Studying Neutrality in Cyber-Space: A Comparative Geographical Analysis of Honeypot Responses

Martin Strohmeier[1]**, James Pavur[2]**, Ivan Martinovic[2], and Vincent Lenders[1]

[1] armasuisse Science + Technology, Thun, Switzerland
firstname.lastname@armasuisse.ch
[2] Department of Computer Science, University of Oxford, Oxford, UK
firstname.lastname@cs.ox.ac.uk

Abstract. Neutrality has long played a central role in the political and strategic stances of many small states. The impact of neutrality on matters of national security and sovereignty have thus been subject to significant academic interest. With the recent emergence of cyber as a fifth dimension of interstate competition, the impact of permanent state neutrality in this domain has not yet been well characterized.

We examine the reality of this concept using countries with a long-standing history and tradition of neutrality in matters of warfare and foreign policy. A theoretical analysis of the complexities of neutrality and cyber-crime is used to motivate a novel data-driven experimental assessment of real-world outcomes for neutral states.

This experimental study leverages low-interaction honeypots distributed across 13 countries. Delving into more than 1.5 billion network sessions made from these honeypots over an 80-day period reveals more than one million malicious attacks originating from information systems in 177 different countries. Through statistical analysis of these attacks, we find little evidence that low-sophistication adversaries target their attacks with consideration of victim location or state neutrality. Beyond the immediate implications of these findings, we believe the method presented in this paper represents a unique data-driven approach to comparative international study of cyber-neutrality and the global dynamics of cyber-security more broadly.

Keywords: cyber-neutrality · cyber-sovereignty · honeypots.

** Both authors contributed equally to this work.

1 Introduction

The strategic value of a national policy of neutrality in the modern era is an oft-debated topic. While much of this discussion has focused on conventional security and diplomatic practice, the intersection between cyber-space and state neutrality is not well understood. Given that cyber-aggression often occurs far below traditional thresholds for interstate conflict, lacks reliable attribution, and exhibits aterritorial properties, it remains an open question as to whether small-state neutrality has any practical effect on a nation's overall information security.

A natural baseline assumption for this hypothesis on neutrality effects may be that low-effort actors are looking for attack infrastructure or opportunistic targets and, hence, agnostic preferences to location seem intuitive. This is clearly in contrast to high-sophistication attacks using zero-day knowledge against high-value targets, which are all but apolitical. However, practitioners have long noted clear preferences among even low-level cyber-crime actors and their customers to avoid certain locations and countries depending on geopolitical allegiances, enforced for example via policies [22] or by checking undesirable locale characteristics of the target platforms [23].

We examine the reality of this concept using the example of Switzerland, a country with a long-standing history and tradition of neutrality. A theoretical analysis of the complexities of neutrality and cyber-crime is used to motivate a data-driven experimental assessment of real-world outcomes for formally neutral states including Switzerland and Singapore.

This assessment consists of a large-scale comparative study using purpose-built low-interaction honeypot installations deployed in 13 countries. The resultant dataset encapsulates more than 1.5 billion connections over a period of about two months. Analyzing the more than one million malicious attacks found within this dataset shows that there is little difference in low-level cyber-attacker behavior with regards to target geography or stance on neutrality. These findings present broader insights into the effect that neutrality may have on exposure to transnational cyber-crime.

The remainder of this work is organised as follows. First, we briefly introduce the literature on the general topic of neutrality in Section 2, before moving on to its application in cyber-space. Section 4 covers both aspects for our case study of Switzerland. Section 5 introduces our experimental setup, which is followed by the results, discussion and conclusion.

2 Background and Related Work

The role and relevance of permanent neutrality is a complex topic which has been subject to centuries of debate and interpretation [1]. A complete analysis thus lies well beyond the scope of this research. Nevertheless, it is worth considering those aspects of traditional neutrality theory which are particularly relevant to cyber-defence.

In this paper, we define neutrality simply as a permanent and public legal position which eschews warfighting as an instrument of foreign policy. This is distinct from the related concepts of non-alignment, and neutralism – which relate to diplomatic practices in multipolar systems and ad-hoc decisions regarding particular conflicts [20]. This narrow definition implies certain obligations and constraints under international law not only for the neutral state, but also for belligerents [3, 20]. Often, neutrality in practice incorporates aspects of non-alignment and diplomatic restraint, but this is not treated as an absolute requirement.

In practice, only a handful of modern states meet this definition. Even within this subset, specific cases are complex. For example, membership in the European Union challenges the ultimate neutrality of states like Austria and Sweden. Similarly, some states have only recently declared neutrality, such as Ghana (in 2012), Mongolia (in 2015), and Rwanda (in 2009). In such instances, these stances may not have accrued meaningful credibility with belligerents.

Historically, territorial sovereignty has been a vital lens for determining the operation of state neutrality. For example, a neutral state invaded by foreign military forces has an absolute right to defend itself and would not lose its neutral status in doing so. On the other hand, an invasion by a neutral state into the territory of another state would jeopardize its status, at which point retaliation by both the invaded state and its allies would be justifiable [3]. These principles extend beyond direct military invasion into even non-violent activities. For example, a neutral state could not permit the establishment of a foreign military base within its borders or sell arms to conflict participants [20, 3]. In return, belligerent parties are obligated to respect the sovereignty of the neutral state by, for example, avoiding incursions into their territorial waters and airspace.

In practice, complex variations on neutrality principles have emerged. Often, these tensions result from changes to the nature of warfighting. For example, military technologies required for effective Cold War deterrence surpassed the production capabilities of most neutrals. In order to sustain armed neutrality, these states paradoxically had to compromise on some of its principles for access to arms. This dynamic ultimately leads to Swiss and Swedish adherence to US-led sanctions regimes in exchange for the ability to import radar systems and other military technologies [26]. Similar compromises may be required in the cyber-context for access to privileged threat intelligence or technical assistance.

Moreover, modern conflict now incorporates non-state actors whose obligations to permanent neutrals is unclear. Recently, Swiss courts ruled that a private Swiss citizen leading a Christian militia in Syria had nevertheless violated the national principle of neutrality [31]. Similar cases have involved Swiss nationals suspected of joining ISIL-affiliated insurgent groups [10]. States seeking to maintain neutrality may be legally obligated to prevent such instances [24]. Moreover, non-neutral states may have a reciprocal obligation to ensure that their own citizens respect the rights of neutrals. Ad-hoc neutrals already justify foreign military operations on the failure of states to exercise this control over domestic radicals. For permanent neutrals, acceptable recourse is less clear. The analogy to the cyber-criminal context is intuitive. Traditional neutrality norms provide little guidance on the degree to which states are obligated to prevent their own citizens from attacking information systems abroad. While it is clear that all states still have an obligation to prevent harm emanating from a state's territory due to the due diligence principle, the specific impact of neutrality on this concept in cyber-space remains unclear both in theory and in practice.

3 Dynamics of Neutrality in Cyber-Space

In short, the meaning of permanent neutrality is not static. Neutrality and its challenges have given rise to a long history of exception and revision. These redefinitions are often political and normative, but they originate in response to positive effects. For example, in the Cold War, neutrality's side effect of technological weakness motivated the political decision to accept sanctions-linked arms deals in Switzerland and Sweden. The focus of this paper is not to answer the normative question as to how neutral states should respond to international cyber-crime. Rather, it is to identify the impact that neutrality has on exposure to international cyber-crime. While it is intuitive that cyber-threats will pose new challenges to neutrality, the form these will take is less clear. In this section, we propose some prominent factors which may impact neutral state cyber-security. Broadly, we consider three possible reasons that neutrality may increase a state's vulnerability to international cyber-crime. These are the difficulty of deterrence by denial, political barriers to deterrence by punishment, and challenges in intelligence sharing and targeted regulation. For realists, permanent neutrality is credible only if backed by significant defensive military force. This is not a universally accepted viewpoint [12, 17]. However, the link between neutrality and denial remains central to the professed foreign policy of many neutrals [6]. In the Swiss context, the combination of terrain, military spending, and mandatory conscription serve to dissuade territorial violations. It is not clear that Swiss efforts to defend domestic computer systems can achieve an equivalent effect. While effective denial is difficult in any domain, it is near insurmountable in cyber-space [5]. This suggests that credible armed neutrality in cyber-space cannot rest on the foundation of absolute defence. One alternative to deterrence by denial in cyber-space is the use of counterattacks and deterrence by punishment. States are increasingly asserting a legal right to retaliate in response to both state and non-state actors [21]. Even individual businesses have voiced interest in "hacking back" as a means to discourage cyber-criminal attacks [19, 25]. This retaliation may occur in cyber-space, but also via conventional or diplomatic channels. However, many domain features of cyber-space, such as attribution difficulties and unclear proportionality metrics, complicate deterrence by punishment [21, 33]. For permanent neutrals, deterrence by punishment is more complex. While neutrals have an indisputable right to self-defence, in conventional contexts designating a legitimate target or proportionate retaliation is much simpler than in cyber-space. Were a neutral to retaliate to a cyber-attack without absolute attribution and credibility, and especially if they were to do so via conventional means, it is unclear if the international community would view this as legitimate. Moreover, as indicated by aforementioned Swiss case law regarding private citizen involvement in the Syrian conflict, “hacking back” operations by corporate entities may similarly threaten the credibility of the state’s neutrality [31]. These dynamics may embolden adversaries, and, to the extent that deterrence by punishment works, neutrals may be unwilling or unable to avail themselves of its benefits. A third constraint for neutrals in cyber-space is political non-alignment. While not a definitional requirement, permanent neutrality frequently comes with some degree of non-alignment in foreign policy. This may impair certain functions of cyber-defence, such as reducing political willingness to engage in foreign intelligence collection [44]. Similarly, absent membership in collective security bodies (e.g. NATO) or alliances with major intelligence powers, access to shared cyber-threat intelligence may be limited. Finally, non-alignment may limit the defensive options available to a state. For example, several states have decided for a mix of political and technical reasons, to ban the use of 5G networking equipment from the Chinese telecommunications company Huawei [15]. Putting aside the specific merits of this action, it is unclear if political neutrals could engage in similar targeted trade actions without threatening the credibility of their overall non-alignment. Even with these constraints, state neutrality may nevertheless decrease exposure to international cyber-crime. We present three reasons that this might be the case. First, neutral states may make less attractive targets. Second, they may focus more effectively on defensive technologies. And third, neutral states may benefit from greater judicial reach in criminal prosecutions. While realists contend that neutrality is only as credible as the army which supports it, neutrality may impart strong cultural and political norms that insulate permanent neutrals from external threats. For example, neutral states often act as mediator in disputes between belligerents, accruing diplomatic capital and insulation from both sides of conflicts [27]. This mediator position may even bolster neutral states access to threat intelligence beyond that which is available within any single alignment-bloc [44]. The strong legal norms around neutrality may further disincentivize state-sponsored cyber-crime. Even non-state actors may be affected by the intangible “soft power” effects of permanent neutrality – especially those adversaries motivated by political objectives. Indeed, the long-standing reputation of Swiss neutrality has been characterized as a one of myriad factors potentially explaining the relatively low incidence of international terrorism within Swiss borders [43]. The realist requirement of credible self-defence is also not necessarily impossible for a neutral state. While defence is challenging in cyber-space, absolute defence is not always required. Criminals seek out the path of least resistance and having even marginally better security than peer states can discourage many attacks. As permanent neutrals have historically prioritized defensive military technology, policymakers may find significant investment in defensive cyber-technology more palatable. Indeed, as the role of great-power militaries shifts back towards territorial disputes rather than crisis management, this neutral advantage has been observed in conventional domains. Permanent neutrals are finding their relative expertise from focused investment in modern territorial defence much sought-after on the international stage [27]. Paradoxically, the very constraints imposed by permanent neutrality may be what enables effective focus on cyber-defence. Finally, in the context of transnational cyber-crime, neutral states may be better able to dissuade attacks with the threat of judicial punishment. One of the principle challenges in combatting cyber-crime is limited avenues for extradition and prosecution of foreign nationals [29]. Often, willingness to cooperate in transnational cases hinges on broader diplomatic relations between two states. In the case of a permanent neutral, these relations are less likely to be hostile. Of course, international extradition is a complex legal topic and difficult to generalize. However, in the Swiss case, law enforcement authorities have had limited success with extraditions to and from a wide array of countries crossing political and strategic blocs [38, 4]. As case law develops, neutrals may find a deterrence effect from greater judicial reach in transnational cyber-criminal prosecutions. Of course, it may also be the case that state neutrality has little to no effect on exposure to international cyber-crime. Cyber-criminals may not be aware of, or concerned with, the political stances of the countries they target. Far from a banal observation, this outcome would raise critical questions for the behavior of neutral states. If the appearance of neutrality in cyber-space, for example, did not cause any damage to the nation's defensive capabilities but bolstered the overall credibility of its neutrality, policymakers may prioritize actions which preserve this appearance. Conversely, if neutrality offers little benefit, but imposes costs on other functions, policymakers may decide that the appearance of neutrality in cyber-space is not an important priority. This brief analysis suggests that theoretical reasoning alone is unlikely to reveal clear answers as to the effect permanent neutrality has on state exposure to transnational cyber-crime. While we have suggested several factors for consideration, these represent only a small portion of the myriad challenges at the intersection of cyber-space and neutrality, many of which are well characterized elsewhere [14, 16, 42]. To bring new information to this debate, the remainder of this paper presents a comparative experimental case-study looking at data from real-world transnational cyber-attacks with a focus on the impact of Swiss neutrality.

4 Switzerland and Cyber-Sovereignty: A Brief Review

In the following, we will briefly review the public discussion around the topics of cyber-defence, cyber-sovereignty and neutrality in Switzerland. As acknowledged by most Swiss citizens and academic onlookers (e.g., [37, 36]), the Swiss stance on neutrality is both a key operating principle in Swiss diplomacy and holds a significant place in the country's identity. Paired with Swiss direct democracy, any development that has the potential to touch upon this neutrality can spark major debate within the legislature, media and the general public.

4.1 The Swiss Notion of Neutrality

Stolz [36] provides a brief history of neutrality in the Swiss context, which in the eyes of some historians reaches as far back as 1515. In the following 300 years, the states, which made up the Swiss confederation managed to avoid armed conflicts outside Swiss territory. As discussed by Suter [37], this state of affairs was made permanent after the events of the Vienna Congress and the Treaty of Paris in 1815. Switzerland's geographic placement at the heart of Europe and the interests of the European Great Powers at the time manifested this situation for the following centuries. Over time, in particular in the lead up towards the World Wars, the notion of neutrality in Switzerland changed towards a concept of "armed neutrality", whereby a substantial army is required in lieu of strong alliances to successfully deter attacks and preserve territorial integrity. While Swiss neutrality in World War 2 is sometimes viewed as far from perfect as secret consultations and forced economic collaborations happened [36], it is held that the concept is in part responsible for the relative peace Switzerland enjoyed during this time [37]. After World War 2, Switzerland continued on this path, notably forgoing UN membership during the Cold War (and until 2002) and not participating in economic sanctions. More recently, there has been the development of a notion of "active neutrality", whereby the Switzerland of the 21st century acts as a trusted broker and mediator between parties and countries, building on its credible image of impartiality. A vast majority of Swiss citizens supports the policy of neutrality, ensuring that its implementation will continue for the foreseeable future [36].

4.2 The National Strategy for the protection of Switzerland against Cyber-Risks

In 2018, the Federal Council of Switzerland published the second version of its National Strategy for the protection of Switzerland against Cyber-Risks (NCS) [39]. This document discusses at a high level the Swiss strategy to "secure and expand welfare [..] for the long term" in the face of digitalization. It takes into account several threat actors, including state actors, which are considered in the areas of cyber-espionage, cyber-sabotage or disinformation and propaganda. Beyond these, the NCS discusses the possibility of cyber-attacks in conflicts, which are acts just short of an all-out cyber war between state actors. Here, it is clearly stated that "Switzerland must therefore include cyber-defence and cyber-diplomacy in its preparations for potential conflict" [39]. In light of this, it is notable that there is explicitly no reference towards neutrality, neither in this context nor in the complete NCS document. It remains speculative whether this is intentional and a direct instantiation of "active neutrality" applied in cyber-space or instead a reflection of the fact that Switzerland's official position on neutrality is still in its infancy when it comes to non-conventional diplomacy. In April 2021, the Swiss Department of Defence announced the continued implementation of the NCS with regards to cyber-defence, the Cyber-Strategy for 2021-2024 [7]. In it, the concept of neutrality is not mentioned specifically, however it is noted that Switzerland has not been a target of attacks on its critical infrastructures yet. Potentially in light of political reality including neutrality, the Cyber-Strategy regards collateral damage as more likely than direct attacks targeting Swiss infrastructure specifically.

4.3 Early Public Debate on Cyber-Sovereignty in Switzerland

While the public debate around the topics of cyber-neutrality and cyber-sovereignty is in very early stages, there have been several events that have shaped the discussion in the past two years. The main discussion on this topic happened around Switzerland's accession as a contributing nation to the NATO Cooperative Cyber-Defence Centre of Excellence in 2018. In acknowledgment of the sensitivity of both engaging with a NATO-led centre of excellence in general and the cooperation with other states on cyber-defence in particular, the Swiss Federal Council argues that cooperation with the CCDCOE is non-problematic with regards the legal and political dimensions of Swiss neutrality [40]. More concretely, it is stated that the CCDCOE was not part of NATO's chain of command nor that it had an operational mandate. No rights or duties under international law could further be derived from participation and the scope of participation remains firmly in Switzerland's hands. This cautious stance has been reflected in Swiss media reports, for example regarding the visit of the 13th Secretary General of NATO, Jens Stoltenberg in 2017 or in the preceding process about participation in the Locked Shields exercise. Here, commentators note the difficult relationship of Switzerland and NATO but that cooperation in the cyber-domain could possibly be strengthened further [28]. Besides the application of the Swiss notion of neutrality towards cyber-space, the concept of "cyber-sovereignty" has seen increased uptake in Swiss government circles. The Federal Council's delegate for Cyber-Security stated that Switzerland must consider focusing on "security, education and neutrality" in order to be successful in the digital world, which would include an increase in security start-ups to protect Swiss ability to act in cyber-space ([28, 18]). Likewise, the head of the Federal IT Steering Unit considers a retreat towards the national arena an infeasible position, but that well-chosen international cooperation is required, which in turn needs to be adaptive to the situation [32]. The Swiss Federal Council's formal point of view on questions of international law and cyber-space is strongly informed by the Tallinn Manual [34]. In a parliamentary statement, the Department of Foreign Affairs takes the position that neutrality is a fundamentally applicable concept in cyber-conflicts. It is further stated that, while Swiss military law allows offensive responses against any networks where attacks originate from, this requires approval by the Federal Council and needs to be both permissible under international law and compatible with Switzerland's neutrality [8]. More recently, academics have begun considering the issue of Swiss neutrality specifically on cyber-operations. Stolz discusses a major challenge in this area, the "clash between national interest and the self-restrictions of neutrality." [36] Among other points, this dichotomy affects the national capacity of cyber-defence, which requires strong international collaboration and knowledge exchange, which is potentially at odds with the requirements of traditional neutrality.

Indeed, despite its state neutrality, Switzerland is not only a member of the CCDCOE but has been a founding member of Interpol, is party to the Council of Europe's Budapest Convention on Cybercrime, and even an active member of the Joint Cybercrime Action Taskforce (J-CAT) within Europol's European Cybercrime Centre (EC3). Finally, it is clear that both public opinion and government policy are constantly developing under the impression of current events.

5 Design of a Cyber-Neutrality Experiment

We will now describe the experimental design that we chose in order to test whether neutrality has an impact on largely automated attacks in modern cyber-space. First, we discuss the concept of Honeynet, a system developed to measure the number of attacks an ordinary Internet end point has to endure. Then we elaborate on the global deployment and distribution of Honeynet installations used to study neutrality's effects and the processing of the collected data.

5.1 Honeynet

Honeynet is a Docker-based collection of technologies that mimic the appearance of common, potentially insecure, web services. In this experimental scenario, Honeynet was configured to appear like an IoT device to attackers and simulate realistic targets for both automated and human cyber-attacks. Concretely, this involved deploying the embedded Linux toolkit BusyBox on each honeypot instance. The main attractive feature for potential attackers of our low-interaction honeypot is a telnet client with weak default credentials. Through exploiting this weakness, it is possible to enter the server and install software of the attacker's choosing. To enable good internet citizenship, all honeypot Docker-images were wiped and redeployed every 3 minutes, at which point the collected network data was transferred and stored at a central processing server in Switzerland in the free and open PCAP format. This was done to prevent any potential of real exploitation by attackers and misuse of the servers against further targets. This approach also ensures that all attempted connections were principally conducted by automated bots looking for easy targets.

5.2 Deployment

Our Honeynet sensors were deployed in 13 different countries around the globe. Besides Switzerland, four were deployed in Western Europe, three each in North America and Asia, two in Eastern Europe and one in the Middle East. Table 1 illustrates the deployment in more detail.

We deployed our Docker-based images on virtual machines (VMs) running Debian 8, 9 and 10. There are several notable insights to be reported from our deployment experiences. First, there is a notable absence of South America, Sub-Saharan Africa, and China. Throughout these regions, there are often stronger identification requirements for renting VMs. Proof of passport and residency requirements made it infeasible to deploy Honeynet in these regions for our initial study, but a more concerted effort in future work may prove beneficial.

Table 1. Deployment of the Honeynet test environment.

City	Country	Region	OS	IP4 Range
Amsterdam	Netherlands	Western Europe	Debian 10	142.93.*
Bangalore	India	Asia	Debian 10	165.22.*
Chișinău	Moldova	Eastern Europe	Debian 9	192.121.*
Frankfurt	Germany	Western Europe	Debian 10	206.189.*
Gravelines	France	Western Europe	Debian 10	137.74.*
London	United Kingdom	Western Europe	Debian 10	134.209.*
New York	United States	North America	Debian 10	165.227.*
San Francisco	United States	North America	Debian 10	157.230.*
Singapore	Singapore	Asia	Debian 10	134.209.*
St. Petersburg	Russia	Eastern Europe	Debian 9	213.183.*
Tel Aviv	Israel	Middle East	Debian 9	193.182.*
Thun	Switzerland	Western Europe	Debian 10	194.209.*
Tokyo	Japan	Asia	Debian 9	194.68.*
Toronto	Canada	North America	Debian 10	68.183.*

The second notable event was the attempt to deploy Honeynet on a commercial provider in Switzerland itself. Within a short time frame, MELANI, the Swiss government's reporting and analysis centre for information assurance, contacted the provider we utilized about the deployed server and the appearance of the IP address in international botnet structures. Thus, to not affect our study, we informed MELANI about our experiments. No other provider/country notified us about any similar issues.

5.3 Comparative Traffic Analysis

In total, the honeypots observed around 300 GB of unsolicited traffic from more than 1.5 billion sessions over an 80-day period. The open source tool Arkime (previously Moloch) was used to identify sessions and extract metadata from this traffic [2]. However, not all unsolicited internet connections are necessarily malicious. For example, many legitimate services perform whole-internet scans of active hosts for the purpose of research. To extract the most relevant data, we employed the Suricata network monitoring engine and several heuristic and signature-based intrusion detection rules to tag malicious traffic [9, 41, 11, 30]. This process enabled us to identify approximately 1.1 million malicious sessions originating from more than 100,000 unique attacker IP addresses in 177 countries. For those countries with multiple honeypots in the study, a random sample of data was taken in proportion to the total number of alerts observed. It is worth noting that attacker IP address does not provide a perfect indication of attacker location. Attackers may choose to purchase overseas cloud services or compromise vulnerable computers anywhere from which they can launch subsequent attacks. Nevertheless, servers located in a particular region may still demonstrate geographic effects either due to preference from local attackers or regulatory differences which impact attacker capabilities.

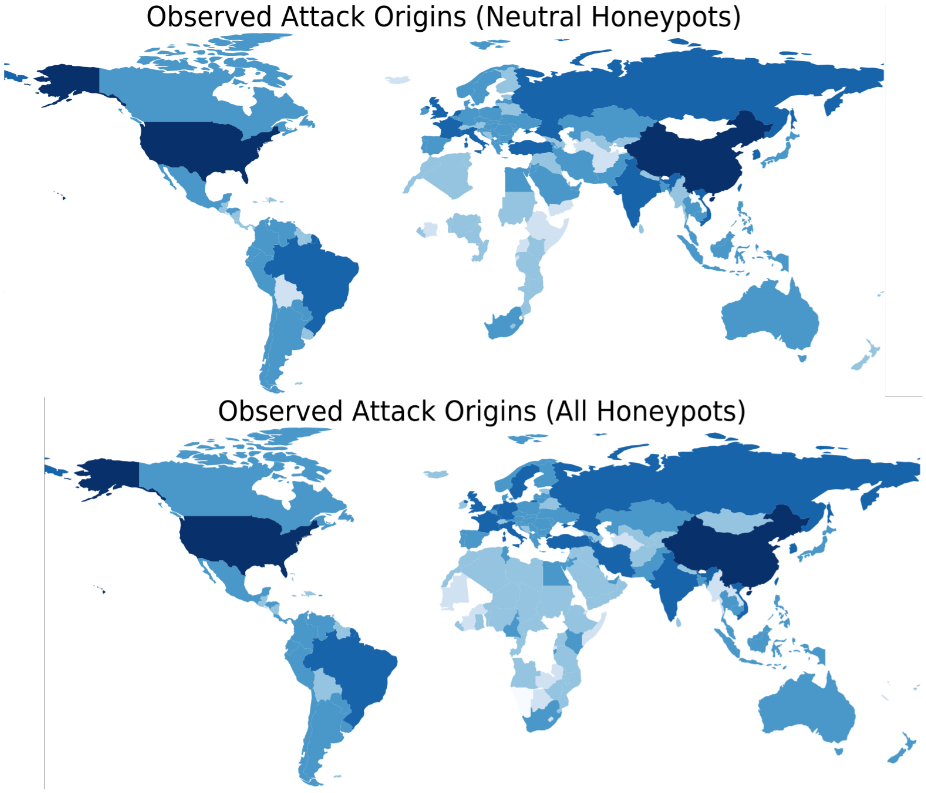

Observed Attack Origins (Neutral Honeypots)

Observed Attack Origins (All Honeypots)

Fig. 1. Comparative mappings of observed attacker origins.

6 Results

Across all honeypots, the distribution of attacker origins was roughly the same, suggesting that observed attacks were largely automated and agnostic to target location. The United States and China represented the principle locations associated with attacker IP address, accounting for roughly half of all observed attacks (Figure 1). At a macro level, only slight differences could be observed between the distribution of attackers targeting honeypots in the three formally neutral countries (Switzerland, Singapore, and Moldova) compared to the non-neutral honeypots (Figure 2).

A clearer sense of the relationship between attack quantity and geography can be achieved through correspondence analysis (Figure 3). The correspondence analysis was implemented using the open-source Python factor analysis library Prince [13]. In Figure 3, the distance between points on the chart is representative of the chi-squared distances between rows in the normalized contingency table associating attacker country to honeypot country. As a result, proximity between labels in the same dimension (e.g. attacker countries) suggests similar-

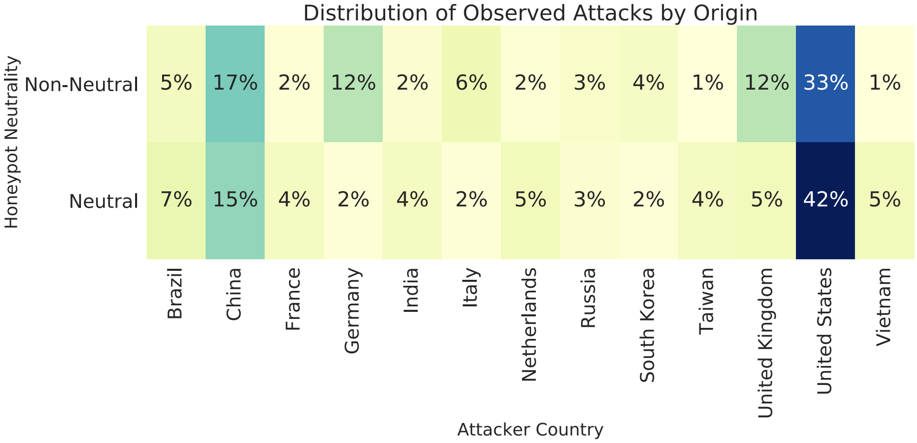

Fig. 2. Comparative distribution of attack origins in neutral vs. non-neutral countries for most observed attackers.

ity in observations. Points which are further from the origin are generally more discriminating/distinct from those which are closer. In this case, the sum of the inertia values for components 0 and 1 is high (89.97) suggesting that the correspondence analysis captures much of the variance in observed frequencies. See [35] for more information on interpreting correspondence analysis.

The clustering suggests that attackers from certain countries (e.g. China, Iran, and Russia) are relatively similar in terms of the honeypots which they targeted, while attackers from other countries (e.g. the United States and Vietnam) selected targets quite differently. Likewise, observed attackers for the Russian, Swiss, French, Israeli and Moldovan honeypots were similar, as were the origins for the cluster containing Canadian, British, and Dutch honeypots. More broadly, this analysis suggests that many attackers (those in the cluster towards the lower-left quadrant of Figure 3), behave similarly regardless of origin IP.

By calculating the uncertainty coefficient on a random sample of 500,000 observations from those attack origins which constituted a meaningful proportion of observed traffic (>0.1%), it is possible to better determine the strength of these relationships (Figure 4). This suggests that, while a bi-directional association between attacker origin and honeypot location exists, this association is quite weak ($U_1 = 0.17$ & $U_2 = 0.13$). The uncertainty coefficient, or Thiel's U, presented in Figure 4 is a measure of nominal association between two variables observed in our dataset. The value ranges from 0 (indicating no association) to 1 (indicating perfect association). For this paper, we consider values above 0.2 as moderately associated and values above 0.5 as highly associated. Unlike other metrics (e.g. Cramer's V), this value is asymmetric. So, for example knowing the attacker_country provides a very high degree of information as to the attacker_language ($U = 1$). However, knowing the *attacker_language* provides a slightly less (but still significant) degree of information as to the attacker_country

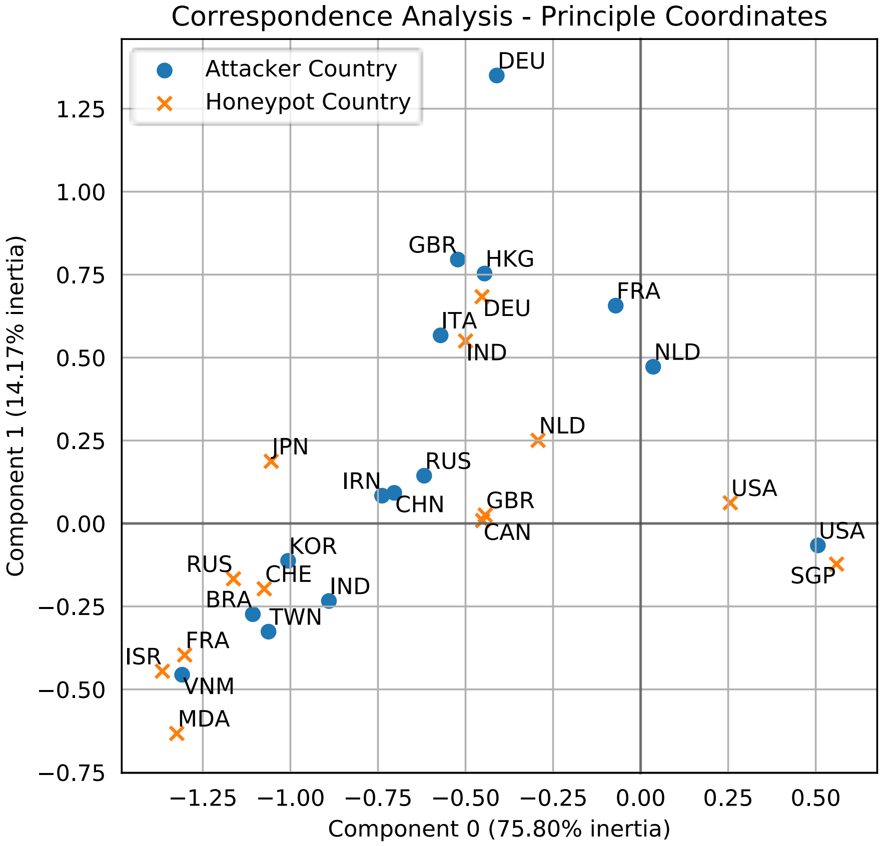

Fig. 3. Correspondence analysis of the contingency table associating attacker country to honeypot country frequencies.

($U = 0.8$). This makes sense as each country is keyed as having only one dominant language in our dataset, but many countries may share the same dominant language.

A moderate association between attacker origin and state neutrality is observed but the inverse association is almost non-existent ($U_1 = 0.20$ & $U_2 = 0.06$). That is, knowing a state is neutral does not provide much information about where its attackers come from, but knowing the origin of attackers may provide information as to whether they attack the neutral honeypots. This suggests that only a subset of attackers consider state neutrality (or some unconsidered third factor) in determining. Given the small number of neutral states both in our dataset (and globally) it is possible that a portion of this effect may be explained by the more general weak association between attackers and targeted countries.

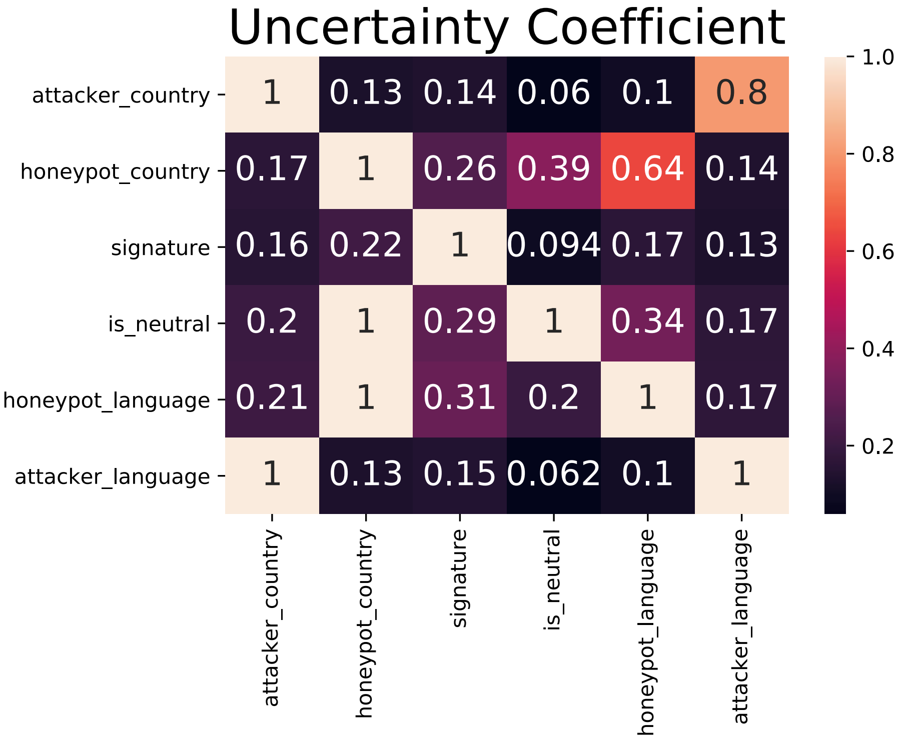

Fig. 4. Uncertainty coefficients.

It is also worth considering the dynamics involved in specific attack types. Here we find that the specific Suricata signature detected demonstrates a moderate degree of association with the target geography ($U_1 = 0.26$ & $U_2 = 0.16$). A similar association is observed for neutral honeypots, but it is difficult to distinguish these two relationships. This suggests that attackers demonstrate a slight preference for certain attack types depending on their own locations and the IP of their targets.

While the focus of this analysis is state neutrality, the method may prove useful for understanding other policy interactions with cyber-attack activity. For example, we found no correlation between attacker GDP per capita and the quantity of attacks ($\rho = 0.12$), a potentially surprising outcome as one might expect strong relationships between GDP per capita and availability of IT infrastructure from which to launch attacks. Inversely, honeypots in relatively wealthy countries were not significantly more likely to experience cyber-attacks ($\rho = 0.19$). Deeper research, especially with honeypots in the developing world, may bolster these insights.

7 Discussion

Reflecting on our experimental analysis, we find little evidence that low-level cyber-criminals using largely automated attacks have meaningful sensitivity to the national policies of their victim's countries – much less to national policies on neutrality. Indeed, the experimental data collected in this study most clearly supports the case that attackers are unaware or unconcerned with the geographic location of their targets altogether.

However, the case presented here is a simplified and cursory look at a complex problem and our methodology has several limitations, which we discuss in the following. We follow up by suggesting future work based on our methodology that can potentially shed more light on these questions.

7.1 Limitations

Most importantly, it must be assumed that by focusing on network traffic from attacks against generic telnet honeypots hosted by commercial virtual private server providers, this experimental data is inherently biased towards low-level threat actors. Intuitively it makes sense that attacks which are rudimentary in their means (e.g. telnet brute-force logins) also lack finesse with respect to their targeting.

While not the focus of this paper, there is of course always significant doubt about the true origin of an attacker. Attribution in cyber-space is hard and compromised machines similar to honeypots are typically used as a jump host for further attacks in order to obscure the true origin (not to mention the options of Tor, proxies or virtual private networks).

Finally, we appreciate that state neutrality as technically defined by international law applies to traditional armed conflict. With Cyber being defined as the 5th domain of warfare nowadays, this definition is notionally being broadened. More importantly, however, location preferences along the lines of state allegiances are commonly seen in the wild. One example is given by Brian Krebs, who cites a malware developer forbidding use of their ransomware tool against targets in the Commonwealth of Independent States (CIS) [22].

7.2 Future Work

Future work would benefit from increasing the study size on all dimensions. Of particular value would be geographical broadening to include non-represented countries (which may require the circumvention of several restrictions on overseas VM deployment) and an increased sample of neutral states. Additionally, a multi-year time horizon could provide deeper insights into the relationships considered here and related dynamics tied to political and economic changes. Addressing the issue of low-level threat actors would require much greater technical and logistical effort in future work. Sophisticated high-interaction, dynamic deception honeypots targeted at advanced persistent threats could show lateral movements of human (rather than automated) attackers and their targeting priorities for sensitive data and systems. This would provide clearer insights into the approaches of nation state actors towards neutral states in cyber-space. As such attacks are rarer than the threats considered in this paper, finding relevant activity at statistically meaningful scale would require a coordinated long-term effort. In the short term, it may be more effective to supplement the method presented here with a comprehensive data analysis of global attack reports describing the signatures of complex advanced persistent threats, such as Stuxnet.

8 Conclusion

For many countries, formal neutrality has defined their approach to diplomacy and warfare. Whereas a long history informs our understanding of neutrality's dynamics on land, air and sea, the picture is much less clear for emergent domains like cyber-space. In the process of developing an understanding of the links between state neutrality and transnational cyber-crime, this work also presents a novel experimental approach for testing general hypothesis at the intersection of cyber-space and national policy. By monitoring more than 1.5 billion connections to a research honeypot network spanning 13 countries over a two-month period, we isolated more than a million malicious cyber-attacks from more than 177 counties. Statistical analysis of this dataset suggests that low-sophistication attackers take little stock in victim state neutrality or geography more generally when executing their attacks. While this finding raises significant questions for policymakers seeking to deter cyber-attacks through political and legal means, it also suggests avenues for future experimental research considering more sophisticated cyber-attack dynamics.

Acknowledgments

The authors want to thank Dr. Luca Gambazzi for his invaluable support.

The views and opinions expressed in this article are those of the authors and do not necessarily reflect the official policy or position of any agency of the Swiss government.

References

1. Agius, C., Devine, K.: ‘neutrality: A really dead concept?’a reprise. Cooperation and Conflict **46**(3), 265–284 (2011)
2. Arkime: Arkime (April 2021), https://github.com/arkime/arkime
3. Bothe, M.: Neutrality, concept and general rules. Max Planck Encyclopedia of Public International Law. http://opil. ouplaw. com/view/10.1093/law:epil/9780199231690/law-9780199231690-e349 (2011)
4. Bradley, S.: Swiss back extradition with assurances. Swissinfo (Jan 2008), https://www.swissinfo.ch/eng/swiss-back-extradition-with-assurances/6376598

5. Brantly, A.F.: The cyber deterrence problem. In: 2018 10th International Conference on Cyber Conflict (CyCon). pp. 31–54. IEEE (2018)
6. Dalsjö, R.: 5 sweden and its deterrence deficit. Deterring Russia in Europe: Defence Strategies for Neighbouring States p. 2010 (2018)
7. Federal Department of Defence, C.P., Sport: Strategie cyber vbs (2021), https://www.vbs.admin.ch/de/verteidigung/schutz-vor-cyber-angriffen.html
8. Dobler, M.: Interpellation 18.3335: Cyberespace et droit international (March 2019), https://www.parlament.ch/en/ratsbetrieb/suche-curia-vista/geschaeft?AffairId=20183335
9. Foundation, O.I.S.: Suricata (April 2021), https://suricata-ids.org
10. Glaus, D., Vidino, L.: Swiss foreign fighters active in syria. CTC Sentinel **7**(7), 8–11 (2014)
11. Green, T.: Tgreen/hunting ruleset (April 2021), https://github.com/travisbgreen/hunting-rules
12. Guo, Y., Woo, J.J.: Singapore and Switzerland: Secrets to small state success. World Scientific (2016)
13. Halford, M.: Prince (April 2021), https://github.com/MaxHalford/prince
14. Healey, J.: When "not my problem" isn't enough: Political neutrality and national responsibility in cyber conflict. In: 2012 4th International Conference on Cyber Conflict (CYCON 2012). pp. 1–13. IEEE (2012)
15. Inkster, N.: The huawei affair and china's technology ambitions. Survival **61**(1), 105–111 (2019)
16. Jensen, E.T.: Sovereignty and neutrality in cyber conflict. Fordham Int'l LJ **35**, 815 (2011)
17. Jesse, N.G.: Choosing to go it alone: Irish neutrality in theoretical and comparative perspective. International Political Science Review **27**(1), 7–28 (2006)
18. Kaat, C.: Mr. cyber sagt, warum die schweiz mehr security-start-ups braucht. Netzwoche (Oct 2019), https://www.netzwoche.ch/storys/2019-10-14/mr-cyber-sagt-warum-die-schweiz-mehr-security-start-ups-braucht
19. Kallberg, J.: A right to cybercounter strikes: The risks of legalizing hack backs. IT Professional **17**(1), 30–35 (2015)
20. Karsh, E.: Neutrality and small states. Routledge (2012)
21. Kesan, J.P., Hayes, C.M.: Mitigative counterstriking: Self-defense and deterrence in cyberspace. Harv. JL & Tech. **25**, 429 (2011)
22. Krebs, B.: Is 'REvil' the New GandCrab Ransomware? (July 2019), https://krebsonsecurity.com/2019/07/is-revil-the-new-gandcrab-ransomware/
23. Krebs, B.: Try This One Weird Trick Russian Hackers Hate (May 2021), https://krebsonsecurity.com/2021/05/try-this-one-weird-trick-russian-hackers-hate/
24. Lloydd, M.: Retrieving neutrality law to consider 'other'foreign fighters under international law. In: European Society of International Law (ESIL) 2017 Research Forum (Granada) (2017)
25. McLaughlin, K.L.: Cyber attack! is a counter attack warranted? Information Security Journal: A Global Perspective **20**(1), 58–64 (2011)
26. Nilsson, M., Wyss, M.: The armed neutrality paradox: Sweden and switzerland in us cold war armaments policy. Journal of Contemporary History **51**(2), 335–363 (2016)
27. Nünlist, C.: Neutrality for peace: Switzerland's independent foreign policy. In: Engaged Neutrality: An Evolved Approach to the Cold War, pp. 161–187. Lexington Books (2017)

28. Nuspliger, N.: Die bedeutung von neutralität wandelt sich. Neue Zürcher Zeitung (Feb 2017), https://www.nzz.ch/schweiz/nato-generalsekretaer-stoltenberg-besucht-die-schweiz-die-bedeutung-von-neutralitaet-wandelt-sich-ld.148152?reduced=true
29. Perloff-Giles, A.: Transnational cyber offenses: Overcoming jurisdictional challenges. Yale J. Int'l L. **43**, 191 (2018)
30. Research, P.: Suricata pt open ruleset (April 2021), https://github.com/ptresearch/AttackDetection
31. Reuters: Ex-soldier is convicted of violating swiss neutrality by fighting isis. The New York Times (Feb 2019), https://www.nytimes.com/2019/02/24/world/europe/switzerland-soldier-isis.html
32. Rickenbacher, F.: Der bund will eine deutlich aktivere rolle übernehmen. Netzwoche (Oct 2019), https://www.netzwoche.ch/news/2019-10-16/der-bund-will-eine-deutlich-aktivere-rolle-uebernehmen
33. Ryan, N.: Five kinds of cyber deterrence. Philosophy & Technology **31**(3), 331–338 (2018)
34. Schmitt, M.N.: Tallinn manual 2.0 on the international law applicable to cyber operations. Cambridge University Press (2017)
35. Sourial, N., Wolfson, C., Zhu, B., Quail, J., Fletcher, J., Karunananthan, S., Bandeen-Roche, K., Béland, F., Bergman, H.: Correspondence analysis is a useful tool to uncover the relationships among categorical variables. Journal of clinical epidemiology **63**(6), 638–646 (2010)
36. Stolz, M.: On neutrality and cyber defence. In: European Conference on Cyber Warfare and Security. pp. 484–XIX. Academic Conferences International Limited (2019)
37. Suter, A.: Neutralität. praxis, prinzip und geschichtsbewusstsein. In: Hettling, M., Schaffner, M., König, M., Suter, A., Jakob, T. (eds.) Eine kleine Geschichte der Schweiz. Suhrkamp, Berlin (1998)
38. of Switzerland, A.G.: Coordinated operation in a cybercrime case. (Jul 2019), https://www.nytimes.com/2019/02/24/world/europe/switzerland-soldier-isis.html
39. of Switzerland, F.C.: National strategy for the protection of switzerland against cyber risks (ncs) 2018-2022 (2018), https://www.swissinfo.ch/eng/swiss-back-extradition-with-assurances/6376598
40. of Switzerland, F.C.: Teilnahme der schweiz am "cooperative cyber defence centre of excellence" (May 2019), https://www.admin.ch/gov/de/start/dokumentation/medienmitteilungen.msg-id-75145.html
41. Threats, E.: Emerging threats open ruleset (April 2021), https://rules.emergingthreats.net/
42. Turns, D.: Cyber war and the law of neutrality. In: Research Handbook on International Law and Cyberspace. Edward Elgar Publishing (2015)
43. Vidino, L.: Jihadist radicalization in switzerland. Tech. rep., ETH Zurich (2013)
44. Wylie, N.: 'the importance of being honest': Switzerland, neutrality and the problems of intelligence collection and liaison. Intelligence and National Security **21**(5), 782–808 (2006)