Banned Books: Analysis of Censorship on Amazon.com

Jeffrey Knockel
Bowdoin College / Citizen Lab
j.knockel@bowdoin.edu

Jakub Dałek*
Citizen Lab, University of Toronto
jakub@citizenlab.ca

Noura Aljizawi
Citizen Lab, University of Toronto
noura@citizenlab.ca

Mohamed Ahmed
Citizen Lab, University of Toronto
mohamed.ahmed@citizenlab.ca

Levi Meletti
Citizen Lab, University of Toronto
25rxd@queensu.ca

Justin Lau
Citizen Lab, University of Toronto
justin@citizenlab.ca

Abstract

We analyze the system Amazon deploys on the US "amazon.com" storefront to restrict shipments of certain products to other regions in the world. We found 17,842 products that Amazon restricted from being shipped to at least one world region. While many of the shipping restrictions are related to regulations involving WiFi, car seats, and other heavily regulated product categories, the most common product category restricted by Amazon in our study was books. Banned books were largely related to LGBTIQ, the occult, erotica, Christianity, and health and wellness. The regions affected by this censorship were the UAE, Saudi Arabia, and many other Middle Eastern countries as well as Brunei Darussalam, Papua New Guinea, Seychelles, and Zambia. In our Common Crawl test sample, Amazon restricted shipment of over 1.1% of the books sold on amazon.com to at least one of these regions. We identified three major blocklists which Amazon assigns to different regions. In numerous cases, Amazon's restrictions were either overly broad or miscategorized. Examples include the restriction of books relating to breast cancer, recipe books invoking "food porn" euphemisms, and Nietzsche's *Gay Science*. To justify why restricted products cannot be shipped, Amazon used varying and often misleading error messages such as by conveying that an item is temporarily out of stock.

We reported our findings to Amazon. A year later, we found that Amazon reduced its use of misleading error messages and had eliminated many of its false positive restrictions. Although we cannot conclusively link Amazon's improvements to our reporting, we believe that our findings serve as evidence that measurement studies can effect change in problematic company behavior.

1 Introduction

The rise in online shopping has led to more global reach into markets that may otherwise be inaccessible for companies through traditional retail channels. This increased reach brings new opportunities but also has its own challenges for global e-commerce retailers. One such challenge is in dealing with different, more restrictive regulatory environments worldwide that seek to control access to various information.

In this paper, we analyze American e-commerce retailer Amazon's system for preventing shipments of certain products to certain world regions as it is implemented on the US storefront — amazon.com. Specifically, we analyze the functionality that Amazon implements to *restrict* shipments of certain products to certain regions even if the product is available and sellers are offering to ship it there. While Amazon normally hides this restriction system from customers using opaque error messages, we employ a novel methodology to uncover and measure on which products and in which regions it is activated by peeling back the layers of Amazon's website and analyzing its internal workings. Notably, our method can distinguish between a product being restricted by Amazon and it being organically unavailable in a region.

Our work presents the following key findings:

- We analyzed the system Amazon deploys on the US "amazon.com" storefront to restrict shipments of certain products to specific regions, finding 17,842 products restricted from being shipped to at least one world region.
- While many of the shipping restrictions are related to regulations involving WiFi, car seats, and other heavily regulated product categories, the most common product category restricted in our study was books.
- Banned books were largely related to LGBTIQ, the occult, erotica, Christianity, and health and wellness. The regions affected by this censorship were the UAE, Saudi Arabia, and many other Middle Eastern countries as well as Brunei Darussalam, Papua New Guinea, Seychelles, and Zambia. In our Common Crawl test sample, 1.1% of the books sold on amazon.com were restricted from being shipped to at least one of these regions.
- Using clustering, we identified three major censorship regimes to which Amazon's system bins different regions. In numerous cases, the resulting restrictions are either overly broad or miscategorized. Examples include the restriction of books relating to breast cancer, recipe books invoking "food porn" euphemisms, and Nietzsche's *Gay Science*.
- To justify why restricted products cannot be shipped, Amazon used varying error messages such as by conveying that an item is temporarily out of stock. A year later, after reporting these and other issues to Amazon, we found that Amazon improved the accuracy of its messages and reduced censorship false positives, providing evidence for how measurement studies can effect change in problematic company behavior.

The remainder of this work is structured as follows. In § 2 we situate our work in the existing literature on network and platform

*The author's contributions were made in 2024 and earlier while at Citizen Lab.

Proceedings on Privacy Enhancing Technologies YYYY(X), 1–15

https://doi.org/XXXXXXX.XXXXXXX

censorship. In § 3, we provide background on Amazon and its regulatory pressures. § 4 summarizes our experiment design, and §§ 5, 6, and 7 detail each phase of our three-phase experiment. We discuss the limitations of our work in § 8. We detail the ethical considerations of our work in § 9 and conclude our work in § 10.

2 Related work

An existing body of work on network censorship develops measurement frameworks for detecting and characterizing access restrictions varying across regions and vantage points, separating intentional interference from benign unavailability, inferring the presence of centralized filtering, keyword triggers, and collateral censorship [10, 17, 33, 40, 45]. This work also underscores that unavailability can arise from restrictions implemented above network protocol layers, such as server-side blocking and geoblocking [21, 31, 44]. Within these bodies of work, multiple studies have specifically analyzed censorship in the Middle East and Iran. The Open Net Initiative found broad filtering in eight countries in the region [19]. Dainotti et al. [13] used BGP and traffic data to study Egypt's 2011 near-total Internet shutdown. Much work on Iran has identified highly centralized and especially broad filtering [8, 9, 18, 30].

Analyzing a layer of enforcement above network and server reachability, prior work has measured the compliance of American multinational corporations to the censorship demands of governments. GreatFire has studied Apple's political censorship of its App Store in China where it has taken down major VPN apps, foreign news apps, LGBTIQ dating apps, and encrypted messaging apps [43]. Knockel and Ruan [27] found Apple using over 1,000 keyword censorship rules on its product engraving service in six world regions, inappropriately applying mainland Chinese censorship outside of the region. Apple was also reported to have taken down an app that allowed protesters to track the police from its Hong Kong App Store [34]. The company has also removed a number of Hong Kong originating songs and artists from Apple Music in mainland China [11] as well as hiding Taiwan's flag emoji for users in Hong Kong and Macau [37].

Microsoft has also been found applying censorship to different world regions on a number of their products. Knockel and Ruan [28] found that Microsoft censored Bing search suggestions for the names of Chinese party leaders, dissidents, and other persons and topics of Chinese-political sensitivity in its service across the globe, i.e., even outside of China. There has also been a number of studies outlining political censorship on Skype [22], Bing search results [26], and Bing Translate [38]. LinkedIn, a Microsoft subsidiary, was also found to censor posts or profiles for political reasons [32].

Several measurement studies have empirically examined how Twitter/X complies with state-imposed censorship by authoritarian governments. Tanash et al. [42] conducted an early large-scale analysis, collecting over 20 million Turkish tweets and demonstrating that Twitter under-reports censored content by two orders of magnitude compared to its transparency disclosures. Following the 2016 failed coup in Turkey, Tanash et al. [41] extended this work and showed a 72% drop in publicly detectable censored tweets, which they attribute partly to increased self-censorship among users, quantifying that 41% of users voluntarily removed content on politically sensitive topics. Recently, Çetinkaya and Elmas [47] presented a longitudinal study of withheld accounts, modeling predictors of takedown compliance.

Our work contributes to these bodies by applying censorship measurement to a new layer of information control enforcement: digitally enforced restrictions of commercial shipments. We measure the cross-border availability of information (e.g., books) and other items via platform-enforced shipping restrictions. As in the earlier cited network and platform measurement studies, we use special signals to distinguish the enforcement of intentional restrictions from benign causes of unavailability, and we identify and characterize centrally and automatically enforced censorship rules that often produce collateral censorship.

3 Background

In this section we briefly describe Amazon's history as it relates to our analysis and Amazon's policies and practices when it comes to compliance with international regulations.

3.1 Amazon background

Amazon is best known for its original website — amazon.com — which serves as the landing page for US customers, although sellers can also permit their items to be shipped globally. As of 2025, there are dedicated storefronts for 24 other regions [6]. Alongside the online expansions to other regions there has been an analogous expansion of physical infrastructure in those regions including shipping hubs, fulfillment centers, sorting facilities, and delivery stations.

Most relevant to our study, Amazon has expanded its dedicated storefronts to include the UAE in 2019 and Saudi Arabia in 2020 [46]. This expansion included opening a regional headquarters in Riyadh, Saudi Arabia, in 2022 [16] and a fulfillment center in Dubai, UAE, in 2023 [35]. These recent expansions into the Middle East create their own unique challenges to the retailer because of the region's distinct regulatory environments, which we detail below.

3.2 Compliance with international regulations

Amazon polices the products sold on its platform, and their own shipping restrictions FAQ provides some guidance on why certain products may be restricted, including the need to "comply with all laws and regulations and with Amazon policies" and that Amazon may be "restricted from shipping to your location due to government import/export requirements, manufacturer restrictions, or warranty issues" [2]. Amazon has adapted its policies to allow for the removal of offensive content including content that Amazon determines is "hate speech, promotes the abuse or sexual exploitation of children, contains pornography, glorifies rape or pedophilia, advocates terrorism", but also "other material [they] deem inappropriate or offensive" [4]. However, Amazon has failed to reveal specifically what categories of content it restricts to comply with content regulations in various regions.

In the UAE, imported media including books are governed by the Federal Decree-Law No. 55 of 2023 on Media Regulation [14], which regulates print, television, as well as online media. As examples, it restricts books that disrespect Islamic beliefs or that offend prevailing values in society. Federal Decree-Law No. 31 of 2021 [15]

also criminalizes the import of "indecent materials" and books or spells relating to "trickery, magic or, jugglery".

In Saudi Arabia, imported media are governed by the 2003 Law of Printing and Publication [36] which regulates print media. Article 9 states that printed media cannot contravene Sharia Law, nor lead to a breach of either public security, public policy, foreign interests, stir up internal discord, or injure the economic and health situation of the country. Article 18 states these regulations should apply to the importation and distribution of printed materials. An approval is required, within the framework of Article 18, in order to certify that content is free from anything deemed objectionable. Printed items found violating Articles 9 and 18 can be withdrawn from circulation on the basis of Article 39.

In most regions Amazon is liable for ensuring that, using its platform, only items that can be legally imported can be shipped to them due to laws relating to smuggling or the facilitation of crime. In countries such as the UAE and Saudi Arabia, Amazon proactively restricts shipment of products prohibited by regulations on imported media. However, products may be unavailable for reasons that are opaque to customers, and Amazon may superfluously restrict legal products due to, e.g., applying an insufficiently fine-graded filter that blocks innocuous content [20] or by misunderstanding its legal requirements in a region altogether [27]. Because Amazon does not disclose which items are restricted in which locations or why, we use the methodology described in the following section to identify when shipment restrictions occur and to distinguish censorship from more mundane import restrictions.

4 Experiment overview

Our methodology for measuring shipping restrictions on Amazon consists of three phases. Our primary motivation was to understand how, in Middle Eastern countries, Amazon enforces *censorship*, which we define as the suppression of ideas and information that some individuals, groups, or government officials find objectionable or dangerous [7]. In our first phase, we focus on studying how shipment restrictions vary across multiple countries in the Middle East. To understand how the censorship that we measured in the Middle East applies more broadly to the world at large, in our second phase we pivot from the results of the first phase and measure product availability in regions across the globe. In the same way that Phase 2 expands our measurements over space, in Phase 3 we expand over time. We look at how censorship changed in response to our disclosure to Amazon of issues that we had found.

5 Phase 1: Censorship in the Middle East

5.1 Methodology

One way to try to measure Amazon's implementation of censorship in Middle Eastern countries would be to visit those Amazon storefronts which are available in the Middle East, namely, the UAE's amazon.ae or Saudi Arabia's amazon.sa, and to try to determine which products are anomalously "missing" from being sold on these two Amazon sites. This approach, however, would be limited. For example, if we saw one book related to LGBTIQ topics that was sold on amazon.com but not amazon.ae, that might be due to the book being censored on amazon.ae, but another possibility is that the book was out of stock or not sufficiently popular to be sold in some countries. However, if we saw a disproportionately large number of books related to LGBTIQ topics that were available on amazon.com but not sold on amazon.ae, then we would have a stronger argument, but this argument would be at best a statistical argument, and for any individual product we would not be able to say whether it was the victim of censorship or was unavailable on that storefront for some other reason.

Given the weakness of the previously described approach, we instead measured whether products on amazon.com, the American storefront, could be shipped to various countries. As an additional benefit, this approach allowed us to study censorship in regions that did not have their own dedicated storefront. For our investigation of censorship in the Middle East, we picked four Middle Eastern countries: the UAE, Saudi Arabia, Qatar, and Yemen. We picked these four countries because we were originally motivated by testing whether Amazon applied more censorship in countries in which it operated storefronts (UAE — amazon.ae and Saudi Arabia — amazon.sa) versus two it did not (Qatar and Yemen), even though this was not the ultimate direction of our study. We also tested a fifth country, Canada, as a screen, which we explain later in our methodology.

To test which products we could ship to these five countries, we required a method for sampling a sufficiently diverse set of Amazon products to test. To address this requirement, we made use of the Common Crawl data set provided by the Common Crawl Foundation [12]. This data set is a diverse, open Internet-wide sample of Web pages scraped beginning in 2008. In April 2023, we downloaded all of the archives up to and including the February/March 2023 archive. To avoid excessive storage requirements of storing the entire data set, we downloaded the archive in streaming fashion, filtering out any Amazon product URL into a file without storing any other data from the data set. We processed the Common Crawl data from 2013 through March 2023, as March 2023 was the most recent data set available at the time that we began our testing. Although we were only interested in products available on the amazon.com storefront, since products are often available on multiple storefronts, we collected products from the 23 Amazon dedicated storefronts that were in use at this time.

Using this method, we collected a list of 114,542,719 Amazon URLs. Since not every Amazon URL is a URL to a product, we processed this URL list by searching each URL with the following regular expression:

```
/(?:dp|gp/product|gp/aw/d|gp/switch-language/product
        |product-reviews|asin|offer-listing
       |kindle-dbs/product)/([\^/]*)(?:/|$)
```

This regular expression was designed to search for and detect a variety of ways that Amazon inserts Amazon Standard Identification Numbers (ASINs), Amazon's unique product identifiers, into URLs and extract them from the URL. The result of this processing was a list of 19,074,613 unique ASINs.

To gather information on the availability of products in our five tested countries, we sequentially test ASINs in each region using an automated program to perform the following steps. First, we load amazon.com. Then we switch our location to the region that we are testing using Amazon's "update location" feature on the top

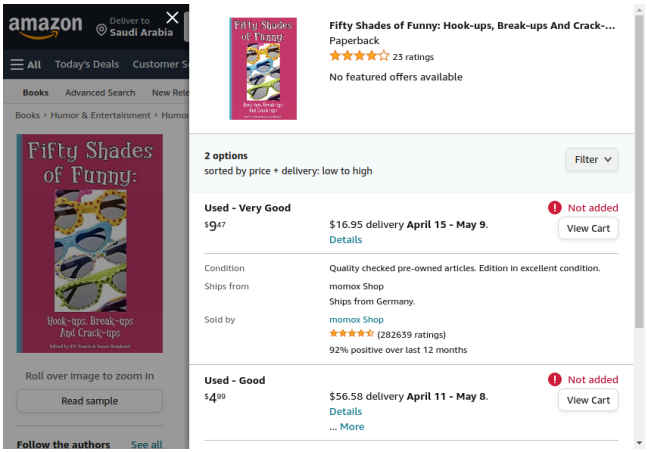

Figure 1: Attempting to add restricted products to the Amazon cart in the all offers display results in "Not added" error messages.

of the page. For each ASIN, we navigate to https://www.amazon.com/dp/**[ASIN]**/ to display that product's detail page. We then parse that page for that product's *availability status*. Note that at no point in our methodology do we sign into any Amazon account. If the product's availability in a region is any of the following, we consider the product *unavailable* in that region:

- This item cannot be shipped to your selected delivery location
- Currently unavailable
- Temporarily out of stock

While product availability on Amazon is influenced by a combination of economic, legal, and prudential processes, our primary motivation is to identify instances of platform-enforced restrictions. We are specifically interested in measuring which products cannot be shipped to a region even if there are shippers who have it in stock and are willing to ship it to that region. We call such products in that region *restricted* products since, even if they were in stock and there were shippers willing to ship them, Amazon would still restrict users from shipping them to that region.

To discern which unavailable products are restricted, we exploit a special side channel, discovered through an ad hoc analysis of amazon.com's purchasing workflow, to reveal if Amazon is preventing an unavailable product from being shipped to us. Namely, we perform the following additional steps via our automated program for any product found to be unavailable. First, we browse to https://www.amazon.com/dp/**[ASIN]**/?aod=1. Note that, compared to the previous URL we had browsed to, this one has appended to it the "?aod=1" query string. Enabling the "aod" parameter signals to Amazon that we want Amazon to render the all offers display (AOD). This advanced display lists all offers from shippers both willing to ship to the user's specified region and who have the product in stock. Next to each shipper's option is a button to add that offer to one's cart (see Figure 1). We programmatically click all of the "Add to cart" buttons on the AOD. We measure the number of buttons whose clicks resulted in the "Added" versus "Not added" messages. If there is at least one offer and all attempts to add offers to our cart result in a "Not added" error, we consider the product *potentially* restricted to our configured location. We schedule another test to run a week after the original, and, if that test has the same result (i.e., that there is at least one offer and all attempts to add offers to our cart result in a "Not added" error), then we consider that product *restricted* to the tested region. If there are no offers (i.e., there are no buttons to click), then we are unable to discern between the product being restricted in the tested region versus being unavailable for benign reasons such as being out of stock (see Table 1 for a summary of possible results). By exploiting how AOD status messages leak whether products are restricted, we are able to learn more about Amazon's system of restricting product shipments to certain regions.

Result from clicking "Add to cart" buttons	Interpretation
At least one "Added"	Product available
All and at least one "Not added"	Product restricted
There are no "Add to cart" buttons to click	Indeterminate

Table 1: Summary of possible results from clicking "Add to cart" buttons and their interpretations.

Note that, although we originally collected the ASINs of Kindle books, due to our methodology's dependence on clicking "Add to cart" buttons, we are unable to use it to test for delivery restrictions relating to digital purchases including Kindle books. As of the time of our study, Kindle books cannot be purchased in this manner.

Since many products on amazon.com do not ship internationally to any region, we perform the following optimization to reduce the number of unavailable products we need to test. When testing unavailable products in each of the four countries for whether they are restricted, we skip testing products that are unavailable in Canada. We use Canada to screen products in this manner because the country has similar values and speech protections compared to the United States [39]. Note though that if Amazon does, e.g., apply censorship to books shipped to Canada beyond whatever restrictions already apply to all books sold on the United States storefront, we will be unable to observe such censorship in our study. Performing this optimization reduced the number of unavailable products that we needed to test by over 85% and increased the amount of products that we could test by a factor of approximately sixfold.

We performed Phase 1 from April 2023 to December 2023.

5.2 Products tested

During our testing period, we were able to test product links collected in the Common Crawl dataset from the February/March 2023 archive working backwards until and partially including the September 2019 archive. Overall we tested 5,870,695 product links during this phase of the experiment. Among these, 2,005,852 (34%) were not (or were no longer) valid product pages, resulting in Amazon "Page not found" errors. Recall that many of the ASINs that we acquired were from dedicated storefronts other than the United States. Therefore, although many of these links may be to products no longer sold, most of these "Page not found" products are

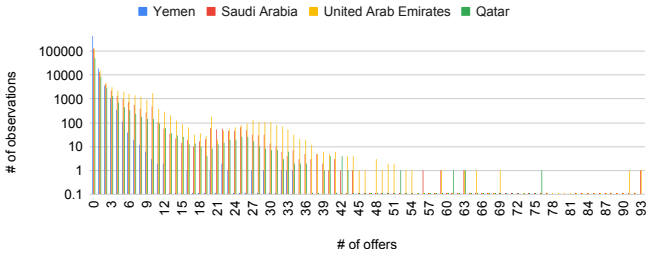

Figure 2: Histogram of the numbers of offers per region.

	# of unavailable products with…		
	zero offers	one offer	> one offer
Yemen	429,571	19,074	5,198
Saudi Arabia	134,737	13,316	11,901
UAE	133,641	14,804	21,943
Qatar	51,455	8,876	6,997
TOTAL	749,404	56,070	46,039

Table 2: Summary of the numbers of offers per region.

likely products that were never available on the US amazon.com storefront, only on the storefronts of other countries. In addition to the aforementioned "Page not found errors", 19,968 product links generated Amazon "Sorry! Something went wrong!" errors. Therefore, among the 5,870,695 product links tested, we tested 3,844,875 actual products.

5.3 Internal consistency of methodology

In our methodology, we aim to measure whether a product cannot be shipped to a region as a property of the product, not the idiosyncratic choices of individual sellers. Although many products had as many as 93 offers (see Figure 2 and Table 2), others had only one. In our methodology, we only consider a product restricted if *all* of its offers result in a "Not added" status. However, as some products only have one offer, we wanted to measure the consistency of results concerning products with multiple offers to gauge the reliability of results concerning products with a single offer. Specifically, we wished to approximate how reliable testing products with a single offer is by measuring how internally consistent testing products with multiple offers is. We do so by looking at the number of products whose offers had statuses which disagreed, namely, where at least one resulted in an "Added" status and at least one resulted in a "Not added" status.

Among our four countries of interest, we observed only 11 conflicting results: three in the UAE, one in Saudi Arabia, and seven in Yemen. We did not observe any noticeable trend in the type of products with conflicting results except that none had any clear motivation for being restricted (see Table 3 for a listing).

Ten out of the 11 products had at least as many "Added" results as "Not added" results. Together with there being no clear rationale for their restriction, the "Not added" results are likely false positives. Since we only consider a product restricted if *all* of its offers result in "Not added" messages, we correctly interpret these cases of mixed

Product	Region	# added	# not added
1944565523	UAE	8	1
B00GEBMPC0	UAE	1	1
0521684188	UAE	4	1
1933662859	Saudi Arabia	33	1
0553152386	Yemen	1	1
B004FTILGC	Yemen	2	1
B00SK73UQQ	Yemen	1	1
1784967335	Yemen	1	1
140819208X	Yemen	2	1
B0052CAZ2O	Yemen	1	2
1905739001	Yemen	2	2

Table 3: Products with conflicting results among offers.

Region	# of known restricted products
UAE	13,604
Saudi Arabia	9,590
Qatar	6,086
Yemen	1,817
TOTAL (unique)	17,842

Table 4: The number of known restricted products in each Phase 1 region studied.

results as being negative cases of restriction. There may be false positives which we were unable to detect, especially for products with only one offer available. However, given the low frequency of these false positives, namely, that among the 46,039 products tested with at least two offers, only 11 showed possible false positives, we can suspect an equally low rate among the products with only one offer. Specifically, if we assume the same false positive rate as what we had measured, then, among the 56,070 products with only one offer, we would expect only between 13 and 14 false positives.

An alternate hypothesis to Amazon implementing these restrictions is that the restrictions result from the idiosyncratic decisions of individual sellers. However, this hypothesis would fail to explain the virtually unanimous consistency in which sellers would have to agree to restrict product shipment to a region. Rather, such consistency in restrictions requires a centrally enforced mechanism.

5.4 Comparison of restricted products

We saw the largest number of restricted products in the UAE, followed by Saudi Arabia, then Qatar. We observed the lowest number of restricted products in Yemen (see Table 4 for details).

We note, though, that this method of comparing the absolute number of known restricted products may be biased. Specifically, if some regions have more products generally available, such regions may appear to be more restricted due to having more restricted products as well. Therefore, a region with a larger number of known restricted products may not necessarily have a higher rate of restriction.

To more fairly compare Amazon's restriction of products across these regions, we chose to generate Venn diagrams in a special

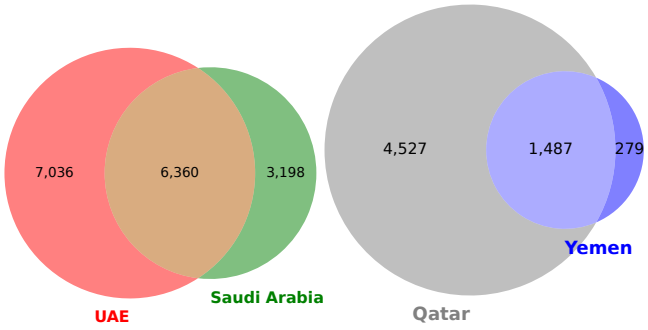

Figure 3: Left, Comparison of overlap of restricted products between the UAE and Saudi Arabia; right, comparison of overlap of restricted products between Qatar and Yemen.

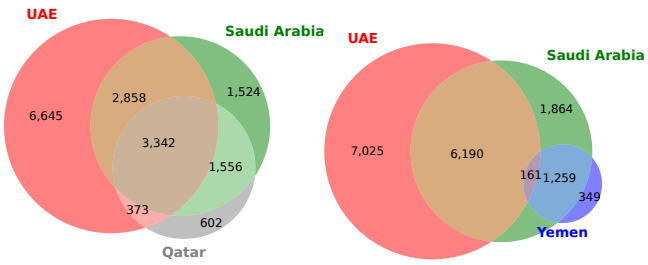

Figure 4: Left, comparison of overlap of restricted products among the UAE, Saudi Arabia, and Qatar; right, comparison of overlap of restricted products among the UAE, Saudi Arabia, and Yemen.

fashion. Namely, we only consider products for which there was at least one offer available in every region. In other words, we only consider products for which we have a clear yes or no result concerning whether it was restricted in every region featured in the diagram. We do this because we do not want our results concerning the number of restricted products to be biased toward regions in which we have more known results. Since each diagram features different regions, the totals therefore may not be consistent across diagrams due to this method of comparison.

Between the UAE and Saudi Arabia, we found that the UAE restricted the largest number of products with around half of the products restricted in UAE also being restricted in Saudi Arabia (see Figure 3). We found that 6,360 products were restricted in common by the UAE and Saudi Arabia.

Comparing Qatar to the UAE and Saudi Arabia, we found that fewer products were restricted in Qatar and that almost all products restricted in Qatar were also restricted in Saudi Arabia and that most of the products restricted in Qatar were also restricted in the UAE (see Figure 4). Comparing Yemen to the UAE and Saudi Arabia, we found that fewer products were restricted in Yemen and that almost all products restricted in Yemen were restricted in Saudi Arabia (see Figure 4).

Finally, comparing Qatar to Yemen, we found that Yemen had the fewest number of restricted products and that almost all products restricted in Yemen were also restricted in Qatar (see Figure 3).

Amazon Error Message	Rate
Currently unavailable	74%
Temporarily out of stock	23%
Cannot ship to location	3%

Table 5: Percentage of messages returned by Amazon for restricted items in the UAE and Saudi Arabia.

Notably, we found different levels of restriction among these four regions despite their cultural and geographic proximity in the Middle East. Namely, the UAE featured the highest level of restricted products, followed by Saudi Arabia, then Qatar, with Yemen having the fewest. Since we found that the UAE and Saudi Arabia feature the most products and because almost all products restricted in Qatar and Yemen were also restricted in either the UAE or Saudi Arabia, we focus the remainder of our analysis on the UAE and Saudi Arabia.

5.5 Analysis of availability messaging

In our dataset, products that are restricted by Amazon presented different and inconsistent messages to the user. The messaging around the rationale for why a certain product is restricted could inform the user as to the reason why items are unavailable. However, among all of the restricted products identified by this study, none presented a message to the user explaining that the items were unavailable due to regulatory reasons. Instead, each item was communicated as either being "currently unavailable", "temporarily out of stock", or that "this item cannot be shipped to your selected delivery location" (see Table 5).

Most of the messaging presented to users for these restricted items is vague and unclear. While "this item cannot be shipped to your selected delivery location" was consistent with the notion that Amazon is restricting shipments based on location, we only observed this message 3% of the time. In the other messages, Amazon used misleading terminology suggesting a transitory condition such as the items being temporarily out of stock or currently unavailable.

5.6 Analysis of restrictions in Saudi Arabia and the UAE by Amazon product category

Figure 5 shows the overall products by categories that are restricted in Saudi Arabia and the UAE. This is based on Amazon's own categorization of a product, which can include both specific categories (e.g., "Book -> Science Fiction") as well as more general categories (e.g., "Book -> Genre Fiction"). The restricted products are dominated largely by book-related categories, mainly "Genre Fiction" and "New Age & Spirituality", which are the top one and two restricted categories respectively in both countries. Japanese Manga and Fantasy book categories are also present in the top 15 restricted product categories in both countries. The book-related categories that are most restricted in Saudi Arabia include "Thrillers & Suspense", "Erotica", "Christian Living", "Anthologies", and "Action & Adventure". Non-book-related product categories restricted in Saudi Arabia include "Cell Phone Cases", "Screen Protectors", and "Groceries". In the UAE the book categories most restricted include "Thrillers and Suspense", "French", "Occult & Paranormal",

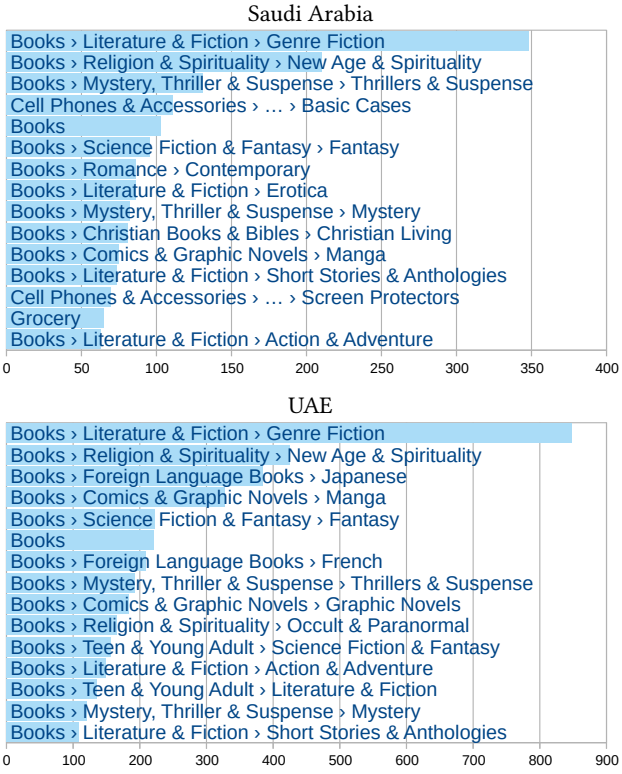

Figure 5: Top 15 restricted product categories in Saudi Arabia (top) and the UAE (bottom). Note that the top categories in the UAE are entirely book-related categories. Saudi Arabia's top categories are all but three book-related categories.

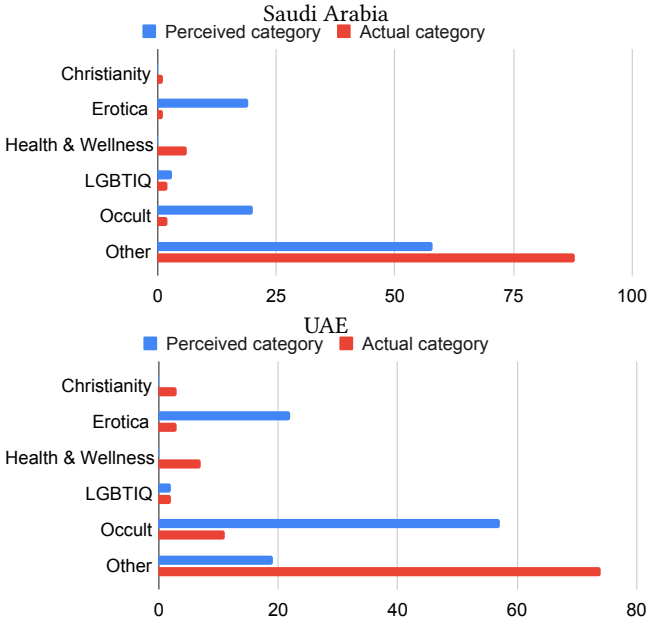

Figure 6: Perceived reason for restriction versus actual category of product in Saudi Arabia (top) and the UAE (bottom) among 100 randomly chosen, restricted products.

"Mystery", and "Short Stories". There are no non-book categories represented in the top 15 restricted categories in the UAE.

Although Phase 1 of our experiment began in April 2023, ending in December 2023, it was only on June 6, 2023, that we began capturing the Amazon category for products that we did not find restricted in any region. Considering only the books that we tested since that time, we find that Amazon restricted shipment of 8,965 out of the 796,081 (1.1%) books in that sample to at least one of Saudi Arabia, the UAE, Qatar, or Yemen.

5.7 Analysis of restrictions in Saudi Arabia and the UAE by motivation

In the previous section, our analysis was limited to understanding the motivation behind Amazon's shipping restrictions by looking at the categories Amazon assigns each restricted product. While this kind of analysis tells us which products are restricted, it does a poor job of describing the motivations for why they are restricted. To address this gap, we conduct a qualitative analysis, employing a more nuanced approach to decipher the underlying reasons for these restrictions.

We selected a random sample of 200 products, 100 items restricted from shipment to the UAE and another 100 items restricted from shipment to Saudi Arabia, to understand the breadth of Amazon's shipping restrictions across various product categories. This random selection sought to minimize the bias and ensure representativeness.

We analyzed items from this random sample based on their titles and descriptions. Following grounded theory, we first went through all of the restricted products to discern broad categories and repeated themes. We then reviewed all of the restricted products again and applied an appropriate category to each product. In some cases when the product information was not listed in the description, we conducted further background research to understand the product. We categorized them based on the *actual* nature of the products and the *perceived* category of the products. The latter refers to a category we inferred based on potentially sensitive keywords within the products' descriptions or titles that may be triggering Amazon's algorithms to perceive it as being under a category. As one example, Nietzsche's *Gay Science* contains the word "gay" in the title, suggesting that it was censored for containing the word "gay", even though the book is a philosophical work that does not speak to LGBTIQ topics. (We found no other works by Nietzsche or other works concerning nihilism restricted.) This dual categorization was designed to uncover discrepancies between a product's apparent content and its perception by Amazon's keyword-sensitive algorithms (see Figure 6).

Both categorizations consisted of the following categories: "LGBTIQ", "Occult", "Erotica", "Christianity", "Health & Wellness", and "Other". The actual category of a product that we assign is based on our analysis of the items' titles and descriptions, whereas the perceived category is also based on the identification of certain keywords believed to trigger Amazon's algorithm and influence

the miscategorization of some items. We attribute to censorship the subset of restrictions that (1) fall into a small number of recurring themes known to be censored (e.g., LGBTIQ, occult, sexuality, religion) or (2) are associated with these themes via keywords in titles or descriptions, including keyword-induced false positives.

Our review of items restricted in the UAE and Saudi Arabia highlights Amazon's possible collateral censorship of items, which imposes unnecessary censorship onto users. In the following sections we highlight select items among each of our categories, including suspected cases of collateral censorship.

5.7.1 LGBTIQ. Examples of LGBTIQ content from our random sample of restricted products include a movie featuring gay characters — and containing the word "gay" in its title — and a book on the history of the persecution of homosexuals in Nazi Germany.

Other types of LGBTIQ content from outside of this sample that we observed include workbooks relating to sexual orientation and gender identity, trans and genderqueer erotica, books on LGBTIQ criminalization in the US and queer activism, and cookbooks based on queer community-led culinary practices.

However, many products containing the word "rainbow" in their descriptions were censored despite not being otherwise related to LGBTIQ themes. Such products included rainbow-colored hair extensions, a travel case, a video game, a movie DVD, a detective novel called *Mister Rainbow*, and Mentos rainbow candy. These findings reveal that Amazon restricts rainbow-related products that have no explicitly stated affiliation with LGBTIQ topics.

As we noted previously, many books were also collaterally censored for containing the word "gay" in their title or description. As a chief example, Nietzsche's *Gay Science* was censored, despite being a philosophical work unrelated to LGBTIQ topics. (None of Nietzsche's other books were restricted, nor were any other books related to nihilism.)

5.7.2 Occult. In our random sample, we found books related to the occult and the paranormal including those on tarot, fairy tales, demons, jinn, witchcraft, astrology, crystals, freemasonry, astral projection, and Bigfoot. Much of the children's books' censorship seemed motivated by censoring the occult, although it is unclear whether these children's books were expressly targeted or collateral damage of some larger censorship strategy. For example, in our random sample are children's books related to jinn, witches, wizards, and necromancy. Our random sample also featured one book describing how to write fictional books relating to monsters as well as a book about Harry Potter.

Although not represented in our smaller random sample, in our larger data set we observed books relating to Thelema, an occultist movement, as well as its founder, Aleister Crowley. Many books related to extraterrestrial aliens and ufology at large were also restricted. We also observed a Dell Alienware laptop that was restricted, although we are unsure if this is due to the product's allusion to aliens or due to electronics or communications regulations. Notably, while we observed multiple restricted products relating to oracles and divination both inside and outside of our random sample, we also found a large number of books outside of our sample related to the Oracle database software that seem to have also been caught up by the filter, suggesting that books related to oracles are over-censored.

5.7.3 Erotica. We identified a significant number of restricted books in our random sample whose censorship was likely motivated by restricting erotica, even though they did not contain erotica, including literature, humor, photography, travel, self-help, classic literature, among others. A notable example is *Never Mind the Botox*, a book that features the lives of four women working in cosmetic surgery and their daily struggles. Whereas cosmetic surgery is legal in both the UAE and Saudi Arabia, we think that certain keywords in the book description have resulted in its miscategorization under "erotica".

Another example is *Sex Addiction Survival Guide: A Practical Workbook for Reconnecting to Yourself and Others*, which is a guide targeting individuals who are struggling with sex and pornography addiction issues, motivating them to move toward a healthier connection to self and others. Censored, it was most likely flagged as an "erotic" book based on certain words in the title and description, such as (sex, porn, sexual, hypersexuality, etc.).

Finally, *Klarissa Dreams Redux: An Illuminated Anthology* is a collection of poetry and other writings in the context of Klarissa Kocsis's paintings. In the description of the book, however, the author is described as a "breast cancer survivor" and as having "earned a reputation for portraiture and nudes", either of which may be triggering Amazon's censorship enforcement. Literature relating to breast cancer is commonly collaterally censored on various platforms.

5.7.4 Christianity. We identified numerous censored books related to Christianity in our random sample. While Christian proselytization is illegal in many Middle Eastern countries, we found that the restricted works related to Christianity often deal with topics such as demons or the devil. Thus, we believe that these books were not caught up by rules intending to target Christian topics per se but rather demons or other sensitive spiritual topics that are also related to magic or the occult.

Many censored books relating to Christianity explicitly mention the devil, such as *The Soul of The Apostolate*. The book advocates for Christian apostolic work, highlighting that it hinges not merely on activity but fundamentally on a robust interior life. However, it also advertises to reveal the "Devil's special temptations for those working for Our Lord". Another example is the censorship of *Get Thee Behind Me, Satan: Rejecting Evil* which raises questions about biblical facts and their relation to devils, although the work is ultimately concerned with motivating the need for vigilance against the pervasive appeal of evil.

There was also one Christian book likely censored due to erotic themes. *Moral Ambiguity* is the fictional story of Kevin Gregory, a celebrated singer, who exposes the corrupt practices and moral hypocrisy of a powerful televangelist. While the book is primarily concerned with the main protagonist revealing the hypocrisy of religious authority, the book's description also alludes to the antagonist being motivated by greed and sex.

5.7.5 Health and wellness. Restricted products include condoms of several brands, sex toys such as vibrators, and sex education, sexual health, and gender identity books and textbooks. Lubricants are also restricted, although automotive greases with the word "lubricant" in their descriptions were also among the restricted products in our random sample.

5.7.6 Other. A large number of products in the "Other" category include those which are heavily regulated, such as products related to WiFi or car seats. This is especially the case in Saudi Arabia, where we also saw a large number of restricted products related to mobile phones. However, the large discrepancy between the number of products in the "Other" *perceived* category and the "Other" *actual* category points to the large amount of overly broad filtering which perceives certain products as being related to restricted categories but in actuality they are unrelated.

For a variety of products that we measured we could not identify any reason for their restriction. As one example, we found that Malala Yousafzai's *She Persisted* was restricted in the UAE and Saudi Arabia. While we might speculate that such a book could be politically sensitive in those countries, it does not fit into any of the above categories of products. Moreover, other Malala Yousafzai titles were available in these countries, so it is unclear if this is intended censorship or if it is collateral damage from some overly broad filtering rule. There may also exist other motivations for product restriction other than those which we could identify.

5.7.7 Other languages. Not all products within our 200 product random sample were in the English language. Thirteen restricted products from non-English language products were also included. These are products in Japanese (six products), French (five products), and German (two products). All restricted non-English media was for written material like books, including many comics, except for one German movie, *Könige der Welt*, a documentary about addiction and success in the music business. Many of the blocked products within other languages can be categorized in themes identified previously such as the occult, erotica or sexual health information.

Many of the restricted Japanese products are Japanese comic books or "manga". Many of these restricted products contain occult themes or contain violent content. Some restricted manga is marketed to a younger audience such as *My Hero Academia* and *Dragon Quest*, which we suspect is restricted due to being caught up in an overly broad filter targeting occult content. Another suspected miscategorization is a book by Japanese author Mayumi Tanimoto called "キャリアポルノは人生の無駄だ" (Career Porn is a Waste of Life). This book is both a humorous and earnest criticism of the work environment and labor conditions in Japan. We suspect that the characterization of this issue in the context of "career porn" has led to an erroneous miscategorization of the product as being pornographic. This is possibly due to the inclusion of the characters "ポルノ" ("porno") in either the title or description.

Restricted French language content includes sexual health information such as a book about the Kama Sutra, a guidebook for performing sexual acts, and a humorous guide to dating. We suspect that two French books were captured by censorship rules targeting the occult: a young adult fantasy novel about the fictitious black magic book the "Necronomicon", and a book about Tarot card interpretation. The one restricted German book is a children's fantasy book about a girl named Willow, who loves nature, according to the title. We suspect that this book was also captured by occult-targeting censorship.

5.7.8 Summary. The resulting restrictions we measured are consistent with Amazon implementing censorship rules targeting content relating to LGBTIQ, erotica, the occult, sexuality, and other religiously adjacent content. Conversely, the presence of an "Other" category containing regulatory and hard-to-explain items suggests that censorship on Amazon is implemented using the same restriction system as conventional import controls implementing mundane regulatory enforcement.

6 Phase 2: Expanding measurements over space

6.1 Methodology

In Phase 1, we measured which products are restricted to the UAE, Saudi Arabia, Qatar, and Yemen. However, it remained unclear to what extent our measurements in Phase 1 fit within Amazon's greater system. Specifically, we were motivated by the following two research questions:

- To what extent are other world regions affected by the restrictions we measured in Phase 1's countries?
- Are the shipping restrictions on books that we measured in Phase 1 performed by sellers individually, by publishers enforcing licensing contracts, or are these restrictions enforced systematically by Amazon?

To answer these questions, in Phase 2, we expanded our measurements to the 239 countries and other regions to which Amazon supports shipment. To do this, we remeasured restricted products from Phase 1 across every other world region.

Specifically, we took the set of products that are restricted in at least one of the four Middle Eastern countries we analyzed. From this set, we created our Phase 2 test list by choosing 1,000 of these products uniformly at random. For each product in our Phase 2 test list, we perform a similar test as we did in Phase 1, except instead of only testing in four Middle Eastern countries we test whether that product is unavailable in and restricted in all 239 regions to which Amazon supports shipping. However, by doing this, we will measure how Amazon restrictions are applied across the globe *insofar as it overlaps with any product restrictions that we had previously measured in the Middle East.* The advantage of taking this approach is that, by testing only a limited number of products in each region, we will be able to scale testing to every Amazon region in the globe. However, our method *will fail to detect forms of censorship or other restrictions not present in the Middle East*, e.g., censorship of Uyghur topics in China. While Phase 2 will prove useful for answering our earlier research questions, *it cannot be interpreted as an absolute measure of the amount of censorship or other regulatory restrictions in a region.*

We performed Phase 2 from May 2024 to June 2024. The following sections detail our results.

6.2 Clustering product availability across regions

To compare how product availability varied across all 239 regions, we clustered each region according to the inter-similarity of the products that are restricted in each region. Specifically, we compare each product's results using their Hamming distance $h(a, b)$, where $h(a, b)$ is 0 if $a = b$, 1 otherwise.

To compare two vectors of results, either across every region keeping the product fixed, or across every product keeping the

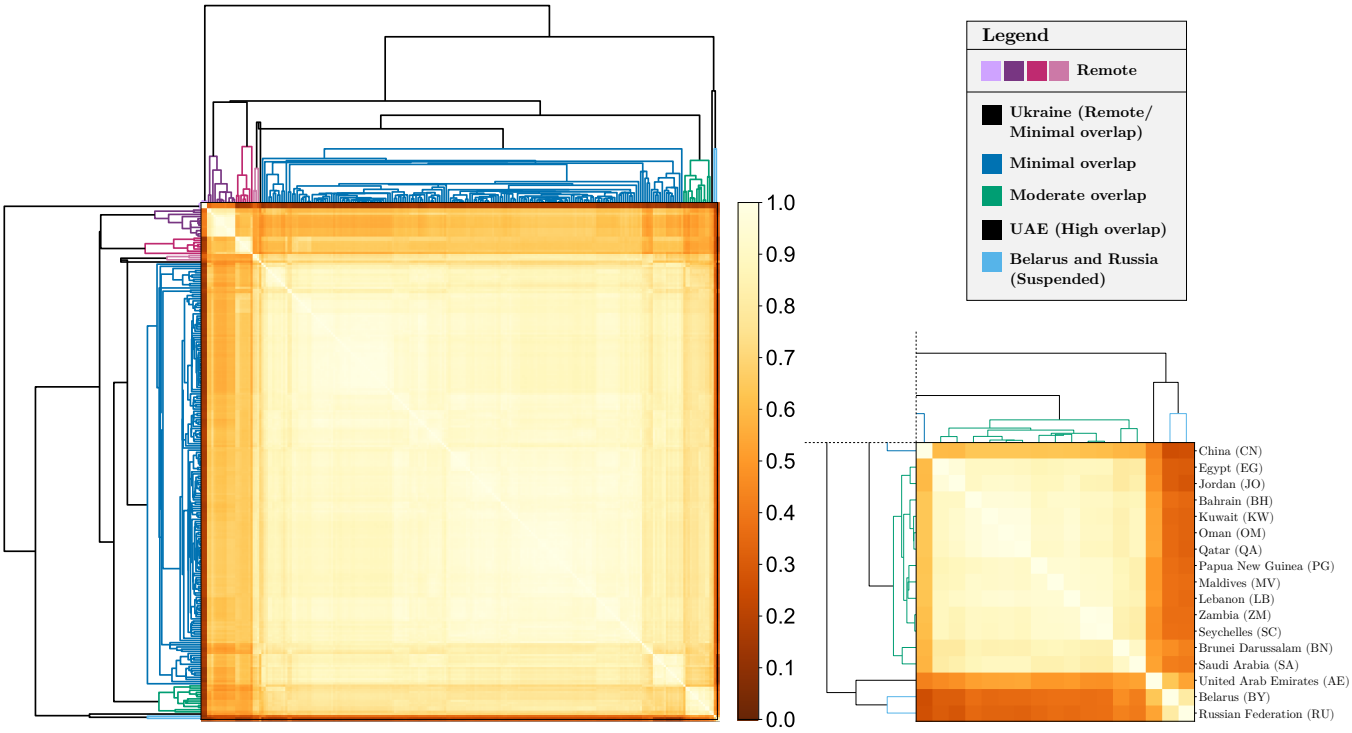

Figure 7: Jaccard similarity matrix of each region's restricted products, hierarchically clustered. Dendrograms help identify seven clusters: the mauve, purple, magenta, pink, blue, green, and cyan clusters as well as the final black singleton cluster. The height of the tree reflects the distance between the two branches joined using the farthest point linkage method. Therefore, the height of the tree is short inside of each cluster and only becomes tall when joining clusters. On the bottom right, the lower right corner of the similarity matrix, zoomed in. See here [25] for the full data set.

region fixed, we use the distance metric

$$D(x, y) = \frac{1}{n} \sum_{i=1}^{n} h(x_i, y_i), \text{ where } n = |x| = |y|.$$

Using this metric, we hierarchically clustered each region. The resulting clustered similarity matrix and dendrogram is in Figures 7. In Figure 7 both axes are regions and each cell represents two regions' restriction similarity. In Figure 8 the Y-axis varies over products, and we hierarchically cluster the rows of different products in the same manner as we did the columns of different regions. Each cell in Figure 8 represents whether a product is artificially restricted in a region.

In Figure 8, we identified the four leftmost clusters as having limited shipping options due to varying degrees of physical or logistical remoteness. At the extreme, some of these locations are remote, unpopulated islands (e.g., Bouvet Island), which would pose obvious shipping challenges. Other locations are not physically remote but are logistically difficult to ship to due to ongoing political instability or military conflicts. As such we **refer to these four** as "remote" clusters.

Moving rightward is a singleton cluster consisting of **Ukraine** followed by a large 195-member cluster with regions we consider to have **minimal overlap** with the Middle Eastern restrictions we measured in Phase 1. Inside of the **minimal overlap** cluster, there is some variation such as which types of groceries or health supplements can be delivered or whether products from Amazon Global Store UK can be delivered. Generally, however, in this cluster, although some products were restricted, we did not find any that could be categorized as religious or political censorship. **Ukraine** being clustered in between the collection of four "remote" clusters and the **minimal overlap** cluster suggests that the 2022 full-scale invasion of Ukraine may have limited courier's access to the region, although not to the same extent as those in the "remote" clusters.

Zooming in on the 17 members in the lower right of Figure 7, beginning from the upper left, we first see China, the final member of the **minimal overlap** cluster. Despite Amazon restricting products in this region, we did not generally observe evidence of censorship to this region in our Phase 2 experiment. We believe that this is because Chinese censorship is derived from largely different motivations than Middle Eastern censorship, and so such censorship would have little overlap with the censorship that we measured in Phase 1. However, we do note that while it was in the **minimal overlap** cluster, it was the region in this cluster that was clustered the closest to the more restricted regions to its right.

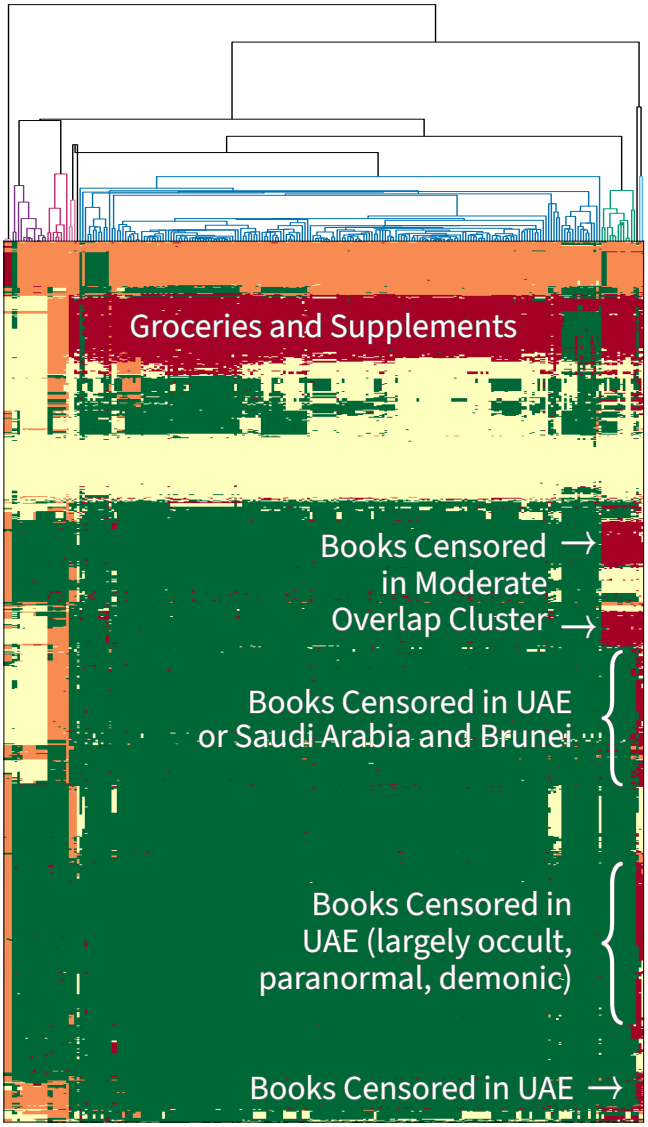

Figure 8: Same *X*-axis and dendrogram legend as Figure 7 but *Y*-axis is all 1,000 tested products where each matrix cell indicates the result of the test; cell color green: "Added", yellow: no offers, red: "Not added". On the far right, Russia and Belarus generally cannot receive shipments. To their immediate left, UAE was the next most restricted. Left of that, the moderate overlap cluster shares, visualized in red blocks, a common set of censored products, largely books. See here [25] for the full data set.

Moving rightward, we see Jordan, Egypt, Bahrain, Oman, Kuwait, Qatar, Lebanon, Papua New Guinea, Maldives, Zambia, and Seychelles forming a cluster. Brunei Darussalam and Saudi Arabia are also in this cluster. Due to this cluster's censorship of LGBTIQ, occult, and other topics that do not reach the UAE's, we refer to this as the **moderate overlap** cluster. However, while the dendrogram tree is shallow between Brunei Darussalam and Saudi Arabia, suggesting their censorship and other product restrictions are highly similar to each other's, the dendrogram link between this pair of countries and the remainder of the "moderate overlap" cluster is taller, suggesting that this pair of countries has less in common with the remainder of the cluster than they do with each other. We found that various phone accessories were restricted in Brunei Darussalam and Saudi Arabia. We are unclear if this is the result of some regulation uniquely affecting these regions or if this is the product of some kind of restriction which we do not presently understand.

Continuing rightward, we have the UAE in a singleton cluster. In Phase 1 we identified it as being the most restricted region among those analyzed, and our Phase 2 results are similar. As such we identify it as the **high overlap** cluster.

Finally, we see Belarus and the Russian Federation forming a cluster. Following the 2022 full-scale invasion of Ukraine, Amazon announced suspension of shipment to these countries [1]. As such we refer to this as the **suspended** cluster.

Our clustering technique was capable of revealing clusters of regions with similar restrictions as well as for uncovering those restrictions themselves. These results also shed light on our previous Phase 1 findings which found that the UAE censored the most, followed by Saudi Arabia and Qatar, with Yemen censoring the least. We can now identify the UAE as belonging to the **high overlap** cluster, Saudi Arabia and Qatar as belonging to the **moderate overlap** cluster (with Saudi Arabia having additional restrictions affecting, e.g., phone accessories), and Yemen as belonging to the **minimal overlap** cluster.

6.3 Censorship regimes

While our analysis fleshed out various reasons for restrictions on shipping, including the remoteness of unpopulated islands or Amazon's commitment to suspend shipments to Russia and Belarus, we note that when we restrict our concern to only restrictions motivated by censorship, we observe three clusters: the **minimal overlap** cluster, the **moderate overlap** cluster, and the **high overlap** cluster.

We believe that these clusters are explained by the application of three different censorship *regimes* which Amazon uses to simplify enforcement of censorship. By binning regions into different censorship regimes, each with a *masterlist* of censorship approximating what products to censor in each regime, Amazon can perform censorship more expeditiously versus applying censorship specifically tailored to each region. This "pigeonholing" approach may help explain the over-broadness of some censorship that we observed.

6.4 Incompletely applied Russian shipping suspensions

Despite Amazon's announcement of suspension of shipments to Russia and Belarus [1], we found multiple products that could be shipped to these regions, including a mystery novel, a Nativity toy play set, and a tube of seafood-flavored cat toothpaste. We could identify no other commonality among these and other products which we could ship to Russia and Belarus, and such products were shippable from a variety of sellers located in diverse regions. However, they point to the incomplete application of Amazon's

suspension of shipments to these regions and signal that Amazon's system of restricting shipments may, in rare occasions, fail.

6.5 Explanatory power of cluster model

At a glance, Figure 8 appears "block-like". We measured how well a small number of region clusters explains the availability matrix by comparing a one-block model (all regions share a single average availability pattern) to a clusters model in which each of nine region clusters has its own template pattern. Using the reduction in summed squared error, the clusters model explains 83.9% of the variance in entries of the availability matrix ($R^2 = 0.839$), relative to the single-block baseline. In other words, most of the variation in whether products are available, or whether they are artificially restricted, can be captured by assigning regions to a small number of clusters with shared availability templates, rather than treating each region as behaving idiosyncratically. Previous work studying censorship systems shows that, absent centralized censorship, even when acting under similar motivations, censors typically censor widely diverging content [23, 24, 29].

Given these results, we reach the following conclusions. First, **the censorship we measured during Phase 1 is applied to many other regions worldwide via a small number of censorship regimes**. While the UAE is unique in being subject to an exceptionally high level of censorship, the same censorship regime that applies to Saudi Arabia and Qatar **includes 11 other world regions spanning three continents**. Second, **the restrictions on book shipments we are measuring are applied by Amazon**, not individually by millions of sellers nor by thousands of independent publishers. In addition to the blockiness of our matrix, this conclusion is also supported by our observation that censorship of a product is predicted by the presence of keywords, even if those keywords generate false positives (e.g., Nietzsche's *Gay Science*). Such false positives are predicted by the use of automated censorship systems based on keywords, such as those that Amazon might employ, not by sellers or publishers making idiosyncratic decisions.

7 Phase 3: Expanding measurements over time

7.1 Methodology

In Phase 3 we measured how Amazon's restriction of product shipments has changed over time, specifically in response to an October 23, 2024, disclosure of our findings to Amazon. One year after our initial experiment, we retested all of the products that we found to be artificially restricted in the UAE and in Saudi Arabia. We also retested each product in Canada as a control.

Since our disclosure, Amazon removed the AOD side channel that we used for determining if a product was being artificially restricted, as Amazon appears to now remove offers which could not be added to users' carts from being listed in the AOD in the first place. Therefore, to restrict our analysis to products which are more likely to be artificially restricted, we only consider products in this phase that are both (1) unavailable in the UAE or Saudi Arabia and (2) available in Canada.

Among such products, for each of UAE and Saudi Arabia, and for each of the "This item cannot be shipped", "Temporarily out of stock", and "Currently unavailable" messages, we randomly sampled 100 products and manually annotated them. These annotations included the reason why the product may have been identified to be restricted and whether those reasons match the actual product category. As one example of a false positive where the reasons fail to match, we found that Nietzsche's *Gay Science* was restricted for containing the word "gay" in its title despite it being unrelated to LGBTIQ issues.

We performed Phase 3 from July to August 2025. The following sections detail our results.

7.2 Changes in availability messaging

In our samples, the "Temporarily out of stock" message consistently allowed backordering the item with the natural caveat that the item will not ship until it is back in stock. Thus, such products do not appear to be actually restricted, and we find that the product description "Temporarily out of stock", which was the most common message used in Phase 1 to describe restricted products, is now used only to indicate product stock and not whether it is restricted. This is a notable improvement in transparency, given that this message's language is the furthest from representing that a product's shipment is being artificially restricted. However, we found that "Currently unavailable" and — "This item cannot be shipped to your selected delivery location. Please choose a different delivery location." — products now consistently show no offers in the all offers display. Thus, further analysis is required to determine if these messages are accurate, with "This item cannot be shipped" being the message most accurately reflecting a product being artificially restricted.

Of the 17,842 products which we retested, there were 9,938 that were both (1) available in Canada but (2) unavailable in either the UAE or Saudi Arabia. The vast majority of these reported the message "This item cannot be shipped". Figure 9 shows the overlap of products that showed each of the availability messages in the Middle East and products that were unavailable in Canada. The fact that the majority of the "This item cannot be shipped" products were available in Canada while almost all of the other unavailable items were consistent with our control set provides evidence that, in addition to the "Temporarily out of stock" message, the "Currently unavailable" message is also no longer being used for products that are being artificially restricted.

While there has been an improvement in the consistency of these messages, they remain vague. During our annotations, we found that Amazon uses the "This item cannot be shipped" message for products that are restricted for mundane regulatory reasons as well as for products that are censored for their ideas. Further transparency in these messages would help inform users why a product may not be available in their region.

7.3 Changes in false positives

We also measured whether there were changes to Amazon's false positive rate. We found that **47%** of the "Temporarily out of stock messages" products and **11%** of the "This item cannot be shipped" products in our sample were categorized as false positives. Since our analysis suggests that "This item cannot be shipped" products are consistently restricted and that "Temporarily out of stock messages" products can now always be backordered and thus are consistently never restricted, the much lower rate of false positives among the

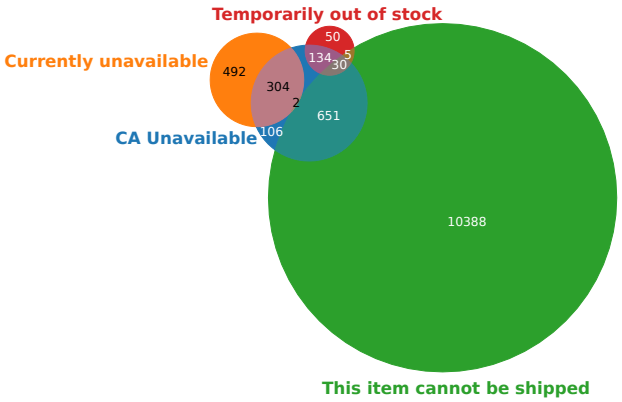

Figure 9: Green, orange, red: products reporting "This item cannot be shipped", "Currently unavailable", and "Temporarily out of stock", respectively, in either the UAE or Saudi Arabia; blue: products that are unavailable in Canada for any reason.

former suggests that currently restricted products have a lower rate of false positives than those which had been restricted during Phase 1 but can now be ordered. Nevertheless, Amazon's censorship mechanism still inappropriately flags products with certain keywords that appear in their titles or descriptions.

8 Limitations

In this section we enumerate and evaluate various limitations of the methodologies we employed in this study.

During Phase 1 of our experiment, we derived our list of products to test by accumulating product URLs from the Common Crawl dataset and by scraping those URLs with a crawler which we implemented. This sampling of products does not take into account sales volume or product popularity, so it is possible that our dataset of products includes obscure or never-sold items. Further analysis would be required to identify the full extent to which these restrictions impact overall product sales. However, this is outside of the scope of our study which seeks to characterize the availability of information on Amazon and identify what content triggers censorship in Amazon's automated system.

During Phase 2 of our experiment, our test set was derived from items that we found restricted in the UAE and Saudi Arabia. As regulatory environments vary between countries, other countries have categories of restricted products not captured by this test set, e.g., censorship of Uyghur topics in China. Therefore, our measurements in Phase 2 must be interpreted as overlap with Middle Eastern censorship and not an absolute measurement of censorship.

In our study we encountered some limitations relating to language diversity. Although our study ultimately included books in multiple languages, such as English, Japanese, and French, it notably did not encompass any Arabic books, which are critically relevant in both the UAE and Saudi Arabia. The lack of representation of such books may be due to a minuscule representation on the American Amazon site — amazon.com — or due to a sampling bias of Common Crawl.

9 Ethical considerations

We uphold the belief that we do not need the consent of a censorship system's operator to measure its censorship. Moreover, both Amazon's "Conditions of Use" and its `robots.txt` file welcome bots to crawl and interact with its product pages [3, 5]. However, we still carefully designed our measurements to not harm other users of Amazon's platform. First, we did not use proxies or any other method of hiding the origin of our measurements. Second, we downloaded pages from Amazon within the rate limits that Amazon allowed, and we did not attempt to circumvent these limits. Third, we did not create or use any Amazon accounts, nor did we conduct any actual purchases over the platform. Finally, all data we collected was available publicly without being logged in. To further minimize our risk of storing sensitive data even if it is publicly available, we processed each downloaded product page's HTML before storing its relevant contents. Specifically, we only stored product titles, categories, and descriptions; we did not store any other strings from the downloaded HTML such as users' reviews.

10 Conclusion

We designed and executed a novel methodology to distinguish cases of Amazon's digital enforcement of product shipment restrictions from benign explanations of product unavailability such as an item being out of stock. Our results reveal that Amazon uses centralized and automated rules to enforce regulations including government-mandated censorship.

While corporate compliance with government censorship demands is not unexpected, we make note of the following problematic behaviors exceeding Amazon's legal obligations: (1) In response to attempting to ship censored products, Amazon used opaque error messages suggesting transitory conditions such as the items being temporarily out of stock; (2) Amazon's use of broad, keyword-sensitive censorship rules resulted in collateral censorship beyond what laws require; (3) Amazon often uses centralized rules to censor books across entire regimes of countries, which can span multiple continents, instead of implementing bespoke, case-by-case enforcement.

On points (1) and (2) we found improvements a year after reporting the issues to Amazon, but Amazon still collaterally censored products. On point (3), in future work we invite legal scholars and other domain experts to scrutinize whether such censorship exceeds that required in affected regions.

Finally, we note that the UAE and Saudi Arabia, two countries in which we measured the most censorship, are two countries in which Amazon operates storefronts. We hypothesize that and leave to future work to investigate whether Amazon's increased legal exposure in these regions has motivated broader compliance, especially in the UAE in which Amazon implemented bespoke censorship rules among the products and regions we analyzed.

Books have a long history of being subject to censorship, and digital tools equip governments and corporations with new means by which to curtail the spread of sensitive ideas. Just as it is important for researchers to study how wholly digital platforms comply with authoritarian information controls, we invite further work to be done scrutinizing the use of digital technologies to censor and control analog media. Our data is publicly available here [25].

Acknowledgments

We would like to thank Jedidiah Crandall, Irene Poetranto, Adam Senft, and our paper's anonymous reviewers for valuable feedback. We are also grateful to Rasha Younes whose conversation with us concerning Amazon's search-based censorship on amazon.ae and amazon.sa inspired us to perform this research. Funding for this project was provided by the Citizen Lab at the University of Toronto's Munk School of Global Affairs & Public Policy.

References

[1] Amazon. 2022. Updates to Amazon's retail, entertainment, and AWS businesses in Russia and Belarus. https://www.aboutamazon.com/news/aws/updates-to-amazons-retail-entertainment-and-aws-businesses-in-russia-and-belarus

[2] Amazon. 2025. About Shipping Restrictions. https://www.amazon.ca/gp/help/customer/display.html?nodeId=G8UUV583T397XQ3R

[3] Amazon. 2025. Conditions of Use. www.amazon.com/gp/help/customer/display.html/?nodeId=GLSBYFE9MGKKQXXM Relevant section is under the heading Agents.

[4] Amazon. 2025. Content Guidelines for Books. https://www.amazon.ca/gp/help/customer/display.html?nodeId=GJUAA7A28KBJE2YV

[5] Amazon. 2025. robots.txt. www.amazon.com/robots.txt

[6] Amazon. 2025. Website (Country/Region). https://www.amazon.com/gp/navigation-country/select-country At the time of submission there are 24 regions available to choose from.

[7] American Library Association. 2023. Intellectual Freedom and Censorship Q & A. https://www.ala.org/advocacy/intfreedom/censorship/faq

[8] Collin Anderson. 2013. Dimming the Internet: Detecting Throttling as a Mechanism of Censorship in Iran. arXiv:1306.4361 [cs.NI] https://arxiv.org/abs/1306.4361

[9] Simurgh Aryan, Homa Aryan, and J. Alex Halderman. 2013. Internet Censorship in Iran: A First Look. In *3rd USENIX Workshop on Free and Open Communications on the Internet (FOCI 13)*. USENIX Association, Washington, D.C., 8 pages. https://www.usenix.org/conference/foci13/workshop-program/presentation/aryan

[10] Sam Burnett and Nick Feamster. 2015. Encore: Lightweight Measurement of Web Censorship with Cross-Origin Requests. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication* (London, United Kingdom) *(SIGCOMM '15)*. Association for Computing Machinery, New York, NY, USA, 653–667. https://doi.org/10.1145/2785956.2787485

[11] Holmes Chan. 2019. Apple Music in China removes Jacky Cheung song with reference to Tiananmen massacre. https://hongkongfp.com/2019/04/09/apple-music-china-removes-jacky-cheung-song-reference-tiananmen-massacre

[12] Common Crawl. 2025. Common Crawl - Open Repository of Web Crawl Data. https://commoncrawl.org

[13] Alberto Dainotti, Claudio Squarcella, Emile Aben, Kimberly C. Claffy, Marco Chiesa, Michele Russo, and Antonio Pescapé. 2011. Analysis of country-wide internet outages caused by censorship. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference* (Berlin, Germany) *(IMC '11)*. Association for Computing Machinery, New York, NY, USA, 1–18. https://doi.org/10.1145/2068816.2068818

[14] United Arab Emirates. 2023. Federal Decree by Law Concerning Media Regulation. https://uaelegislation.gov.ae/en/legislations/2145

[15] United Arab Emirates. 2023. Federal Law by Decree Promulgating the Crimes and Penalties Law. https://uaelegislation.gov.ae/en/legislations/1529

[16] Marco Ferrari. 2022. Amazon opens new corporate office in Riyadh. https://english.alarabiya.net/business/2022/10/17/Amazon-opens-new-corporate-office-in-Riyadh

[17] Arturo Filastò and Jacob Appelbaum. 2012. OONI: Open Observatory of Network Interference. In *2nd USENIX Workshop on Free and Open Communications on the Internet (FOCI 12)*. USENIX Association, Bellevue, WA, 8 pages. https://www.usenix.org/conference/foci12/workshop-program/presentation/Filast%C3%B2

[18] Alireza Hashemzadegan and Ali Gholami. 2022. Internet Censorship in Iran: An Inside Look. *Journal of Cyberspace Studies* 6, 2 (2022), 183–204. https://doi.org/10.22059/jcss.2022.349715.1080

[19] Helmi Noman and Elijah Zarwan. 2007. *Internet Filtering in the Middle East and North Africa in 2006–2007.* Technical Report. Open Net Initiative. https://opennet.net/studies/mena2007

[20] Nguyen Phong Hoang, Arian Akhavan Niaki, Jakub Dalek, Jeffrey Knockel, Pellaeon Lin, Bill Marczak, Masashi Crete-Nishihata, Phillipa Gill, and Michalis Polychronakis. 2021. How Great is the Great Firewall? Measuring China's DNS Censorship. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, virtual event, 3381–3398. https://www.usenix.org/conference/usenixsecurity21/presentation/hoang

[21] Ben Jones, Tzu-Wen Lee, Nick Feamster, and Phillipa Gill. 2014. Automated Detection and Fingerprinting of Censorship Block Pages. In *Proceedings of the 2014 Conference on Internet Measurement Conference* (Vancouver, BC, Canada) *(IMC '14)*. Association for Computing Machinery, New York, NY, USA, 299–304. https://doi.org/10.1145/2663716.2663722

[22] Jeffrey Knockel, Jedidiah R. Crandall, and Jared Saia. 2011. Three Researchers, Five Conjectures: An Empirical Analysis of TOM-Skype Censorship and Surveillance. In *USENIX Workshop on Free and Open Communications on the Internet*. USENIX Association, San Francisco, CA, USA, 8 pages. https://www.usenix.org/legacy/event/foci11/tech/final_files/Knockel.pdf

[23] Jeffrey Knockel, Masashi Crete-Nishihata, Jason Q. Ng, Adam Senft, and Jedidiah R. Crandall. 2015. Every Rose Has Its Thorn: Censorship and Surveillance on Social Video Platforms in China. In *5th USENIX Workshop on Free and Open Communications on the Internet (FOCI 15)*. USENIX Association, Washington, D.C., 10 pages.

[24] Jeffrey Knockel, Masashi Crete-Nishihata, and Lotus Ruan. 2018. The effect of information controls on developers in China: An analysis of censorship in Chinese open source projects. In *Proceedings of the First Workshop on Natural Language Processing for Internet Freedom*, Chris Brew, Anna Feldman, and Chris Leberknight (Eds.). Association for Computational Linguistics, Santa Fe, New Mexico, USA, 1–11.

[25] Jeffrey Knockel, Jakub Dałek, Noura Aljizawi, Mohamed Ahmed, Levi Meletti, and Justin Lau. 2026. Banned Books public data set. https://github.com/citizenlab/chat-censorship/tree/master/amazon

[26] Jeffrey Knockel, Ken Kato, and Emile Dirks. 2023. *Missing Links: A comparison of search censorship in China.* Technical Report. The Citizen Lab. https://citizenlab.ca/2023/04/a-comparison-of-search-censorship-in-china

[27] Jeffrey Knockel and Lotus Ruan. 2021. *Engrave Danger: An Analysis of Apple Engraving Censorship across Six Regions.* Technical Report. The Citizen Lab. https://citizenlab.ca/2021/08/engrave-danger-an-analysis-of-apple-engraving-censorship-across-six-regions

[28] Jeffrey Knockel and Lotus Ruan. 2022. *Bada Bing, Bada Boom: Microsoft Bing's Chinese Political Censorship of Autosuggestions in North America.* Technical Report. The Citizen Lab. https://citizenlab.ca/2022/05/bada-bing-bada-boom-microsoft-bings-chinese-political-censorship-autosuggestions-north-america

[29] Jeffrey Knockel, Lotus Ruan, and Masashi Crete-Nishihata. 2017. Measuring Decentralization of Chinese Keyword Censorship via Mobile Games. In *7th USENIX Workshop on Free and Open Communications on the Internet (FOCI 17)*. USENIX Association, Vancouver, BC, 9 pages.

[30] Felix Lange, Niklas Niere, Jonathan von Niessen, Dennis Suermann, Nico Heitmann, and Juraj Somorovsky. 2025. I(ra)nconsistencies: Novel Insights into Iran's Censorship. In *Free and Open Communications on the Internet (FOCI)*. Proceedings on Privacy Enhancing Technologies, Washington D.C., USA, 7–12.

[31] Allison McDonald, Matthew Bernhard, Luke Valenta, Benjamin VanderSloot, Will Scott, Nick Sullivan, J. Alex Halderman, and Roya Ensafi. 2018. 403 Forbidden: A Global View of CDN Geoblocking. In *Proceedings of the Internet Measurement Conference 2018* (Boston, MA, USA) *(IMC '18)*. Association for Computing Machinery, New York, NY, USA, 218–230. https://doi.org/10.1145/3278532.3278552

[32] Paul Mozur and Vindu Goel. 2014. To Reach China, LinkedIn Plays by Local Rules. https://www.nytimes.com/2014/10/06/technology/to-reach-china-linkedin-plays-by-local-rules.html

[33] Arian Akhavan Niaki, Shinyoung Cho, Zachary Weinberg, Nguyen Phong Hoang, Abbas Razaghpanah, Nicolas Christin, and Phillipa Gill. 2020. ICLab: A Global, Longitudinal Internet Censorship Measurement Platform. In *2020 IEEE Symposium on Security and Privacy (SP)*. 135–151. https://doi.org/10.1109/SP40000.2020.00014

[34] Jack Nicas. 2019. Apple Removes App That Helps Hong Kong Protesters Track the Police. https://www.nytimes.com/2019/10/09/technology/apple-hong-kong-app.html

[35] Government of Dubai Media Office. 2023. Amazon opens a new Fulfillment Center in Dubai South. https://mediaoffice.ae/en/news/2023/March/21-03/Amazon-opening-a-new-Fulfillment-Center-in-Dubai-South

[36] The Embassy of the Kingdom of Saudi Arabia. 2003. Law of Printing and Publication. https://web.archive.org/web/20250415140052/https://www.saudiembassy.net/law-printing-and-publication

[37] Jay Peters and Nick Statt. 2019. Apple is hiding Taiwan's flag emoji if you're in Hong Kong or Macau. https://www.theverge.com/2019/10/7/20903613/apple-hiding-taiwan-flag-emoji-hong-kong-macau-china

[38] Samuel Ruo, Jeffrey Knockel, and Zoë Reichert. 2024. Lost in Translation: Characterizing Automated Censorship in Online Translation Services. In *Free and Open Communications on the Internet (FOCI)*. Proceedings on Privacy Enhancing Technologies, Bristol, UK, 93–101. https://www.petsymposium.org/foci/2024/foci-2024-0018.pdf

[39] Robert A Sedler. 1988. The Constitutional Protection of Freedom of Religion, Expression, and Association in Canada and the United States: A Comparative Analysis. *Case W. Res. J. Int'l L.* 20 (1988), 577.

[40] Ram Sundara Raman, Prerana Shenoy, Katharina Kohls, and Roya Ensafi. 2020. Censored Planet: An Internet-wide, Longitudinal Censorship Observatory. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security* (Virtual Event, USA) *(CCS '20)*. Association for Computing Machinery,

New York, NY, USA, 49–66. https://doi.org/10.1145/3372297.3417883

[41] Rima Tanash, Zhouhan Chen, Dan Wallach, and Melissa Marschall. 2017. The Decline of Social Media Censorship and the Rise of Self-Censorship after the 2016 Failed Turkish Coup. In *7th USENIX Workshop on Free and Open Communications on the Internet (FOCI 17)*. USENIX Association, Vancouver, BC, 7 pages. https://www.usenix.org/conference/foci17/workshop-program/presentation/tanash

[42] Rima S. Tanash, Zhouhan Chen, Tanmay Thakur, Dan S. Wallach, and Devika Subramanian. 2015. Known Unknowns: An Analysis of Twitter Censorship in Turkey. In *Proceedings of the 14th ACM Workshop on Privacy in the Electronic Society* (Denver, Colorado, USA) *(WPES '15)*. Association for Computing Machinery, New York, NY, USA, 11–20. https://doi.org/10.1145/2808138.2808147

[43] AppCensorship team. 2024. *Isolation By Design: Censorship in Apple's App Store in China and Human Rigths Implications.* Technical Report. The Great Fire Project. https://appcensorship.org/files/Isolation-By-Design.pdf

[44] Michael Carl Tschantz, Sadia Afroz, Shaarif Sajid, Shoaib Asif Qazi, Mobin Javed, and Vern Paxson. 2018. A Bestiary of Blocking: The Motivations and Modes behind Website Unavailability. In *8th USENIX Workshop on Free and Open Communications on the Internet (FOCI 18)*. USENIX Association, Baltimore, MD, 9 pages. https://www.usenix.org/conference/foci18/presentation/tschantz

[45] Benjamin VanderSloot, Allison McDonald, Will Scott, J. Alex Halderman, and Roya Ensafi. 2018. Quack: Scalable Remote Measurement of Application-Layer Censorship. In *27th USENIX Security Symposium (USENIX Security 18)*. USENIX Association, Baltimore, MD, 187–202. https://www.usenix.org/conference/usenixsecurity18/presentation/vandersloot

[46] Wikipedia contributors. 2025. Amazon (company). https://en.wikipedia.org/wiki/Amazon_(company)

[47] Yusuf Mücahit Çetinkaya and Tuğrulcan Elmas. 2025. State & Geopolitical Censorship on Twitter (X): Detection & Impact Analysis of Withheld Content. arXiv:2508.13375 [cs.SI] https://arxiv.org/abs/2508.13375