it will pull LinkedIn from the Chinese market in October 2022 [citing](#) a “challenging operating environment,” LinkedIn was found to [censor](#) posts or personal profiles considered sensitive to the Chinese government. Similarly, Microsoft has [censored results](#) on Bing in China since 2009. However, the company is under growing scrutiny concerning whether it will expand censorship of Chinese political sensitivity beyond China to advance its commercial interests. In 2021, Bing was [found](#) to censor image results for the query “tank man” in the United States and elsewhere around the 1989 Tiananmen Square Movement anniversary. Microsoft said the blocking was “due to an accidental human error,” dismissing concerns about possible censorship beyond China. In December of 2021, the Chinese government [suspended](#) Bing’s search autosuggestions for users in China for 30 days.

## Methodology

Our analysis found that, other than the search query that the user has typed so far, at least three variables affect the autosuggestions provided by Bing: the user’s region setting, language setting, and geolocation as determined by the user’s IP address. For purposes of measuring censorship of popularly searched names of individuals, we found it primarily relevant whether a user’s IP address is inside or outside of mainland China. In the remainder of this report, we will call a combination of (1) a region, (2) a language, and (3) whether one’s geolocation is inside mainland China a *locale*. Outside of geolocation, the other aspects of a locale can be easily set using Bing’s Web UI or by manually setting URL parameters. For instance, to switch Bing to the “en-US” (United States) region and “fr” (French) language, users can visit the following URL:

```
https://www.bing.com/?mkt=en-US&setlang=fr
```

We found that these settings affect the entire browsing session, and so setting them will affect other browser tabs and windows, unless those tabs or windows were specifically created in a separate browsing session. We set these URL parameters to automatically test different Bing regions and languages.

To test a variety of locales, we chose a subset of the regions documented by Microsoft in Bing's [API documentation](#), namely: "en-US" (United States), "en-CA" (Canada), and "zh-CN" (mainland China). For each region tested, we tested two different languages, English ("en") and simplified Chinese ("zh-hans"). Geolocation as determined by IP address can affect autosuggestions provided by Bing, such as Bing suggesting local restaurants when searching for dining options. However, while we found that whether one's IP address was inside or outside of mainland China dramatically affected the level of censorship Bing applied to autosuggestions, we are not otherwise aware of IP address affecting Bing's autosuggestion filtering. Accordingly, we test from two different networks, a network in North America and a network in mainland China.

Since Bing only allows the mainland China region to be selected when accessing the site from a mainland China IP address, we are not able to test other regions from a mainland China IP address. Thus, instead of testing different regions from this address, we test a feature specific to the mainland China region, whether the search engine is configured to use the 国内版 (domestic version) versus the 国际版 (international version). While we were unable to find documentation clearly elucidating the differences between these versions, we generally found that the domestic version was more likely to interpret English letters as Chinese pinyin whereas the international version generally interpreted English letters as English words. For nomenclature purposes, we consider the "domestic version" and the "international version" to be two different languages of the mainland China region, even though they are not languages per se.

To test for censorship in each locale, we used sample testing. We generated queries to test in each locale using the following method. From English Wikipedia, we extracted the titles of any article meeting all of the following criteria:

- After stripping parentheticals from its title, the article's title consisted entirely of English alphabetic characters or spaces.
- The article received at least 1,000 views during September 2021.
- The article contained either a "person" or an "officeholder" infobox (English Wikipedia articles generally use a special infobox for political officeholders).

The resulting list of article titles we henceforth refer to as *English letter names*.

From Chinese Wikipedia, we also extracted the titles of any article meeting all of the following criteria

- After stripping interpunct (".") symbols and parentheticals from its title, the article's title consisted entirely of simplified Chinese characters or spaces. (In Chinese, interpuncts are often used to mark the separation of first, last, and other names in names transliterated from other languages.)

- The article received at least 1,000 views during September 2021.
- The article contained either a “person” or an “藝人” (artist) infobox (Chinese Wikipedia articles generally use a special infobox for artists such as actors or singers).

The resulting list of article titles we henceforth refer to as *Chinese character names*.

For each of these names, we generated three queries from the name as follows:

- the name minus the last letter (e.g., “Xi\_Jinpin”\*)
- the name itself (“Xi\_Jinping”)
- the name followed by a space (“Xi\_Jinping\_”)

If, for a given name, none of its queries’ autosuggestions contain the original name (“Xi\_Jinping”), including if there were no autosuggestions at all, then we say that the name was *suggestionless*.

\* In this report, we use the “open box” symbol “\_” to unambiguously render spaces when describing Bing test queries. We do this because often spaces appear at the end of our test queries, which might be difficult to display using the traditional space character “ ”. For instance, instead of rendering “xi” followed by a space as “xi ”, we render it as “xi\_” so that it is obvious that a single space character trails “xi” in this query.

Just because a name is suggestionless does not necessarily mean that any censorship is occurring, as it might be a name that is uncommonly searched for using Bing or otherwise too obscure. While we previously restricted our names to only those whose corresponding Wikipedia articles had over 1,000 views in a month, Wikipedia article views are not necessarily a predictor of how often a term is searched on Bing. Thus, to help ensure that a name is sufficiently popular on Bing to justify the conclusion that it is censored, we utilized Bing’s [Keyword Research API](#), which [provides](#) search “query volume data” in units called *impressions*. Bing [describes](#) these “impressions” as “based on organic query data from Bing and is raw data, not rounded in any way.”

For each locale, for all suggestionless names, we used the Keyword Research API to determine how many times that name had been searched in that locale’s region. If the name reported at least 35 impressions in the last six months, we concluded that the suggestionless name had been *censored* in that region. We chose the number 35 qualitatively, as suggestionless names with at least 35 impressions tended to fall into predictable categories such as being related to eroticism, misinformation, or Chinese political sensitivity, whereas names with fewer than 35 impressions more often had no identifiable motivation for being censored and were thus more likely to be suggestionless due to their search unpopularity.

# Experimental setup

In our methodology, we describe testing from two different networks: a North American network and a mainland Chinese network. For the North American network, we tested from Toronto, Canada, and, for the Chinese network, we tested from Shaoxing, China. The Toronto testing occurred from a machine hosted on a DigitalOcean network, and the Shaoxing testing was performed using a VPN server provided by a popular VPN service whose Chinese vantage points we had confirmed to be in China and subject to censorship from China's national firewall. We performed the experiment described above during the week of December 10–17, 2021.

# Findings

Performing this experiment, we collected 7,186 Chinese character names from Chinese Wikipedia and 97,698 English letter names from English Wikipedia.

Geolocation	Mainland China		North America					
	Mainland China		Mainland China		USA		Canada	
	Domestic	International	Domestic	International	English	Simplified Chinese	English	Simplified Chinese
<b>Suggestionless</b>	441	444	394	374	374	106	262	280
<b>Censored</b>	146	142	73	59	32	33	39	47

Table 1: Among 7,186 Chinese character names, for each locale, how many were suggestionless and how many of the suggestionless were censored (had at least 35 search volume impressions).

Geolocation	Mainland China		North America				
	Mainland China		Mainland China		USA		Canada
	International	International	English	Simplified Chinese	English	Simplified Chinese	
<b>Suggestionless</b>	2,183	2,183	2,183	1,699	1,380	2,333	
<b>Censored</b>	162	159	653	636	456	579	

Table 2: Among 97,698 English letter names, for each locale, how many were suggestionless and how many of the suggestionless were censored (had at least 35 search volume impressions).

Each locale tested had suggestionless names as well as names that we found to be censored (see Tables 1 and 2). Most of the suggestionless names had little to no search volume and thus were most likely suggestionless due to being insufficiently popular search queries. However, as we explain in our methodology, for each locale, we consider a

suggestionless name censored only if its locale had at least 35 Bing search volume impressions in the six months prior to our experiment. Resultantly, we found between 32 and 146 Chinese character names censored in each locale tested and between 162 and 653 English letter names censored in each locale tested.

Since there may exist names that are both targeted by Bing for censorship and that are low search volume, one consequence of our requirement that a name have at least 35 impressions in a region before being considered censored is that we would expect to discover more censored words in regions with more Bing users and thus more search volume. Thus, with our data, we cannot use the absolute number of censored names across different regions as the means to say that one region is more censored than another. To state this another way: Bing may be targeting names for censorship in regions whose censorship we were unable to detect due to those names having insufficient search volume in those regions. Thus, if we detect a name censored in one region but not another, we may have not detected the name censored in the other region merely due to it having insufficient search volume in that region.

Geolocation	Mainland China		North America					
	Mainland China		Mainland China		USA		Canada	
Region	Domestic	International	Domestic	International	English	Simplified Chinese	English	Simplified Chinese
Language	Domestic	International	Domestic	International	English	Simplified Chinese	English	Simplified Chinese
<b>Chinese political</b>	119 (81.5%)	110 (77.5%)	63 (86.3%)	49 (83.1%)	30 (93.8%)	5 (15.2%)	30 (76.9%)	39 (83.0%)
<b>Entertainment</b>	16 (11.0%)	17 (12.0%)	6 (8.2%)	6 (10.2%)	1 (3.1%)	6 (18.2%)	5 (12.8%)	5 (10.6%)
<b>Eroticism</b>	8 (5.5%)	10 (7.0%)	3 (4.1%)	3 (5.1%)	0 (0.0%)	18 (54.5%)	3 (7.7%)	3 (6.4%)
<b>Historical figure</b>	0 (0.0%)	0 (0.0%)	0 (0.0%)	0 (0.0%)	0 (0.0%)	0 (0.0%)	1 (2.6%)	0 (0.0%)
<b>International politician</b>	1 (0.7%)	1 (0.7%)	1 (1.4%)	1 (1.7%)	1 (3.1%)	2 (6.1%)	0 (0.0%)	0 (0.0%)
<b>Public figure</b>	2 (1.4%)	4 (2.8%)	0 (0.0%)	0 (0.0%)	0 (0.0%)	2 (6.1%)	0 (0.0%)	0 (0.0%)
<b>Total</b>	146 (100.0%)	142 (100.0%)	73 (100.0%)	59 (100.0%)	32 (100.0%)	33 (100.0%)	39 (100.0%)	47 (100.0%)

Table 3: For each locale, the number of Chinese character names which have at least 35 impressions in that region and that have no autosuggestions, according to each name's content category.

Geolocation		Mainland China	North America				
Region	Mainland China	Mainland China	USA	Canada			
Language	International	International	English	Simplified Chinese	English	Simplified Chinese	
<b>Chinese political</b>	8 (4.9%)	6 (3.8%)	11 (1.7%)	4 (0.6%)	1 (0.2%)	3 (0.5%)	
<b>Collateral</b>	35 (21.6%)	32 (20.1%)	160 (25.0%)	170 (27.3%)	89 (19.7%)	173 (30.1%)	
<b>Entertainment</b>	1 (0.6%)	1 (0.6%)	15 (2.3%)	35 (5.6%)	3 (0.7%)	15 (2.6%)	
<b>Eroticism</b>	112 (69.1%)	114 (71.7%)	433 (67.6%)	395 (63.4%)	351 (77.7%)	371 (64.5%)	
<b>Historical figure</b>	0 (0.0%)	0 (0.0%)	1 (0.2%)	1 (0.2%)	0 (0.0%)	2 (0.3%)	
<b>International politician</b>	0 (0.0%)	0 (0.0%)	1 (0.2%)	7 (1.1%)	0 (0.0%)	1 (0.2%)	
<b>Overshadowed</b>	6 (3.7%)	6 (3.8%)	10 (1.6%)	2 (0.3%)	2 (0.4%)	0 (0.0%)	
<b>Public figure</b>	0 (0.0%)	0 (0.0%)	10 (1.6%)	9 (1.4%)	6 (1.3%)	10 (1.7%)	
<b>Total</b>	162 (100.0%)	159 (100.0%)	641 (100.0%)	623 (100.0%)	452 (100.0%)	575 (100.0%)	

Table 4: For each locale, the number of English letter names which have at least 35 impressions in that region and that have no autosuggestions, according to each name’s content category.

To better understand the names we discovered to be censored, we categorized them based on their underlying context (see Tables 3 and 4 for more details). We began by reviewing all of the names and abstract common themes among them. We then reviewed the names again and categorized them into the common themes that we discovered, which include “Chinese political” (e.g., incumbent and retired Chinese Communist Party leaders, dissidents, political activists, and religious figures), “historical figure” (e.g., ancient philosophers and pre-PRC thinkers), “international politician”, “entertainment” (e.g., singers, celebrities), and “eroticism”. We categorized as “eroticism” anyone who meets or has met any of the following criteria: anyone participating in pornography or its production, glamor models, gravure models, burlesque dancers, drag queens, and anyone who has been a famous victim of a nude photo or sex video leak. As with the other categories, we based our criteria for this category not on our own intuitions but rather on what we found Bing to censor.

Finally, we created two special categories, “collateral” and “overshadowed”. We assigned the “collateral” category to names that do not appear to be directly targeted for censorship but rather appear to be collateral censorship from some other censorship rule. We found that the most common reason for a name being collaterally censored was containing the name “Dick”, e.g., “Dick Cheney”. The “overshadowed” category is similar in that we assign names to this category which do not appear targeted for censorship. However, instead of being collaterally censored, we believe that overshadowed names do not have autosuggestions because they are overshadowed by the autosuggestions of someone with a similar name. For example, we found that suggestions for actor “Gordon Ramsey” were overshadowed by suggestions for the more famous “Gordon Ramsay”, celebrity chef. However, suggestions for the less famous “Gordon Ramsey” could still be found if we used a more specific query, such as “Gordon Ramsey actor”.

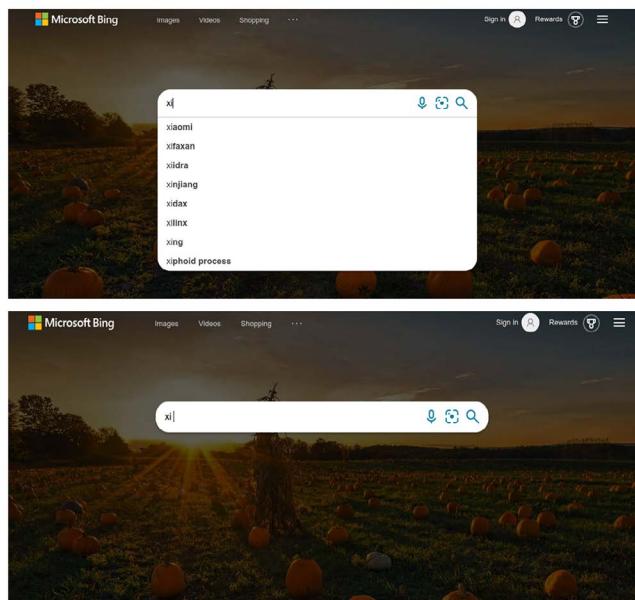


Figure 2: Top, October 2021 autosuggestions for “xi” in the United States English locale which contain no mention of Xi Jinping; bottom, a complete absence of autosuggestions for “xi\_”. Our December 2021 measurement findings were analogous.

In nearly every locale we tested, censored Chinese character names were most likely to belong to the “Chinese political” category, whereas censored English letter names were more likely to belong to the “eroticism” category, although each locale also censored “Chinese political” English letter names such as “Xi Jinping”, “Liu Xiaobo”, and “Tank Man” (see Figure 2 for an example). Overall, when considering all censored Chinese character and English letter names in aggregate, the largest number of names were in the “eroticism” category followed by the “Chinese political” category, excluding the two special “collateral” and “overshadowed” categories.

Already this might seem like compelling evidence that Bing performs Chinese political censorship both inside and outside of China. However, how can we be certain? After

all, while many sensitive Chinese political names are censored on Bing, many are not. Moreover, in locales like United States English, there were 11 “Chinese political” English letter names censored – is that even significant? Perhaps these names are censored simply due to some defect in Bing that fails to show suggestions for names uniformly at random. In the following section, we use statistical techniques to answer whether Bing is performing any targeted censorship at all, whether any of the targeting is for names of Chinese political sensitivity, and whether such censorship extends outside of China.

## Is there Chinese political censorship in North America?

Up to this point, we have found the names of people who are popularly read about on Wikipedia and searched about on Bing but yet, for some reason, have no Bing autosuggestions. These names appear to often fall into certain categories, such as being associated with eroticism, being politically sensitive in China, or containing certain English swear words in them such as “dick”. Still, how do we really know that these names are being specifically targeted for censorship and that they are not just random failures of Bing’s autosuggestion system? For instance, it could be that, because of the way we selected popular names, we found a lot of names related to eroticism and Chinese politics because those are the names that are most popularly searched for on Bing or read about on Wikipedia.

In this section, we further explore the nature of Bing’s censorship using statistical techniques. Since it is already well known that Bing implements censorship in mainland China to comply with legal requirements, we focus in this section on whether Bing implements Chinese political censorship in North America. Particularly, we look at the United States English locale, as we presume that that locale is the most common locale of users in North America.

## Is there any targeted censorship of Chinese character names in the United States?

In our methodology, recall that we select names to test by only looking at sufficiently popular Wikipedia articles about people. We then further filter our results by only considering suggestionless names with a sufficiently high Bing search volume. Thus, in two ways we select for names that are popular. It is therefore possible that, even if the names for which Bing failed to show autosuggestions were not really censored but rather chosen by some random process, we may still see themes such as Chinese political sensitivity and eroticism occur with high frequency in our censored data set because these may simply be popular topics commonly viewed on Wikipedia and searched on Bing.

To test the hypothesis that we may only be seeing these results by chance, we use statistical significance testing. If Bing does not target any types of content for censorship and

the names that we call censored are resulting from a uniformly random process, then we would expect the proportions of categories in both names that are censored and names that are non-censored to be the same. With the censored words already categorized, we chose 120 non-censored names at random using the same thresholds concerning Wikipedia article views and Bing search volume that we applied to the censored names.

	Non-censored	Censored
<b>Chinese political</b>	11 (9.2%)	30 (93.8%)
<b>Entertainment</b>	69 (57.5%)	1 (3.1%)
<b>Eroticism</b>	5 (4.2%)	0 (0.0%)
<b>Historical figure</b>	12 (10.0%)	0 (0.0%)
<b>International politician</b>	3 (2.5%)	1 (3.1%)
<b>Public figure</b>	20 (16.7%)	0 (0.0%)
<b>Total</b>	120 (100.0%)	32 (100.0%)

Table 5: Among Chinese character names, a contingency table for whether a name is censored versus its category.

Just by eyeballing the results of this comparison (see Table 5), we can see that the proportions of categories between censored and non-censored names are radically different. For instance, among the censored names, 93.8 percent are Chinese political, whereas among the non-censored, only 9.2 percent are in that category. Since we chose the non-censored names at random, if Bing were also choosing names to censor at random, we would expect these proportions to be the same.

Nevertheless, using statistical significance testing, we need not rely merely on our intuitions. Using [Fisher's test](#), we can test this hypothesis, called in statistical significance testing the *null hypothesis*. Specifically, our null hypothesis under question is that the proportional category sizes in both groups are the same. The result of Fisher's test is something called a *p value*, which, in this case, is the probability that we would see at least as extreme of differences in their proportions by chance. Applying Fisher's test to Table 5, we find that  $p = 5.35 \cdot 10^{-19}$ , confirming our intuitions that these differences are not random chance and that Bing is targeting specific categories of content in the United States for censorship. With a *p* value this small, other considerations beyond observing as extreme of results by chance become more likely such as that we mistyped the number into this document.

## Is there any targeted Chinese political censorship of Chinese character names in the United States?

Above we established that Bing is censoring specific categories of Chinese character names in the United States. However, are they specifically targeting Chinese politically sensitive names for censorship?

	Non-censored	Censored
<b>Chinese political</b>	11 (9.2%)	30 (93.8%)
<b>Not Chinese political</b>	109 (90.8%)	2 (6.2%)
<b>Total</b>	120 (100.0%)	32 (100.0%)

Table 6: Among Chinese character names, a contingency table for whether a name is censored versus whether it is Chinese politically sensitive.

To answer this question, we are concerned with a smaller, 2x2 table containing only two categories, “Chinese political” and “not Chinese political”, a category that is the aggregate of all other categories (see Table 6). While again it seems obvious looking at the data that censored names are disproportionately Chinese political compared to names chosen at random, we nevertheless perform Fisher’s test. In this case, we formulate our null hypothesis to be that the proportion of censored Chinese political names is no greater than the proportion of Chinese political names that are not censored. We find that  $p = 2.61 \cdot 10^{-20}$ , all but confirming that Bing is targeting Chinese politically sensitive names in the United States for censorship.

## Is there any targeted censorship of English letter names in the United States?

In the previous section, we used statistical techniques to test whether Chinese character names such as “张高丽” (Zhang Gaoli) were disproportionately targeted for censorship in the United States or whether their lack of autosuggestions might somehow be the result of random failure. In this section, we will apply the same techniques to test this question with respect to censored English letter names such as “Xi Jinping”.

While previously we categorized censored English letter names which were censored for containing words like “Dick” as “collateral” censorship and names that were overshadowed by much more famous people as being “overshadowed”, we have no way of making such categorizations in the non-censored group. Thus, for purposes of our statistical testing, we recategorize such names according to how we would categorize them a priori, i.e., as we would if we did not know whether they were censored or not into ordinary categories such as “entertainment” and so on.

	Non-censored	Censored
<b>Chinese political</b>	1 (0.8%)	11 (1.7%)
<b>Entertainment</b>	73 (60.8%)	110 (17.2%)
<b>Eroticism</b>	0 (0.0%)	436 (68.0%)
<b>Historical figure</b>	5 (4.2%)	12 (1.9%)
<b>International politician</b>	10 (8.3%)	12 (1.9%)
<b>Public figure</b>	31 (25.8%)	60 (9.4%)
<b>Total</b>	120 (100.0%)	641 (100.0%)

Table 7: Among English letter names, a contingency table for whether a name is censored versus its category. Since names categorized as “collateral” censorship and “overshadowed” are recategorized *a priori*, the proportion of censored names in categories such as “entertainment” and “public figure” is much higher than in Table 5.

Just by eyeballing the proportions between censored and non-censored names (see Table 7), we can again see that the proportions of categories between censored and non-censored names are radically different. For instance, among the censored names, 68.0% are related to eroticism, whereas among the non-censored, zero are. Using [Fisher’s test](#), we again test whether the proportions between censored and non-censored are the same, finding that  $p = 4.64 \cdot 10^{-51}$  probability we would see such extreme differences in proportions by chance if the proportions were the same.

	Non-censored	Censored
<b>Eroticism</b>	0 (0.0%)	436 (68.0%)
<b>Not eroticism</b>	120 (100.0%)	205 (32.0%)
<b>Total</b>	120 (100.0%)	641 (100.0%)

Table 8: Among English letter names, a contingency table for whether a name is censored versus whether it is eroticism.

Due to the large differences in proportions in the “eroticism” category, in Table 8, we divide the names into “eroticism” and “not eroticism” categories using the same method as we did to construct Table 6. While again it seems obvious looking at the data that censored names are disproportionately associated with eroticism compared to names chosen at random, we nevertheless perform Fisher’s test. Our null hypothesis is that the

proportion of censored Eroticism names is no greater than the proportion of Eroticism names that are not censored. We find that we would expect to see only a  $p = 9.50 \cdot 10^{-52}$  probability of seeing as great of a number of Eroticism names that are censored by chance, all but confirming that Bing is targeting eroticism-related English letter names for censorship in the United States.

## Is there any targeted Chinese political censorship of English letter names in the United States?

Although thus far we have used statistical hypothesis testing in cases where the data already lent to an obvious conclusion, we will now explore a question in which our intuitions may be less able to come to a conclusion by eyeballing the data. Namely, with the assistance of statistical hypothesis testing, we will investigate whether there is Chinese political censorship of English letter names in the United States. Looking at Table 7, it is not obvious.

	Non-censored	Censored
<b>Chinese political</b>	1 (0.8%)	11 (1.7%)
<b>Not Chinese political</b>	119 (99.2%)	630 (98.3%)
<b>Total</b>	120 (100.0%)	641 (100.0%)

Table 9: Among English letter names, a contingency table for whether a name is censored versus whether it is Chinese politically sensitive.

Table 9 is the resulting 2x2 contingency table comparing whether a name is censored versus whether it is Chinese politically sensitive. We formulate our null hypothesis to be that the proportion of censored Chinese political names is no greater than the proportion of Chinese political names that are not censored. Applying Fisher's test, we find that  $p = 0.412$ , which is an inconclusive result. Looking at the table, this result is not too surprising, as, even though the proportion of censored Chinese political names (1.7 percent) is over twice the proportion of non-censored Chinese political names (0.8 percent), the absolute number of Chinese political names in both the non-censored and censored columns is small, and thus the test is lacking in statistical power. In the following section we perform a different test looking only at Chinese pinyin names, which achieves a more conclusive result concerning whether Bing performs censorship of English letter names in the United States.

## Is there any targeted Chinese political censorship of pinyin names in the United States?

In our tests of English letter names, there exists one possible confounding variable that we have yet to consider. One might argue that, even if Bing's censorship of "Chinese

political” names was the result of a random process, such a process might be disproportionately targeting foreign names such as names written in Chinese pinyin, especially if Bing’s autosuggestion algorithms were more poorly suited to such names. While not all English letter names categorized as “Chinese political” were pinyin (specifically, Tank Man, Gedhun Choekyi Nyima, and Rebiya Kadeer), the remaining eight were (e.g., Xi Jinping, Li Wenliang, etc.). Thus, it may only appear that there is Chinese political censorship of English letter names if Bing is simply bad at providing autosuggestions for pinyin names but not necessarily censoring names politically sensitive in China.

	Non-censored	Censored
<b>Chinese Political</b>	18 (15.0%)	8 (100.0%)
<b>Entertainment</b>	64 (53.3%)	0 (0.0%)
<b>Historical Figure</b>	12 (10.0%)	0 (0.0%)
<b>Public Figure</b>	26 (21.7%)	0 (0.0%)
<b>Total</b>	120 (100.0%)	8 (100.0%)

Table 10: Among pinyin names, a contingency table for whether a name is censored versus its category.

To eliminate such a hypothesis, we compare the eight pinyin names to 120 non-censored pinyin names, finding that all eight (100%) censored are “Chinese political” versus only 15% of the non-censored being Chinese political (see Table 10). Although there is a large difference in proportions in the Chinese political category, with the proportion of censored being 6.67 times the non-censored, we are estimating the censored Chinese political proportion from only eight censored pinyin names. Thus, despite the large difference in proportions, it may not be intuitively obvious whether there is a statistically significant difference between the two. However, Fisher’s test already includes the absolute numbers of samples as part of its calculus, so we can nevertheless be confident in its results without making any additional consideration for the sample sizes.

	Non-censored	Censored
<b>Chinese Political</b>	18 (15.0%)	8 (100.0%)
<b>Not Chinese Political</b>	102 (85.0%)	0 (0.0%)
<b>Total</b>	120 (100.0%)	8 (100.0%)

Table 11: Among Pinyin names, a contingency table for whether a name is censored versus whether it is Chinese politically sensitive.

Applying Fisher's test to Table 11 under the null hypothesis that the proportion of censored Chinese political names is no greater than the proportion of Chinese political names that are not censored, we find that  $p = 1.09 \cdot 10^{-6}$ . Even though in choosing to look only at pinyin names we were initially motivated by eliminating a confounding variable, because the experiment was more powerful we were also able to achieve a conclusive result, almost certainly showing that Bing targets English letter names for Chinese political censorship in the United States just as it similarly targets simplified Chinese character names there.

## Content analysis

Across the three regions we tested (i.e., mainland China, the United States, and Canada), we observed overwhelming censorship of Chinese character names relating to Chinese politics. These names predominantly pertain to names of top-level Chinese government leaders and party figures, including incumbent leaders (e.g., 习近平, “Xi Jinping”), retired officials (e.g., 温家宝, “Wen Jiabao”, a former Chinese Premier), historical figures (e.g., 李大钊, “Li Dazhao,” a co-founder of the Chinese Communist Party), and party leaders involved in political scandals or power struggle (e.g., 周永康, “Zhou Yongkang,” a former Party leader).

United States English locale				Canada English locale		
#	Name	Translation	Category	Name	Translation	Category
1	张高丽	Zhang Gaoli	Chinese political	张高丽	Zhang Gaoli	Chinese political
2	江泽民	Jiang Zemin	Chinese political	习近平	Xi Jinping	Chinese political
3	王岐山	Wang Qishan	Chinese political	江泽民	Jiang Zemin	Chinese political
4	胡锦涛	Hu Jintao	Chinese political	傅政华	Fu Zhenghua	Chinese political
5	周永康	Zhou Yongkang	Chinese political	陈独秀	Cheng Duxiu	Chinese political
6	曾庆红	Zeng Qinghong	Chinese political	王岐山	Wang Qishan	Chinese political
7	汪洋	Wang Yang	Chinese political	薄熙来	Bo Xilai	Chinese political
8	赵紫阳	Zhao Ziyang	Chinese political	李大钊	Li Dazhao	Chinese political
9	胡春华	Hu Chunhua	Chinese political	桃乃木香奈	Kana Momonogi	Eroticism
10	王沪宁	Wang Huning	Chinese political	林彪	Lin Biao	Chinese political

Table 12: Ordered by decreasing search volume, the top 10 Chinese-character names that are censored in the United States English and Canada English locales.

The censorship of Chinese leaders' names in the domestic and international versions of Bing in China may be due to Microsoft's compliance with Chinese laws and regulations. However, there is no legal reason for the names to be censored in Bing autosuggestions in the United States and Canada. In Table 12, we highlight the top 10 highest search volume names censored in each of these two North American regions. Although the two regions do not appear to censor exactly the same names, most of the names on both lists are references to Chinese political figures. In the United States English locale, all the top 10 names reference incumbent or recently retired Chinese politicians. In the Canada English locale, three of the top 10 names are referencing people in relation to the history of the Chinese Communist Party such as its founding members; interestingly, “桃乃木香奈” (Kana Momonogi), the name of a Japanese pornographer, also appears on the list.

During our data collection period, former Chinese Vice Premier Zhang Gaoli received the highest search volume on Bing among the names we found censored in the United States and Canada English locales. The high international search volume of Zhang, a retired Chinese politician, is likely due to a [scandal](#) in which he is alleged to have sexually assaulted Chinese tennis star Peng Shuai. Peng first published her allegations on Weibo on November 2, 2021, which were quickly [censored](#) on all Chinese platforms. She then disappeared from the public for almost three weeks, prompting [worldwide media](#) attention on Zhang as well as [an international campaign](#) calling for information with regards to Peng's whereabouts.

United States English locale			Canada English locale	
#	Name	Category	Name	Category
1	Riley Reid	Eroticism	Mia Khalifa	Eroticism
2	Brandi Love	Eroticism	Brandi Love	Eroticism
3	Mia Khalifa	Eroticism	XXXTentacion	Collateral (“XXX”)
4	XXXTentacion	Collateral (“XXX”)	Adriana Chechik	Eroticism
5	Mia Malkova	Eroticism	Nina Hartley	Eroticism
6	Dick Van Dyke	Collateral (“Dick”)	Kendra Lust	Eroticism
7	Xi Jinping	Chinese political	Jenna Jameson	Eroticism
8	Nina Hartley	Eroticism	Dick Van Dyke	Collateral (“Dick”)
9	Adriana Chechik	Eroticism	Julia Ann	Eroticism
10	Jenna Jameson	Eroticism	Asa Akira	Eroticism

Table 13: Ordered by decreasing search volume, the top 10 English letter names that are censored in the United States English and Canada English locales.

Regarding English letter names, we found that most censored names were related to eroticism in each locale that we tested (see Table 13). Even though Bing's censorship of

eroticism may be out of concern over explicit content, such censorship nevertheless has impacts beyond sexual content. For instance, we found that Bing censors the name of “Ilona Staller” in all regions we tested. [Ilona Staller](#) appears to be a former Hungarian-Italian porn star who has turned to politics and ran for offices in Italy since 1979. Censoring autosuggestions containing this name may affect the person’s political career as well.

Notably, not only were autosuggestions containing politically sensitive Chinese character names censored, names of Chinese leaders, dissidents, political activists, and religious figures in English letters were also censored in the United States English and Canada English locales. Whereas Chinese political censorship of Chinese character names in the same locales pertain predominately to names of incumbent, retired, and historical Chinese Communist Party leaders, Chinese political censorship of English letter names appears to have a greater variety, some of which relate closely to current events.

United States English locale		
#	Name	Note
1	Xi Jinping	Incumbent Chinese president
2	Tank Man	Nickname of an unidentified Chinese man who stood in front of a column of tanks leaving Tiananmen Square in Beijing on June 5, 1989
3	Li Wenliang	Chinese ophthalmologist who warned his colleagues about early COVID-19 infections in Wuhan
4	Jiang Zemin	Former Chinese president
5	Guo Wengui	Exiled Chinese billionaire businessman
6	Liu Xiaobo	Deceased Chinese human rights activist and Nobel Peace Prize awardee
7	Gedhun Choekyi Nyima	The 11th Panchen Lama belonging to the Gelugpa school of Tibetan Buddhism, as recognized and announced by the 14th Dalai Lama on 14 May 1995
8	Li Hongzhi	Founder and leader of Falun Gong
9	Li Yuanchao	Former Chinese vice president
10	Rebiya Kadeer	Uyghur businesswoman and political activist
11	Chai Ling	One of the student leaders in the Tiananmen Square protests of 1989

*Table 14: Ordered by decreasing search volume, each of “Chinese political” category English letter names censored in the United States English locale.*

Table 14 shows the 11 “Chinese political” category names censored in the United States English locale by descending search volume. One of the highest volume censored Chinese political names was that of the late Chinese doctor Li Wenliang. Dr. Li [warned](#) his colleagues about early COVID-19 infections in Wuhan but was later forced by local police and medical officials to sign a statement denouncing his warning as unfounded and illegal rumors. Dr. Li died from COVID-19 in early February 2020. References to Dr. Li have been regularly [censored](#) on mainstream Chinese social media platforms including WeChat.

In the United States English locale, we also found three names of influential figures spreading COVID-19 misinformation and anti-vaccine messages: [Ali Alexander](#), [Pamela Geller](#), and [Sayer Ji](#). These names are too few in number to be useful to generalize whether and to what extent Bing targets COVID-19 misinformation for censorship. However, Microsoft publicly [acknowledges](#) that Bing, like many other Internet operators, applies “algorithmic defenses to help promote reliable information about COVID-19.” Such control of information has also [proven controversial](#), as Internet operators attempt to balance the facilitation of free speech with controlling potentially deadly misinformation concerning COVID-19 and its prevention and treatment.

## Cross-region comparison

Previously we have only looked at Bing’s autosuggestion censorship in a few specifically chosen locales. To understand how Bing’s autosuggestions vary across the world, we expanded our testing to each region documented by Microsoft in Bing’s [API documentation](#), although we found in our testing that the documented “no-NO” (Norway) region was invalid, and so we replaced it with the valid “nb-NO” region. For each of these regions, we tested the default language for that region as obtained by not manually choosing any language for that region—for instance, the default language of Japan (“ja-JP”) when no language is manually chosen is Japanese (“ja”). While in this experiment we tested a greater number of regions and a greater number of languages, we tested a more limited set of queries compared to before. Specifically, in each locale, we tested every name that we found to be censored in any locale in our earlier experiment above. Overall, our test set consisted of 1,178 English letter names and 342 Chinese character names across 41 locales.

To compare results across locales, we used hierarchical clustering. [Hierarchical clustering](#) is a method of clustering observations into a hierarchy or tree according to their similarity. By flattening it, this tree can be represented as an ordered list where items in the list will tend to be near other items which are similar.

Hierarchical clustering requires some metric to measure the similarity or, more specifically, the dissimilarity between observations. For both English letter and Chinese character names, we hierarchically cluster according to the following dissimilarity metric. For any name query in a locale, Bing provides an ordered string of between zero and eight autosuggestions. To compare two strings of autosuggestions, we compute their [Damerau-Levenshtein distance](#), a common distance metric used on strings, where we consider two individual autosuggestions equal only if they are identical. Finally, to compare two regions’ autosuggestions for a set of names, we sum over their Damerau-Levenshtein distances with respect to their autosuggestion strings for each of these names.

In the following two sections, we look at how Bing’s autosuggestions for English letter and simplified Chinese character names vary across locales.

## Comparison of English letter names’ suggestions

In this section, by using hierarchical clustering, we look at how Bing’s autosuggestions for English letter names, including their censorship, vary across 41 different locales.

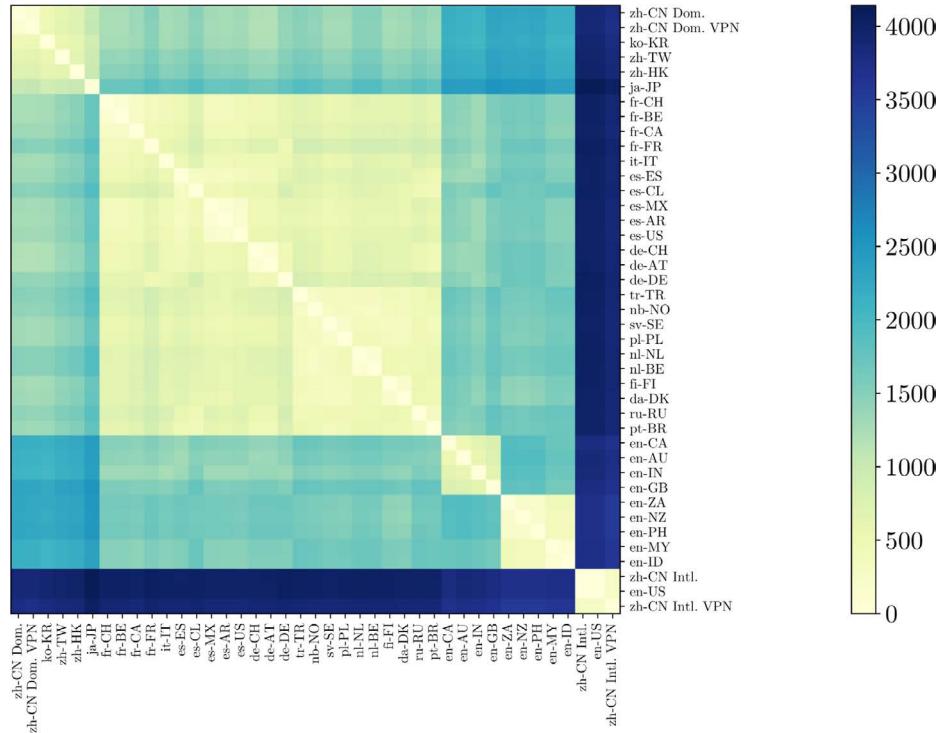


Figure 3: The distance between each locale’s autosuggestions for English letter names hierarchically clustered according to the centroid method.

As we might expect, we found autosuggestions largely clustered around both geography and language (see Figure 3). Beginning at the top left, we find a cluster of East Asian locales (“zh-CN Intl.” through “ja-JP”) as identified by a yellow square. Next, below and to the right, we find a large cluster of European-language but non-English locales (“fr-CH” through “pt-BR”) spanning Europe and Latin America in a large yellow square. Further below and to the right, we find two small clusters of English-speaking locales (“en-CA” through “en-GB” and “en-ZA” through “en-ID”). It is unclear why these locales form two clusters and not one. For instance, both clusters include British commonwealth nations and both include Eastern Hemisphere nations. However, as a whole, it is not surprising that English-language locales would generally cluster in their autosuggestions for English letter names. Finally and most noteworthy, in the bottom right corner is a cluster containing the United States (“en-US”) and the China international version as accessed from both inside mainland China (“zh-CN Intl. VPN”) and outside (“zh-CN Intl.”). This cluster is remarkable in not only how similar the United States’ autosuggestions are to those of the China international version but also how different they are from

every other locale as evident from the dark blue bars across the top side and across the left side of the plot.

Given the large number of United States Bing users, the fact that Microsoft is based in the United States, and the United States' status as an Internet hegemon, it is perhaps not surprising to see its autosuggestions different from other locales. However, what is less clear is why the United States' autosuggestions are so similar to that of China's, including in their Chinese political censorship. While we might imagine, since the international version of Bing's China search engine was developed as an English language search engine for Chinese users, that it might be a thin wrapper around the search engine in the United States, what is less obvious is how Chinese politically motivated censorship is moving in the other direction, from the international China search engine to the United States.

## Comparison of Chinese character names' suggestions

In this section, we apply hierarchical clustering toward looking at how Bing's autosuggestions for Chinese character names, including their censorship, vary across 41 different locales.

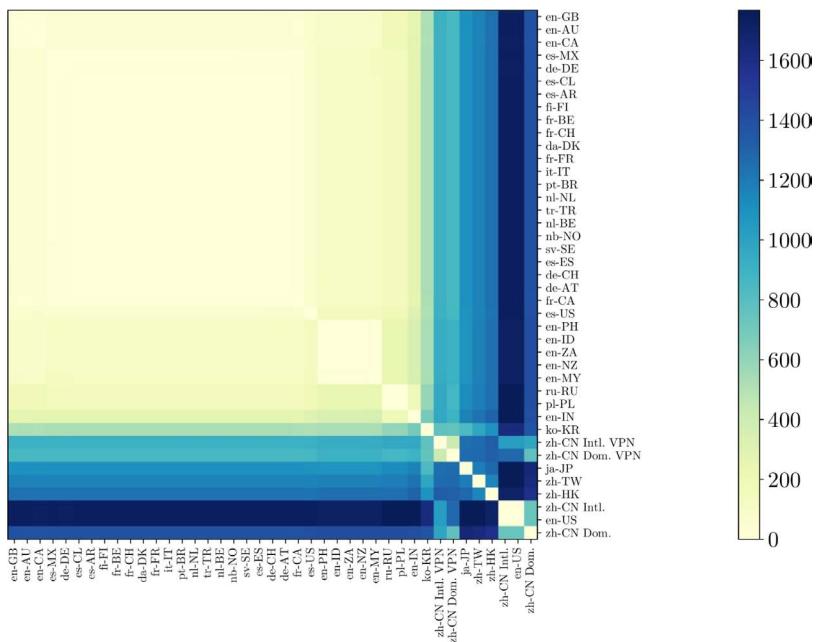


Figure 4: The distance between each locale's autosuggestions for Chinese character names hierarchically clustered according to the centroid method.

As with autosuggestions for English letter names, we found Chinese character names' autosuggestions clustered around both geography and language (see Figure 4). Beginning again at the top left, we find a large cluster of non-East-Asian language locales ("en-GB" through "en-IN"). As, outside of East Asia, Chinese characters are primarily used only by native Chinese speakers, it may be unsurprising that these regions form such a large cluster. Next, moving toward the bottom right, we find South Korea ("ko-KO") in a cluster

by itself, China's domestic and international versions as accessed from mainland China ("zh-CN Dom. VPN" and "zh-CN Intl. VPN") in a cluster, and then Japan ("ja-JP"), Taiwan ("zh-TW"), and Hong Kong ("zh-HK") each in their own singleton clusters. Finally, in the bottom right corner, we have a three member cluster containing the United States ("en-US") and both China's domestic and international versions as accessed from outside of China ("zh-CN Dom." and "zh-CN Intl.").

Unlike with English letter names, where we might imagine that Bing is reusing the United States' autosuggestions to implement their English language international version of their China search engine, it is unclear why, with Chinese character names, the United States also shares a large number of autosuggestions with mainland China's domestic search engine. However, we believe that the answer to this question might play a part in answering how the United States experiences Chinese political censorship of Chinese character names.

## Glitches in the matrix

In analyzing Bing's censorship of autosuggestions across different locales, in addition to the name censorship that we described above, we also encountered other strange anomalies. We describe a few here to help characterize the inconsistency we often observed in Bing's censorship and in the hope that such descriptions may be otherwise helpful for understanding Bing's autosuggestion censorship system.

### Jeff Widener

For some names, we found evidence of censorship, even though they did not meet all of our criteria to be considered censored. For instance, consider photographer Jeff Widener, who is famous for his photos of the June 4 Tiananmen Square protests.

"Jeff_Widene"	"Jeff_Widener"	"Jeff_Widener_"
-	-	jeff widener ou center for spatial analysiseff widener pics  jeff widener photography

*Table 15: Autosuggestions for three queries for Jeff Widener in the United States English language locale, one missing the last letter, one his full name and nothing more, and one his full name followed by a space.*

Using our stringent criteria, we do not consider his name censored in the United States English language locale, as our criteria require that three specific queries fail to

autosuggest his name, whereas in our testing only two failed to autosuggest him (see Table 15 for details). Nevertheless, his name shows signs of censorship, as two of the three queries failed to provide any autosuggestions, and it was only by adding a space to the end of his name that we were finally able to see autosuggestions. This observation is especially curious as one would generally expect fewer suggestions as one types, since typing more can exclude autosuggestions if they do not begin with the inputted text.

Looking at autosuggestions for his name in other locales, his name fails to provide any autosuggestions when queried from a mainland China IP address, whereas in the Canada English language locale, there are autosuggestions for his name in all three of the tests in Table 15. Thus, we suspect that names such as his are for some reason being filtered only on some inputs and not others in the United States English language region. The reasons for this inconsistent filtering are not clear. However, we suspect that, if we understood them, then they may shed light on why Bing censors autosuggestions in the United States and Canada for Chinese political sensitivity at all.

## Hillary Clinton and Alex Jones

In many locales, including the United States English language locale, the queries “alex\\_jone”, “alex\\_jones”, “hillary\\_clinto”, and “hillary\\_clinton” also show signs of censorship following a pattern resembling Jeff Widener. However, instead of these queries being completely censored, we found that their autosuggestions only contained their name immediately followed by punctuation such as “hillary clinton's eyes” or “alex jones+infowars” (see Figure 5).

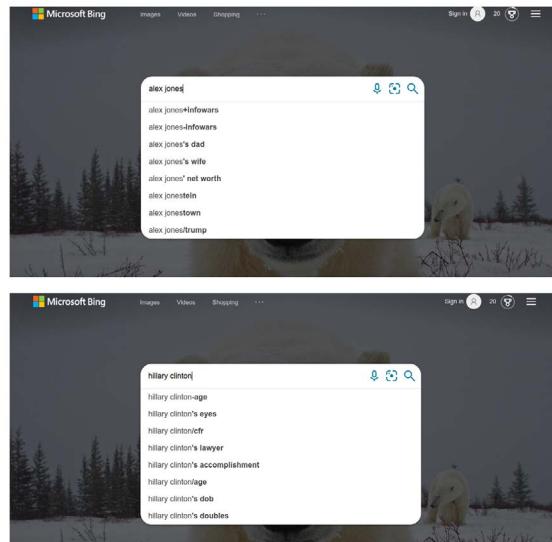


Figure 5: November 2021 autosuggestions for “alex\_jones” “hillary\_clinton” in the United States English locale never contain their bare names nor their names followed by spaces.

One might imagine how Bing could be using some poorly written regular expression which unintendedly fails to filter their names when followed by certain punctuation (e.g., by searching if `/(^|_)hillary_clinton(_|$/)` is present). Curiously, as with Jeff Widener, we found that when their names are followed by spaces (i.e., “alex\_jones\_”, “hillary\_-clinton\_”), we see autosuggestions consistent with our expectations.

Alex Jones may have been targeted for censorship for propagating COVID-19 anti-vaccinationism or other misinformation. It is unclear why Hillary Clinton may be targeted.

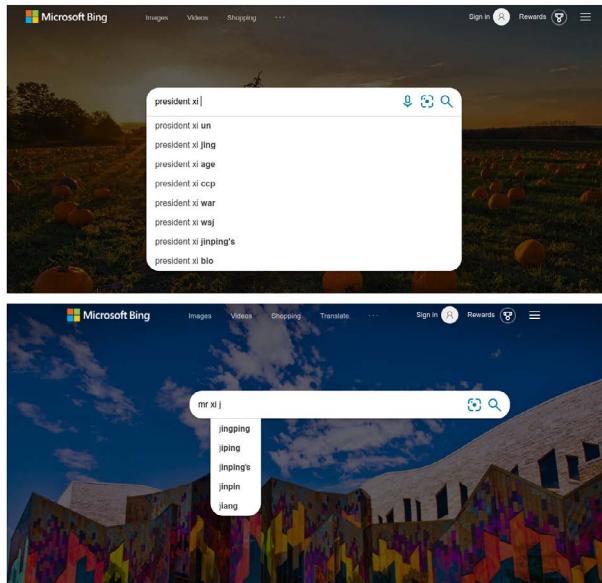


Figure 6: Top, October 2021 autosuggestions for “president\_xi\_” in the United States English locale suggested “president xi jing” and “president xi jinping’s” but seemed unable to utter “president xi jinping”; bottom, May 2022 autosuggestions for “mr\_xi\_j”, again only suggesting typos and “jinping” followed by an apostrophe.

We also found that in many regions, including the United States, the query “xi\_jinping” yielded no autosuggestions. However, once an apostrophe or other punctuation symbol is inputted after his name, then Bing displayed autosuggestions (e.g., “xi\_jinping” has no autosuggestions but “xi\_jinping’s” has autosuggestions). Sometimes the system could be coerced into providing an apostrophe-containing autosuggestion without typing one in (see Figure 6).

## Other censorship findings

While our report focuses on testing people’s names, some casual testing also reveals that other categories of proper nouns politically sensitive in China are also censored in the United States English language locale. Examples include “falun”, a reference to the Falun Gong political and spiritual movement banned in China and “tiananmen” and “june\_-fourth”, the place and day of the June 4 Tiananmen Square massacre.

Motivated by the discovery of “june\\_fourth” being censored, on May 18, 2022, we performed an experiment testing for the censorship of dates. We tested all 366 possible days of the year written as a month followed by an ordinal number, e.g., “january\\_first”, “january\\_second”, etc. We used a similar methodology as before, testing three different queries for each date. We tested from North America the following locales: China international, United States English, and Canada English.

China international version	United States English	Canada English
june\_fourth	june\_fourth	-
june\_fourth\_	june\_fourth\_	

*Table 16: In each locale tested, queries which had no suggestions.*

China international version	United States English	Canada English
june\_fourt	june\_fourt	june\_fourt
june\_fourth	june\_fourth	july\_twenty\_fourth\_
june\_fourth\_	june\_fourth\_	

*Table 17: In each locale tested, queries which were “suggestionless.” Recall that we use the term suggestionless to refer to queries where none of the autosuggestions contain original query, including if there are no autosuggestions.*

We found that in both the China international version and the United States English locale, two of the three tested queries related to “june\\_fourth” had zero autosuggestions (see Table 16). No other days had zero autosuggestions in these regions, and no days had zero autosuggestions in the Canada English locale. We also found that in both the China international version and the United States English locale, all three of the tested queries related to “june\\_fourth” were suggestionless, i.e., none of their autosuggestions contained the original query, including if there were no autosuggestions (see Table 17). No other days were suggestionless in these regions. In Canada, two queries were suggestionless: “june\\_fourt” and “july\\_twenty\\_fourth\\_”.

While we are unaware of any Chinese political significance to the date July 24, which had a single suggestionless query in Canada, the remainder of these results pertain to June 4. Due to the extreme sensitivity surrounding the June 4 Tiananmen Square Massacre, previous research has found that references to this day are one of the most [commonly censored references](#) on the Chinese Internet. Like with our investigation into Bing’s censorship of names, we are aware of no explanation for why in North America Bing would disproportionately censor this day among all days outside of Chinese-motivated political censorship.

# Censorship implementation

While our report thus far has concentrated on understanding what Bing censors, in this section we speculate on how Bing censors, specifically the filtering mechanism that Bing applies toward censoring autosuggestions. We believe that the basic censorship mechanism may work something as follows:

1. Given an inputted string, retrieve up to the top  $n$  (for some  $n > 8$ ) autosuggestions for that string.
2. Apply regular expression or other filters to these autosuggestions to remove autosuggestions which match certain patterns.
3. Among the remaining autosuggestions, display the top 8.

Such a mechanism would explain how, for instance, when one is typing in “xi\_jinping”, at first, after typing in only “xi”, there are still eight results, although, contrary to expectation, none of them containing Xi Jinping’s name. However, as one continues to type “xi\_jinping”, such as reaching as far as “xi\_j”, all of the available autosuggestions now contain Xi Jinping’s name, and thus all of the available autosuggestions are all censored by a filter targeting him, and so Bing displays no autosuggestions. While such a filtering mechanism is compatible with our findings, we are unaware of any test to conclusively determine whether this is the exact mechanism used. Moreover, this explanation does not attempt to address whatever complex ways autosuggestion censorship from some locales may be bleeding into other locales.

# Affected services beyond Bing

Due to how Bing is built into other Microsoft products and due to how other search engines source Bing data, even users who do not use Bing’s search website may still be affected by the reach of its autosuggestion censorship. In this section, we set out some of the other products and services that we found affected.

## Windows Start Menu

We found Windows’ Start Menu search restricted by Bing’s censorship in both Windows 10 and Windows 11. This feature is accessed in Windows by opening the Start Menu and then typing, and it is used by Windows users to not only find search results from the Web, but it is also a common way to search for locally installed apps or locally stored documents.

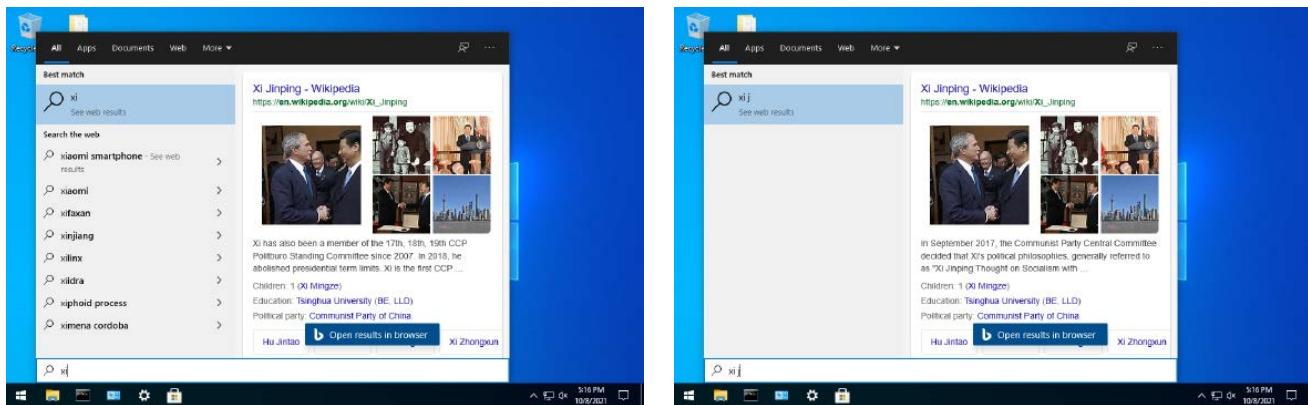


Figure 7: In the United States English locale, Windows Start Menu search shows autosuggestions for “xi” not related to Xi Jinping and not any autosuggestions for “xi\_j”

Windows Start Menu censorship varies depending on how the Windows region settings are configured. We found that the Bing autosuggestions displayed in the Start Menu appeared consistent with the region selected as the Windows “Country or region” in Windows’ “Region” settings. For instance, when Windows is configured with the United States region, we found all autosuggestion results for “xi\_j” censored (see Figure 7) but not when configured with the Canada region.

Of note, the Windows Start Menu appears to strip trailing spaces from queries. Thus, we were unable to directly test “xi\_”. However, in our experience, for purposes of autosuggestions, Bing interprets underscores as equivalent to spaces, and we found that “xi\_” (“xi” followed by an underscore) produced no autosuggestions in consistency with the Web interface.

## Microsoft Edge

Microsoft Edge is a cross-platform browser available for Windows, MacOS, and Linux that comes installed in the latest versions of Windows. We found that, by default, Edge uses Bing for its built-in search functionality and that its autosuggestions are also censored.

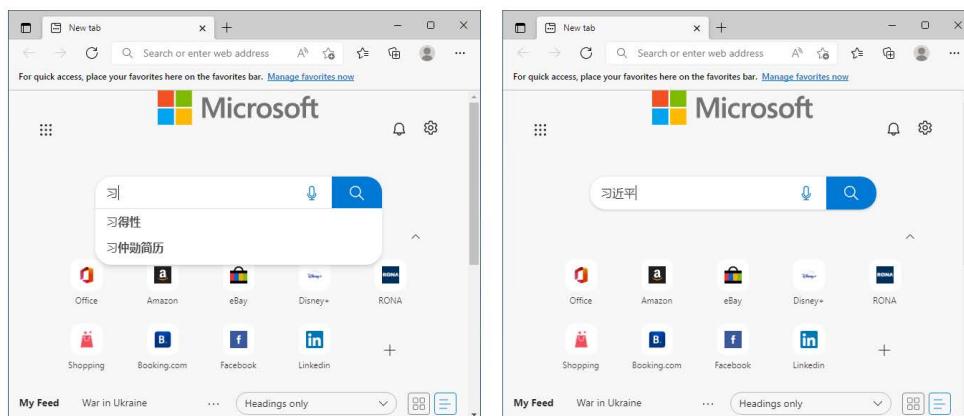


Figure 8: In the Canada English locale, Microsoft Edge shows no autosuggestions containing “习近平” (Xi Jinping).

Edge appears to use one's IP geolocation to select a default region, and we were unable to find any setting to change the region governing the search autosuggestions of Edge's built-in Bing implementation.

## DuckDuckGo

Billed as a privacy-protecting search engine, we found that [DuckDuckGo](#) autosuggestions are nevertheless affected by Bing's autosuggestion censorship. This finding is likely due to DuckDuckGo having a [close partnership](#) with Bing for providing data.

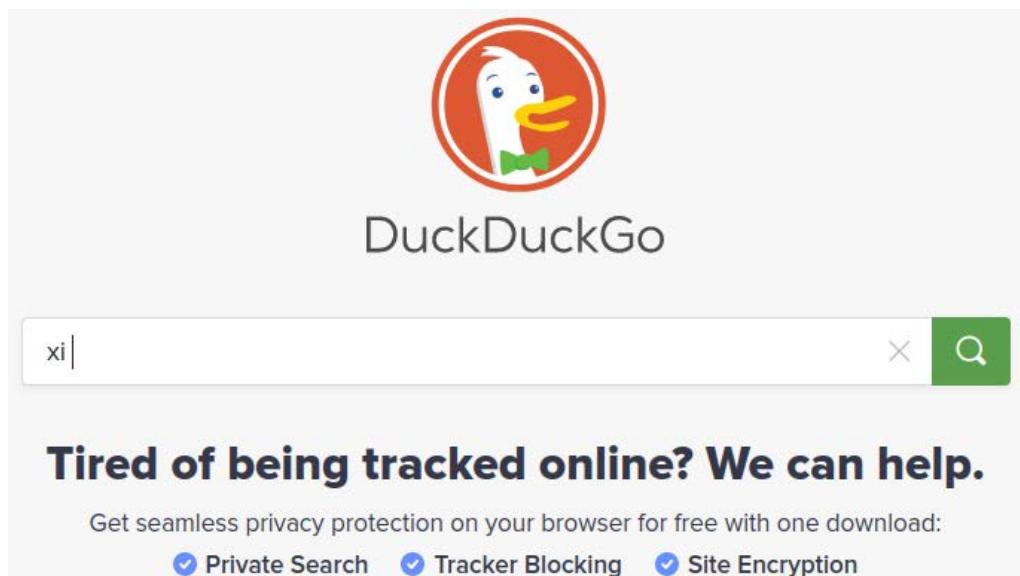


Figure 9: DuckDuckGo provides no autosuggestions for “xi\_”

Although we did not extensively analyze censorship on DuckDuckGo, we found, for instance, that DuckDuckGo provides no autosuggestions for “xi\_” in its default region setting of “All regions” when we browsed from Canada. When explicitly setting DuckDuckGo's region to that of the United States or Canada, the autosuggestions appeared to track the autosuggestions that Bing's Web interface directly provides in those regions, including any censorship.

## Autosuggestion censorship impact

How do autosuggestions and their censorship impact search behavior? One way to approach this question is to find a large sample of people and divide them into two groups, providing autosuggestions to one and not providing them to another and measuring how searches differed in each group. Another approach is to offer one group spurious autosuggestions while giving organic autosuggestions to another. While these approaches would work, they are unnecessary, as Bing, when it disables/enables autosuggestions in a country or offers autosuggestions that no one would organically search for,

is already performing such an experiment for us. Moreover, by using Bing’s historical search volume data, we can measure such an experiment’s results. Thus, to better understand how autosuggestions and their censorship influence search behavior, we analyze Bing’s historical search volume data with respect to the following three data points: (1) the spurious autosuggestions for Hillary Clinton and Alex Jones and the shutdown of autosuggestion functionality in China in mid-December 2021 to early January 2022 and its effect on (2) benign searches for food and (3) searches for the controversial Falun Gong movement.

The spurious autosuggestions for Hillary Clinton and Alex Jones allow us to examine how spuriously introduced autosuggestions influence search trends, whereas the autosuggestion shutdown in China allows us to examine how the absence of autosuggestions influence search trends. Unlike previously in our report when we took a rigorous statistical approach, in the remainder of this section we will only introduce different cases where autosuggestions seem to have influenced what users search for, as a more rigorous analysis of how autosuggestions shape search behavior is outside the scope of this work.

In our research for this section, we also discovered a misleading flaw in the way that Bing’s Keyword Research tool visualizes search volume. As a result, in this section we relied entirely on the raw data returned by Bing’s API. We detail this issue in the Appendix.

## Hillary Clinton and Alex Jones

In this section, we look at how artificially introduced autosuggestions alter search behavior. Specifically, we look at social trend data surrounding the spurious autosuggestions for “Hillary\_Clinton” and “Alex\_Jones” to look at how artificially introduced autosuggestions influence search results. During our December 2021 measurements, Bing provided the following autosuggestions for “Hillary\_Clinton” in the United States English locale:

- hillary\_clinton-age
- hillary\_clinton's\_age
- hillary\_clinton's\_boyfriend
- hillary\_clinton's\_meme
- hillary\_clinton's\_accomplishment
- hillary\_clinton's\_women's\_rights\_speech
- hillary\_clinton's\_college
- hillary\_clinton/age

As we discussed earlier in the report, Bing’s autosuggestions for Hillary Clinton are anomalous, as Clinton’s name is never suggested by itself nor followed by a space and in

every suggestion her name is followed by punctuation. While many searches beginning with “hillary\_clinton’s” (“hillary\_clinton” followed by an apostrophe and “s”) such as “hillary\_clinton’s\_age” may be naturally typed, we find it unlikely that there are many queries typed in with her name followed by hyphens or slashes such as “hillary\_clinton-age” or “hillary\_clinton/age”. Nevertheless, we find these spurious autosuggestions suggested by Bing.

Idea	Query	Symbol following name	Suggested?	Search volume
Hillary Clinton's age	hillary_clinton-age	Hyphen	Yes	13,710
	hillary_clinton_age	Space	No	10,212
	hillary_clinton's_age	Apostrophe	No	1,637
	hillary_clinton/age	Slash	Yes	51

Table 18: Queries relating to Hillary Clinton’s age, whether they were autosuggested by Bing upon inputting “Hillary\_Clinton”, and their search volume between October 2021 and April 2022.

In Table 18, we can see that the spurious autosuggestions resulting from Clinton’s name were competitive against the organic search queries. Notably, the query where her name was followed by a hyphen received more search volume than the query where her name was followed by a space or apostrophe, although we do find that the query where her name was followed by a slash received the least search volume. While this finding provides evidence that users do make use of autosuggestions when searching, it does not necessarily tell us to what extent autosuggestions influence users’ behavior at a high level. After all, it is possible that most of the users who clicked on “hillary\_clinton-age” intended to search for Hillary Clinton’s age anyways, just using more typical punctuation. Nevertheless, the search volume dominance of the hyphen-containing query suggests that, to whatever end, users do click on autosuggestions.

We found that Alex Jones’s autosuggestions in the United States English locale follow a similar pattern:

- alex\_jones
- alex\_jones+infowars
- alex\_jones-infowars
- alex\_jones-youtube
- alex\_jones's\_house
- alex\_jones's\_dad
- alex\_jones's
- alex\_jones/911

Idea	Query	Symbol following name	Suggested?	Search volume
Alex Jones & Infowars	alex_jones+in-fowars	Plus sign	Yes	41,165 (trend data cannot distinguish between spaces and plus signs)
	alex_jones_in-fowars	Space	No	
	alex_jones-in-fowars	Hyphen	Yes	19,438
Alex Jones & YouTube	alex_-jones-youtube	Hyphen	Yes	150
	alex_-jones_youtube	Space	No	20
Alex Jones & September 11	alex_jones/911	Slash	Yes	0
	alex_jones_911	Space	No	0

Table 19: Queries for Alex Jones, whether they were autosuggested by Bing upon inputting “Alex\_Jones,” and their search volume between October 2021 and April 2022.

In Table 19, we can again see spurious autosuggestions motivating a large amount of search volume. Due to limitations in Bing’s API in which spaces and plus signs are treated identically, we cannot distinguish between search volume for queries in which his name is followed by a space versus those where his name is followed by a plus sign. However, for queries related to Jones and his show “Infowars,” the query where his name is followed by a hyphen has almost half as much search search volume as those of his name followed by either a space or plus sign combined. Regarding queries for Jones and his YouTube videos, the query where his name is followed by a hyphen has 7.5 times as much search volume as the one followed by a space. Again, we cannot say if these spurious hyphen-containing queries are pulling users away from other searches that they might have typed concerning Alex Jones or if users already intending to find Jones’s YouTube videos are simply clicking on the oddly punctuated autosuggestion instead of typing out a more naturally punctuated one. Nevertheless, as with our findings concerning Clinton, these findings concerning Jones show that users do click on Bing’s autosuggestions.

## Shutdown of autosuggestion functionality in China

Microsoft’s shutdown of autosuggestion functionality in China offers us the opportunity to test how autosuggestions affect search behavior by allowing us to compare search behavior in their presence versus their total absence. We begin by looking at the autosuggestions for three benign words: 洋葱 (onion), 西瓜 (watermelon), and 榴莲 (durian).

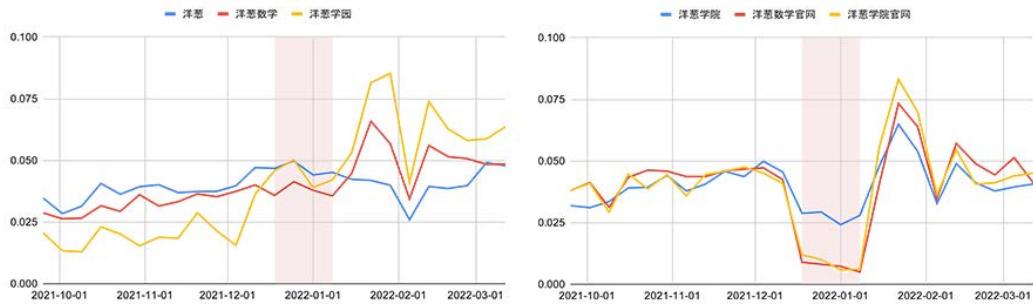


Figure 10: Search trends for autosuggestions for 洋葱 (onion) not affected (left) and affected (right) by the shutdown.

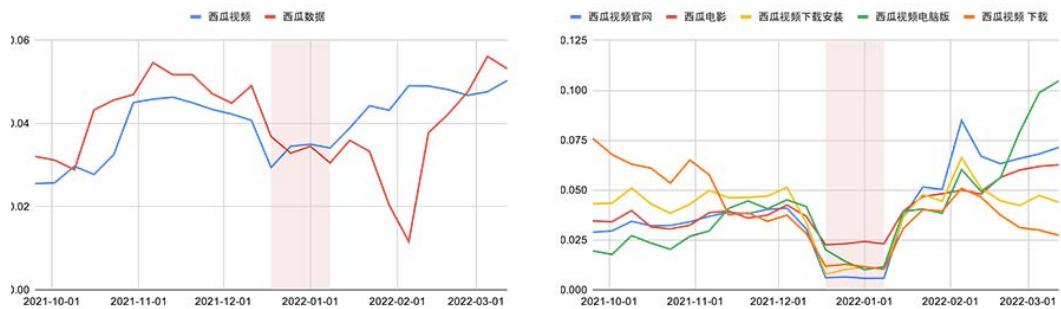


Figure 11: Search trends for autosuggestions for 西瓜 (watermelon) not affected (left) and affected (right) by the shutdown.

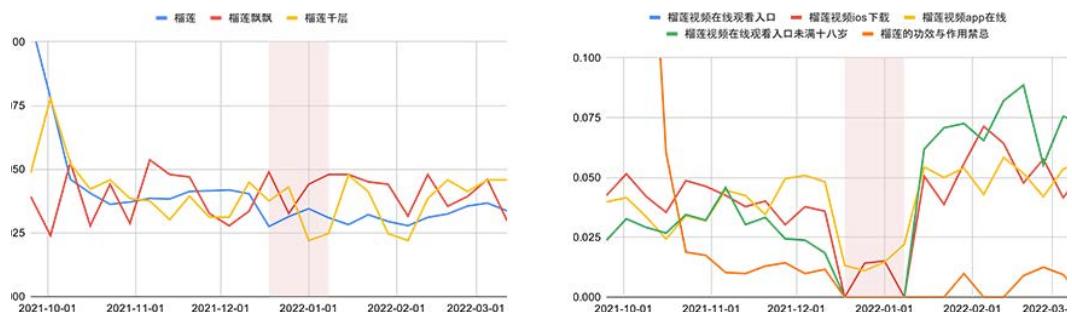


Figure 12: Search trends for autosuggestions for 榴莲 (durian) not affected (left) and affected (right) by the shutdown.

In general, we find that the autosuggestions most affected by the shutdown are longer (see Figures 10, 11, and 12). We hypothesize that these are less likely to be naturally searched and thus benefit more from being autosuggested.

## Falun Gong

We now turn our attention to Bing's autosuggestions in China for the controversial 法轮功 (Falun Gong) spiritual and political movement. We found that, in addition to “法轮功/flg”, where “flg” is an abbreviation of Falun Gong, China Bing provided the following autosuggestions for “法轮功” (Falun Gong):

- “法轮功\_危害” (Falun Gong dangers)

- “法轮功常见非法宣传品” (Falun Gong Common Illegal Propaganda Materials)
- “法轮功什么时候被国家取缔” (When was Falun Gong banned by the state?)
- “法轮功邪教组织什么时候被取缔” (When was the Falun Gong cult banned?)



Figure 13: Search trends for 法轮功 (Falun Gong) (left) and for its autosuggestions (right).

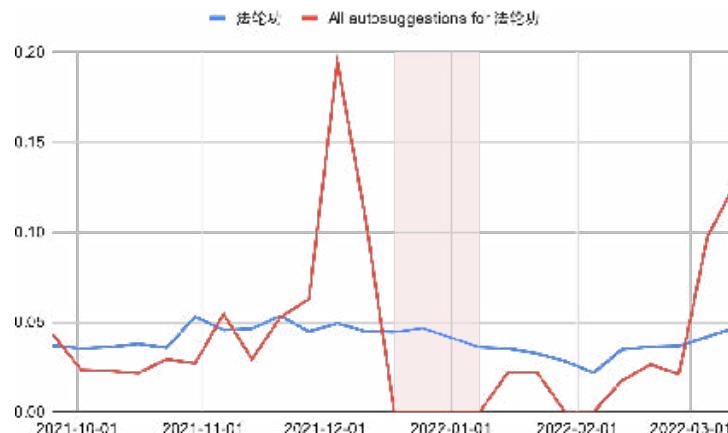


Figure 14: In blue, the trend data for 法轮功 (Falun Gong), and in red, the aggregate (sum) data for all of its suggestions.

While Bing reports erratic trends data for Falun Gong related autosuggestions, we find that none of these autosuggestions for Falun Gong had any trend data during the four weeks of Bing's autosuggestion shutdown in China (see Figures 13 and 14). This finding in addition to the seemingly artificiality and one-sidedness of these autosuggestions suggests that these autosuggestions are influencing searches rather than being induced by them.

To compare Bing's autosuggestions for “法轮功” to other regions, we find that Bing provides no autosuggestions in the United States English and Canada English locales, consistent with the politically motivated censorship in these locales that we measured earlier in our report. Chinese search engine Baidu also provides no autosuggestions for “法轮功”. However, Google in Canada provides the following autosuggestions:

- 法轮功\_英文 (Falun Gong in English)
- 法轮功\_自焚 (Falun Gong self-immolation)
- 法轮功官网 (Falun Gong official website)
- 法轮功是什么 (What is Falun Gong)
- 法轮功真相 (The truth about Falun Gong)
- 法轮功创始人 (Founder of Falun Gong)
- 法轮功媒体 (Falun Gong media)
- 法轮功\_资金来源 (Falun Gong sources of funding)
- 法轮功现状 (Status of Falun Gong)
- 法轮功电影 (Falun Gong Movies)

We must be careful in comparing the autosuggestions of Bing versus Google, as these companies use different algorithms and are used by different populations, and thus we would already expect them to have different autosuggestions for completely benign reasons. However, we do find that Google's autosuggestions are shorter, consistent with something that might be organically typed in by a user, and that the autosuggestions are well balanced, with many supporting of and many critical of the Falun Gong movement.

While it makes sense to ask whether Bing introduces autosuggestions at the behest of the Chinese government, the question may be of little practical significance. Whether Bing introduces autosuggestions or whether they censor all but certain autosuggestions, the practical result would be the same, that Bing is artificially influencing searches in keeping with China's propaganda requirements.

## Limitations

We found that in each region the names censored by Chinese political censorship varied over time. To measure a consistent snapshot of Bing's censorship, we tested during a short period of time (one week in December, 2021). However, many of the examples we illustrate may no longer be censored, or other examples which we found not to be censored may now be. Later in this report, we discuss what the inconsistency of the Chinese political censorship across different regions says about why such censorship is affecting regions outside of China.

In our statistical analysis, there may exist confounding variables that we failed to account for. For English letter names, we tested whether Bing was merely failing to provide autosuggestions for sensitive Chinese political names in the United States because such names were more likely to be written in pinyin, a type of name that Bing may have struggled

to provide autosuggestions for. However, we found that pinyin names failed to account for Bing's Chinese political censorship. Moreover, we are unaware of any confounding variables that might explain why Bing politically censors Chinese character names. Even if such a confounding variable existed that lent some innocuous explanation for Bing's Chinese political censorship, it would not change our finding that Bing disproportionately fails to provide autosuggestions for the names of people who are politically sensitive in China, regardless of explanation.

Using our methodology, a small number of individuals' names appeared censored by Bing under no motivation that we could identify. One reason for this could be that we simply failed to identify some straightforward motivation. Another reason is that the name was collaterally censored and that we failed to recognize that the individual's name included letters consisting of profanity in English or some other language. Although we did discover names collaterally censored for containing letters commonly considered profane in non-English languages, we had limited ability to exhaustively recognize such names. Finally, the name may have simply appeared censored due to being a false positive, somehow not having any autosuggestions despite having both high Wikipedia traffic and large Bing search volume but yet also not being targeted for censorship by Bing in any region. However, since such false positives are, by definition, not being targeted by Bing, we would expect such failures to not be disproportionately politically sensitive in China compared to names chosen uniformly at random, and thus such names would be unable to explain our statistical findings.

## Discussion

Our research shows that Bing's Chinese political censorship of autosuggestions is not restricted to mainland China but also occurs in at least two other regions, the United States and Canada, which are not subject to China's laws and regulations pertaining to information control. To our knowledge, Bing has not provided any public explanation or guidelines on why it has decided to perform censorship in various regions or why it has censored autosuggestions of names of these individuals.

On May 10, 2022, we sent a letter to Microsoft with questions about Microsoft's censorship practices on Bing, committing to publishing their response in full. Read the [letter here](#) and the email [response here](#) that Microsoft sent on May 17, 2022.

Search engines are a major interface between users and the Internet. They serve as a gateway to online information, which to a large extent [influences](#) user attention and knowledge. The [impacts](#) of search engine results on the visibility and identity of a person, an organization, and even a country have been [documented](#) by many. Bing's censorship, therefore, not only affects how users perceive certain entities but also dictates whether

users get to know the existence of these entities. What might explain Bing's censorship decisions in and beyond mainland China? We evaluate multiple hypotheses below.

## Does Bing's development team follow Chinese principles and norms concerning political expression?

After Microsoft's 2021 blocking of images relating to the 1989 Tiananmen Square Movement, journalists speculated whether much of Bing's development team being [based in China](#) contributed to that mistake. There are previous instances of China-headquartered companies and companies with development teams in China implementing censorship targeting Chinese political sensitivity on global-facing products. ByteDance's TikTok, for example, was [found](#) to instruct moderators to censor videos referencing Tiananmen Square, Falun Gong, and other topics considered politically controversial in China despite it not offering services to China-based users. In 2020, a TikTok executive [admitted](#) that the app censored content [critical](#) of the Chinese government but insisted that it has terminated relevant content moderation policies favoring the views of the Chinese government. Since much of Bing's development team is based in China, Bing's developers, who are likely trained to follow Chinese laws and regulations in their everyday practices, may have similarly believed it appropriate to censor such content globally or may have been less likely to recognize if such content were being restricted accidentally.

## Is Microsoft censoring outside of China to appease the Chinese government?

We might suspect that Microsoft performs Chinese political censorship globally in order to appease the Chinese government, perhaps as part of the conditions in which Microsoft is allowed to continue operating Bing and other Microsoft services inside of China. However, accepting this hypothesis requires caution for the following reasons. First, there is no explicit evidence suggesting Bing's censorship behavior is resulting from concessions made to the Chinese government. While [much is known](#) about the concessions that fellow North American tech giant Apple made to be allowed to operate in China, comparatively little is known about those made by Microsoft. Second, in a previous instance in which Bing was [found](#) to censor image results for the query "tank man" in North America, Microsoft attributed the censorship to "accidental human error" and quickly ceased performing the censorship. Finally, the inconsistency over time of which sensitive Chinese political names are censored in North America may suggest that this censorship is unintentional and the emerging property of some complex system rather than solely the deterministic application of a set of rules.

Nevertheless, while Microsoft may not be performing Chinese political censorship globally to appease the Chinese government, they are certainly performing Chinese political censorship in China to appease the Chinese government. Bing's global Chinese

political censorship would seem inextricably linked to Microsoft’s censorship operations in China. While it is unclear why Microsoft’s Chinese political censorship is leaking outside of China, it seems that it could not have leaked out had Microsoft never engaged in such censorship operations in the first place.

## **Is Bing’s censorship of users globally an unintentional result of their censorship of Chinese users?**

We might suspect that Bing’s censorship of users in China is impeding these users from making certain search queries, and, since historical search behavior is used to inform autosuggestions, the search behavior of censored Chinese users is being used to inform the autosuggestions for users globally. Search engines’ autosuggestions are based largely on ranking signals which include, among other indicators, how many users have submitted a certain keyword for search and how many times the suggestions have been selected in the past. In our earlier cross-region comparison, we discovered that the autosuggestions between mainland China and the United States shared a large amount of overlap, suggesting that a shared data source was used to inform autosuggestions in both regions. Thus, it is possible, since Bing’s censorship of autosuggestions in China affects the search behavior of Chinese users, and since autosuggestions in other regions are based, in part, on the search behavior of Chinese users, that this is how Bing’s autosuggestion censorship is leaking globally.

However, even if this mechanism played a role in explaining the behavior that we observed, it would not seem to be a complete explanation. While Bing’s autosuggestion censorship may effectively hide sensitive political names such as those of dissidents from Chinese users, it also has a second purpose which is to hide sensitive suggestions concerning Communist Party leaders. While autosuggestion censorship might prevent a Chinese user searching for “Xi Jinping” from learning about negative narratives surrounding the paramount leader such as the Winnie the Pooh mockery, it would not generally prevent a user from searching for Xi Jinping in general. Thus, this hypothesis seems unable to explain the global autosuggestion censorship of any reference to Communist Party leaders such as Xi Jinping.

## **Is Chinese users’ self-censorship affecting autosuggestions globally?**

We might also wonder if Bing’s global Chinese political censorship is resulting from users in China self-censoring their searches and if this self-censorship search behavior is being used to inform the autosuggestions of users outside of China. However, this hypothesis appears to be, at best, an incomplete explanation for similar reasons as the previous hypothesis. While some names are inherently sensitive, such as Liu Xiaobo, others are not inherently sensitive, such as Xi Jinping or the names of other Communist Party leaders. Thus, it does

not explain why we do not see autosuggestions for Xi Jinping in the United States when Chinese users would have no reason to self-censor their searches for Xi Jinping.

## Can Microsoft even permanently address this problem?

The idea that Microsoft or any other company can operate an Internet platform which facilitates free speech for one demographic of users while intrusively applying political censorship to another demographic of its users may be fundamentally untenable. Our prior research has previously documented failures in such attempts.

In previous work, we discovered that not only did the Chinese version of Microsoft's [Skype](#) perform extensive keyword-based political censorship and surveillance but that such surveillance logs containing users' private messages were inadvertently [exposed to the public](#). Skype's censorship and surveillance also applied globally to users outside of China when they were communicating with Chinese users, with or without their knowledge, and it finally ended when Microsoft abandoned bundling censorship and surveillance in its Chinese version of Skype.

Tencent's WeChat also operates a platform in which it politically censors users in China while attempting to facilitate uncensored speech to users outside of China. Like with Skype, WeChat users out of China can have their communications unknowingly captured by WeChat's censorship and surveillance system when they are communicating with China-based users. However, [recent work](#) also discovered that, even when two WeChat users who are both outside of China are communicating between each other, their communications were silently being surveilled for Chinese political sensitivity and used to train and build up WeChat's political censorship apparatus that it applies to users in China

[Recent work](#) documents Apple's failure to contain their Chinese political censorship of users' product engraving text to only China. In our work we discovered that Apple applied mainland Chinese political censorship to users in Taiwan for reasons most likely stemming from negligence. Since our [first report](#), Apple has ceased political censorship of Taiwan users' engravings. However, [we found](#) that Apple continues to politically censor users in Hong Kong. As other North American tech companies do not perform similar political censorship in Hong Kong, we speculated over possible motivations Apple may have for performing it, including appeasement of the Chinese government.

We may hope that resolving whichever series of decisions, errors, bugs, or glitches which gave rise to the issues we discovered in this report might fix these issues. However, this is not the first time that Microsoft has allowed Chinese political censorship on Bing to be applied to users globally, and there is little assurance that it would be the last. In light of our past research, **the findings in this report again demonstrate that an Internet platform cannot facilitate free speech for one demographic of its users**

**while applying extensive political censorship against another demographic of its users.** One method for Microsoft to assure its users outside of China that they are free from its Chinese political censorship would be for Microsoft to shutter its Chinese political censorship operations in their entirety. However, Microsoft may find this solution undesirable because doing so would mean violating content control laws and regulations in China and upsetting the Chinese government, which would fundamentally affect Microsoft's operations in the Chinese market. A second approach would be for Microsoft to bifurcate their search operations into two completely separate operations, similar to what Microsoft has recently done with their LinkedIn platform. While bifurcating a globalized product also further splinters the Internet and risks cutting off China-based users access to external information, absent such changes, it is unclear whether Microsoft can feasibly operate an Internet platform in which it can provide free speech for some of its users while violating such rights for the rest.

## Availability

The list of names that we discovered to be censored in each locale that we tested is available [here](#).

## Appendix: Bug in Keyword Research tool

In this section, we describe a subtle but potentially misleading flaw in the way that Bing's Keyword Research tool visualizes data.



Figure 15: Top, Bing's Keyword Research Tool results for “李文亮的英雄事迹” (Li Wenliang's heroic deeds) which fails to plot weeks with zero search volume; bottom left, our plot of the raw data returned by Bing's API attempting to reproduce the Research Tool's bug by excluding zero-volume weeks; bottom right, our plot of the raw data returned by Bing's API correctly including zero-volume weeks.

We found that the plot generated by the tool incorrectly collapsed weeks without any search volume, making the plot look continuous even if a large gap existed in search volume. The only hint that the tool gives that the plot is missing data is that, despite a certain range (e.g., six months) being requested, a shorter range is visualized (e.g., in Figure 15, mid-October to mid-January, as opposed to the six month range of mid-October to mid-April that was requested). Erroneously, the plot thus shows search volume for weeks which actually had none.

## Appendix: Comments on Microsoft's response

We would like to thank Microsoft for considering and replying to our letter. With respect to some of our findings from our December 2021 experiment being no longer reproducible in May 2022, we recognize in our report that the censorship of autosuggestions which we characterize through our research fluctuates over time. However, we have observed that the direction of fluctuation is not always in the direction of reducing censorship.

With respect to Microsoft's discovery and resolution of a misconfiguration preventing valid autosuggestions from appearing, we are happy that our research led to the discovery and resolution of such a misconfiguration. However, aside from general fluctuations, we are unaware of any change in Bing's tendency to censor autosuggestions which are politically sensitive in China. For example, in the "Other censorship findings" section, we perform an experiment testing all 366 possible days of the year, finding only the most sensitive day in China, "june\_fourth", the day of the Tiananmen Square Massacre, was censored in the United States English locale. This experiment was performed on May 18, 2022, after the receipt of Microsoft's response.

