

What Motivates Cyberattacks: Lack of Consequences or Abundance of Attack Vectors?

Matthias Schulze and Florian Erdle

Institute for Peace Research and Security Policy (IFSH), Hamburg, Germany

schulze@ifsh.de

erdle@ifsh.de

Abstract: This study examines whether cyber-attacks are motivated by attacker impunity due to lack of deterrence or ease-of-attack due to offense dominance. We empirically measure whether ease-of-attack, measured through Common Vulnerabilities and Exposures (CVEs), drives cyberattack activity. Using global CVE and cyberattack data from 2000 to 2024, we find a statistically significant—though modest—correlation, with the strongest alignment appearing at a one-year lag. This suggests attackers typically take about a year to exploit new vulnerabilities. The findings lend conditional support to deterrence-by-denial, indicating that reducing vulnerabilities can meaningfully influence adversary.

Keywords: Deterrence by denial, Common Vulnerabilities and Exposures (CVEs), Cyberattacks, Correlation analysis

1. Introduction

There is limited agreement on why cyberattacks occur as frequently as they do and how they can be constrained. Cyber conflict scholarship offers two explanations: attackers strike because they face few consequences (impunity), or because cyber operations are structurally cheap and easy due to abundant vulnerabilities (ease-of-attack). This paper investigates whether an expanding technical attack surface helps explain the frequency of cyber operations. It asks two core research questions: Are annual trends in disclosed software vulnerabilities systematically associated with annual trends in high-salience cyber operations? Does this association differ over time and between state-centric and broader incident datasets?

The paper proceeds in four steps. First, it situates the research questions in the existing literature on cyber deterrence, offense–defense dynamics, impunity, and resilience, and derives an observable implication of the ease-of-attack perspective: if an expanding attack surface matters, periods with more disclosed vulnerabilities should, on average, coincide with more politically salient operations. Second, it develops a simple research design that operationalizes ease-of-attack via annual CVE counts (2000–2024) and “politically significant cyber operations” via annual incident counts from CFR and EuRepoC. Third, it examines both co-movement in levels - using Pearson, logPearson, and rank based (Spearman, Kendall) correlations to capture long-term association, phases specific shifts. Fourth, it interprets the temporal patterns in light of the theoretical debate, drawing out what they imply for claims about ease-of-attack, offense dominance, actor heterogeneity, and layered deterrence.

The paper makes three contributions. First, it offers a long-term empirical examination of the temporal association between CVEs and politically significant cyber operations. Second, it shows that the relationship between vulnerabilities and incidents is positive and robust in aggregate but heterogeneous across time and actor type: tightly aligned with state-linked operations in some periods, weaker or even negative in others. Third, it draws out the implications of these patterns for current debates on offence dominance, resilience, and layered deterrence, arguing that ease-of-attack is an important driver of observed activity but is modulated by institutional and strategic adaptations on both attacker and defender sides.

2. State of Research

Within the cyber-conflict literature, there is a vivid discussion on the causes and prevention of adversarial cyber-attacks. The discussion is conceptually based on deterrence studies and the offense-defense balance (Malone 2012, Nye 2017, Slayton 2017). These frameworks try to assess the relative ease, costs, and benefits of initiating an attack versus defending against it (Jervis 1978, van Evera 1998, Gafinkel Dafoe 2019).

The offense-defense balance in cyberspace remains contested because offensive and defensive cyber capabilities are hard to distinguish because they rely on similar tools and are hard to quantify (Valeriano 2022). Many argue that offense dominates due to abundant vulnerabilities, automation, and imperfect defenses (Saltzman 2013; Slayton 2017; Huntley & Shives 2024). Others note that defenders’ system knowledge and network control offer advantages given sufficient investments (Campell & Donahoo 2024). Our focus, however, lies on why cyberattacks occur in the first place.

There are two competing explanations. The first argues that cyber-attackers are motivated because of attacker impunity, i.e., they must fear no consequences. Attacker impunity is facilitated by several factors. Cyber-operations can provide a degree of anonymity, and attackers can obfuscate their identities and hide their traces, making attribution of cyber-attacks to perpetrators slow and complicated (Kello 2017, Brantly 2018, Libicki 2009). If cyberattacks can be conducted with impunity, the attacker has little reason to stop. Because attribution is time-consuming and not always feasible and law enforcement and apprehension of cyber-operators is relatively weak, some scholars make the case for deterrence by punishment: Deterrence by punishment is a strategy aimed at discouraging an adversary from initiating or escalating a conflict by convincing them that the costs of such actions will outweigh any anticipated benefits. Aggressors must be convinced psychologically, through communication and credibility, that they will be identified and punished. Cyber-deterrence by punishment can occur through *in-kind* retaliation using offensive capabilities (Healey 2019) to raise the costs of attacks (Borghard & Lonergan 2023) or *cross-domain* measures such as sanctions, legal actions, public attribution, or limited conventional strikes (Libicki 2009; Van de Velde 2023). Yet, effectiveness is disputed: technical limits and delayed attribution undermine credibility (Borghard & Lonergan 2017; Brantly 2018; Soesanto 2022). Consequently, research finds cyber deterrence by punishment largely impractical.

The other school of thought argues that cyberattacks are driven by the ease of attacking, due to offense dominance (Borghard & Lonergan 2023). First, cyberattacks are comparatively cheap and easy to execute (compared to conventional attacks or covert operations) and thus have a low barrier for entry (Libicki 2011, Bendiek & Metzger 2015, Lonergan and Montgomery 2021). Second, due to architectural vulnerabilities of the global internet, there is an abundance of software and hardware vulnerabilities, in addition to human weaknesses, that can be exploited for cyberattacks (Slayton 2017, Taddeo 2018).

The anticipated strategic gains of many cyberattacks, whether they are financially motivated cyber-crime (ransomware & fraud) or intelligence-driven cyberespionage operations, often outweigh the assumed costs (including punishment) in the calculus of the attacker (Fischerkeller, Goldman & Harknett 2022, 37). Research from other disciplines stresses reward structures as well, whether it is profit, psychological thrill, fame, power and status as primary motivators for cyberattacks (Aldasoro et al. 2022, Lim & Thing 2022). In other words, it is lucrative to attack, and having more to gain than to lose is a powerful motivator for attackers.

These dynamics encourage adversaries to exploit vulnerabilities and launch cyber operations. If ease-of-attack is the root cause, defenders have two goals: strengthening defenses to hinder intrusions and discouraging attacks by convincing adversaries that success is unlikely (Nye 2017). Deterrence by denial thus seeks to prevent attacks externally by changing the attackers calculus and psychology, while defense focuses inward on damage reduction. Defense is generally more feasible as it depends on resource investment rather than influencing attacker psychology (Fischer 2019).

Fischerkeller and Harknett argue that there is little evidence for successful deterrence by denial. i.e., demotivating advanced persistent adversaries from attacking, given the prevalence of continuous cyberattacks against systems (Fischerkeller & Harknett 2017). In addition, nowadays defenders must “assume breach”, that networks already have been compromised. Therefore, many experts and scholars argue for cyber-resilience: the ability to rapidly recover from successful attacks, in case both defense and deterrence by denial fail (Fischer 2019, Valeriano et al. 2022). However, there is no empirical research that tests these assumptions.

3. Methodology

This study examines whether the ease-of-attack argument is systematically and empirically related to politically significant cyber operations over time. The empirical strategy focuses on one observable implication of the ease-of-attack perspective rather than on attacker psychology. Ease-of-attack is determined by available software vulnerabilities (IV). These represent technical weaknesses that can be exploited by malicious actors to compromise systems. In principle, a larger stock of unpatched vulnerabilities increases the attack surface and expands the set of feasible operations for capable attackers. Under this logic, more publicly known vulnerabilities should be associated with more opportunities for exploitation and, ultimately, with higher recorded activity in incident datasets that track politically relevant cyber operations (DV).

Formally, the main hypothesis and its null counterpart are:

H1: Annual counts of publicly disclosed vulnerabilities are positively associated with annual counts of politically significant cyber operations.

H0: If attack surface is not a key driver, vulnerability counts and operation counts will be largely uncorrelated once general time trends are taken into account.

These hypotheses concern temporal association, not causation. The analysis therefore focuses on whether vulnerability and incident series move together in systematic ways, rather than on estimating causal effects or modeling attacker decision-making directly.

3.1 Data

The primary data source documenting the annual number of publicly disclosed software vulnerabilities is the Common Vulnerabilities and Exposures (CVE). The CVE system provides a standardized dictionary of unique identifiers for publicly disclosed software vulnerabilities, enabling consistent tracking and communication across security tools, vendors, and researchers. Launched in 1999 and originally maintained by MITRE Corporation under the U.S. Department of Homeland Security, funded via CISA, it relies on CVE Numbering Authorities (CNAs)—such as vendors, researchers, and organizations like Microsoft—to validate reports from stakeholders, assign IDs, and publish entries to a central list enriched by data publishers. For our purpose, annual CVE counts are aggregated by the year of disclosure. We analyze the 2000–2024 period.

CVE entries underpin both defensive and offensive cybersecurity practices, supporting patch management, penetration testing, and red-teaming. Public vulnerabilities feed into exploit kits (Metasploit) and automated scanning tools, making CVEs operationally significant for state and non-state actors alike. Yet CVE volumes also reflect defensive activity—the extent of research, discovery, and disclosure within the security community.

This dual nature complicates interpretation: CVE counts measure research intensity as much as exploitability. Expanding disclosure practices, vendor compliance with reporting norms, and regulatory pressures have increased the likelihood of CVE assignment, inflating counts without necessarily reflecting greater code weakness. Likewise, bug bounty programs, crowdsourced testing, and advances in automated scanning have accelerated vulnerability discovery, driving CVE growth independent of shifts in actual attack surface.

The analysis therefore treats CVEs as an imperfect but informative proxy for the evolving attack surface. The ease-of-attack perspective expects that, despite these measurement dynamics, an expanding and catalogued vulnerability landscape should, on average, be associated with more opportunities for politically significant cyber operations.

Politically significant cyber operations (DV) are measured using two complementary datasets: 1) The Council on Foreign Relations Cyber Operations Tracker records publicly known, state-linked cyber operations with foreign policy relevance since 2005, focusing on incidents attributed to governments or industry (Council on Foreign Relations n.d.). 2) The European Repository of Cyber Incidents (EuRepoC) documents cyber incidents with political dimensions and critical-infrastructure impact since 2000, being initiated by both state and non-state actors and a broader range of targets (Zettl-Schabath et al. 2025). It captures 3400 incidents, 60 coded variables based on a daily collection of 220 news sources.

For each dataset, annual incident counts are measured over the period 2000–2024 (EuRepoC) and 2005–2024 (CFR). CFR emphasizes state-sponsored, foreign policy-driven operations, whereas EuRepoC also includes incidents by non-state actors such as hacktivists and criminals. Divergent patterns across the two series can therefore shed light on whether any association with vulnerabilities is specific to particular actor types or incident profiles. From an ease-of-attack perspective, one might expect that an expanding vulnerability landscape affects both types of actors, but possibly with different timing and intensity. For instance, non-state actors may exploit widely known vulnerabilities opportunistically and quickly, while state actors may selectively use both known and undisclosed vulnerabilities for more targeted campaigns. The analysis therefore treats divergence between CVE–CFR and CVE–EuRepoC correlations across periods as substantively meaningful, not merely as noise.

Both datasets are subject to selection and reporting biases. They cover only publicly known incidents, rely on media and open sources, and are shaped by evolving coding practices and geopolitical attention. The analysis thus treats their counts not as comprehensive measures of global cyber activity but as indicators of the temporal dynamics of high-salience, politically relevant operations.

All variables are aggregated to the annual level. This choice reflects both data availability and the study's focus on long-term trends rather than short-lived spikes, recognizing that reporting practices can vary within years and that many politically significant operations unfold over weeks or months. Yearly aggregation smooths some of this volatility and makes broad temporal patterns more visible.

3.2 Measurement


The core quantitative strategy is correlation analysis between annual CVE counts and annual incident counts. To address skewed distributions and long-run growth, and concerns about spurious correlation from common trends. The analysis proceeds in four steps:

1. Pearson correlations between raw annual counts of CVEs and incidents are computed to capture linear association over the full period and within 5-year windows.
2. Pearson correlations on log-transformed values are calculated to attenuate the influence of extreme values, especially the sharp rise in CVE counts in later years, and to reflect proportional rather than absolute changes. The goal is to account for possible skewed distributions, which is very often the case within cyber data and better captures underlying log-linear relationships.
3. Rank-based correlations - Spearman's rho and Kendall's tau - are computed for both the full period and five-year windows to assess whether monotonic relationships exist even when functional forms are non-linear or affected by outliers. All three measures range between -1 and 1 , indicating direction and strength of association, where values close to ± 1 indicate strong relationships and values near 0 indicate weak or no relationships, and conventional significance levels are reported for the full-period correlations. For the five-year windows, coefficients are interpreted descriptively, as the small sample size precludes strong inferential claims.
4. The analysis considers concordance in growth rates to address residual concerns about shared trends and to approximate dynamic responsiveness to changes in the vulnerability landscape. For each series, year-on-year changes (first differences) are computed, and concordance scores between 0 and 1 are used to summarize the share of years in which CVEs and incidents move in the same direction and jointly accelerate or decelerate. A value of 1 indicates perfect alignment; 0 indicates perfectly opposing movement, and 0.5 indicates no systematic association. Because attackers may respond to newly disclosed vulnerabilities with a delay, concordance is calculated for three scenarios: no lag, a one-year lag, and a two-year lag of incident changes behind CVE changes.

This combined strategy offers a transparent, data-driven assessment of whether and when publicly disclosed vulnerabilities and politically significant cyber operations exhibit systematic temporal alignment and how this alignment differs across state-focused and broader incident datasets. Within the constraints of annual data, common upward trends can mechanically generate strong associations and relatively short time series. The results are interpreted as evidence of co-movement consistent with the ease-of-attack perspective, not as proof of a causal effect of CVEs on cyber operations.

4. Empirical Findings

The empirical analysis reveals pronounced temporal and actor-specific shifts in the relationship between CVEs and politically significant cyber operations. Over the full period, both CFR and EuRepoC show strong positive associations with CVE counts, but the strength and sign of these relationships vary across subperiods and between state and non-state actors.



Long-run correlations

Over the full observation period, correlations between CVEs and both CFR and EuRepoC are positive, relatively large, and statistically significant across all correlation measures. Total Pearson correlations between CVEs and CFR incidents (0.624), and between CVEs and EuRepoC incidents (0.854), are strong and change slightly when log transforms are used, indicating that the association is not driven solely by a few extreme years. Non-parametric rank correlations (Spearman and Kendall) yield similarly robust positive coefficients in aggregate, suggesting that the positive relationship holds for monotonic co-movement as well.

These long-run patterns clearly reject a strict null hypothesis of “no relationship”, in which vulnerability counts and operation counts would be largely uncorrelated once general time trends are considered. At this coarse temporal level, more disclosed vulnerabilities coincide with more recorded cyber-operations, which is consistent with the notion that an expanding attack surface is associated with more politically relevant cyber activity.

Temporal phases and actor heterogeneity

The five-year correlations highlight substantial temporal heterogeneity. In the early years (2000–2009), correlations between vulnerabilities and cyberattacks are very weak (Pearson(log): -0.073) for CFR (2005–2009) and weakly negative (Pearson(log): -0.364 2000-2004 and -0.340 2005-2009) for EuRepoC, indicating little synchrony between vulnerability disclosures and the rate or visibility of documented cyber-operations. During this phase, CVE volumes rise, but incident reporting - especially for geopolitically cyber-operations - remains sparse, plausibly reflecting underreporting, lower global awareness, and immature tracking infrastructures.

From 2010 to 2014, correlations between CVEs and CFR increase sharply, reaching values close to 0.9 for both Pearson and log-Pearson measures, while correlations with EuRepoC remain very weak to weak (Pearson: 0.083 and Pearson(log):0.227) In the 2015–2019 period, these high CVE–CFR correlations persist (Pearson(log):0.914) , indicating sustained alignment between the volume of disclosed vulnerabilities and recorded state-sponsored activity, as both CVE counts and CFR incidents grow markedly. This pattern suggests a phase in which expanding attack surfaces and state-linked operations move in close temporal step, at least in the observable record.

By contrast, the relationship between CVEs and EuRepoC turns strongly negative (Pearson(log): -0.718) in 2015–2019, even as CVE x CFR correlations remain strongly positive (Pearson(log): 0.914). Possible explanations include a tactical shift among non-state and mixed-motivation actors towards attack vectors less directly tied to known software vulnerabilities, as well as a denominator effect: annual CVE counts increase steeply while EuRepoC incident counts grow slowly and partly plateau, which can mechanically generate negative correlations despite an underlying positive link.

In the most recent period (2020–2024), a clear divergence emerges: correlations for CVE x EuRepoC become strongly positive (Pearson(log): 0.900), whereas those for CVE x CFR turn sharply negative, particularly for the Pearson coefficients (Pearson(log): -0.769). This pattern is consistent with a saturation effect in vulnerability reporting, in which rapidly growing CVE counts no longer translate into higher rates of recorded state-linked operations, while broader, non-state and mixed-motivation activity captured by EuRepoC remains more closely aligned with disclosed weaknesses. Taken together, the temporal phases point to an association that is positive and robust in aggregate but heterogeneous across time and actor type.

Concordance of growth rates

The concordance analysis of year-on-year changes provides a complementary view focused on dynamic alignment rather than levels. Across all three datasets, a one-year lag consistently yields the highest concordance between changes in CVE counts and changes in incident counts, with EuRepoC non-state and others reaching a concordance score of 0.619. This implies that in roughly 62% of years, CVEs and non-state incidents move in the same direction and jointly accelerate or decelerate when non-state incidents are allowed to respond with a one-year delay to changes in vulnerabilities. For EuRepoC total and CFR, lag-1 concordance also exceeds the contemporaneous values, but remains only modestly above the 0.5 baseline that would indicate no systematic association.

	CVEs		
	lag 0 concordance	lag 1 concordance	lag 2 concordance
EuRepoC	.348 (n=23)	.565 (n=23)	.522 (n=23)
EuRepoC (non-state and others)	.238 (n=21)	.619 (n=21)	.476 (n=21)
CFR	.474 (n=19)	.579 (n=19)	.368 (n=19)

These values point to moderate, not tight, dynamic alignment. They suggest that changes in the vulnerability landscape and changes in recorded cyber incidents move together more often than would be expected by chance, especially for non-state actors, but far from perfectly. In other words, the concordance patterns are compatible with a view in which some actors adjust activity with a delay to shifts in the vulnerability environment, yet they do not show strong or systematic responsiveness across all actors and periods.

The graphical comparison of annual changes further highlights differences in temporal alignment across datasets.. For CFR (state-affiliated and state-sponsored attacks), similarity in growth rates is largely confined to a narrow window around 2015–2019, with little sustained correspondence outside this interval. For total EuRepoC incidents, synchrony is most between roughly 2012 and 2020 but is less consistent beyond that. A similar pattern holds for EuRepoC non-state incidents in which both CVEs and non-state incidents rise, flatten, or decline in broadly similar fashion.

Overall, the concordance analysis reinforces the aggregate correlation findings by showing modest dynamic co-movement, particularly for non-state actors and when a one-year lag is allowed, but it also underscores the limits of the association and cautions against strong claims about tight intertemporal coupling.

5. Discussion

Over the full period, strong positive correlations between CVEs and both CFR and EuRepoC reject the null hypothesis of no relationship. In the long run, more disclosed vulnerabilities align with increased high-salience cyber operations, supporting the view that a growing attack surface drives politically significant activity.

However, the temporal patterning complicates any simple reading that more CVEs automatically mean more attacks. The early 2000s with weak or negative correlations suggest that, when vulnerability disclosure and offensive cyber-postures were still in development, increases in CVEs did not immediately translate into a higher visible rate of cyber-operations. Only once both vulnerability ecosystems and offensive cyber-capabilities matured—roughly from 2010 onward—do CVE and incident series move in closer synchrony, particularly state-linked operations in CFR. Note for example, that many countries created their first cyber-command structures, able to launch offensive operations, around the year 2010.

This staged pattern fits an ease-of-attack story that is mediated by institutional capacity: vulnerabilities may expand opportunity structures, but they only show up in the observable record of political cyber operations especially once they became a tool of international statecraft around 2010, represented by high profile operations such as (Stuxnet 2010, Operation Aurora 2010, Shamoon 2013 and others).

The late-period divergence—where CVE–EuRepoC correlations become strongly positive while CVE–CFR correlations turn negative— is puzzling. One explanation could be that the EuRepoC dataset is more reactive than the CFR dataset. CFR only lists cases with clear state-attribution. EuRepoC also lists unattributed incidents, which is especially the relevant since the Russian invasion of Ukraine in 2022, which coincided with a splurge in hacktivist cyber-operations. Future research should evaluate whether this due to the delay in attribution activity or represents a change in attack patterns. It is also possible that changes in political signalling, attribution capabilities, and the credibility of punishment selectively affect state behaviour, while a large pool of non-state actors continues to exploit readily available vulnerabilities, but these mechanisms cannot be identified directly from the present design.

The temporal dynamics connects to the offense–defense and deterrence literature. The strong positive correlations during 2010–2019 between CVEs and CFR incidents align with arguments that offense enjoys structural advantages in cyberspace as long as vulnerabilities are abundant and patching is imperfect. In those years, the data resemble a world in which expanding attack surfaces and growing offensive programmes move

in lockstep, strengthening the intuitive link between exposed weaknesses and the volume of state-sponsored operations.

The later divergence, however, suggests that defensive investments, changes in patching practices, and evolving norms might reshape this relationship, as vulnerability counts continue to grow rapidly while state-linked operations do not keep pace and in some windows even move in the opposite direction —this is consistent with at least four possibilities highlighted by the literature:

- Improved denial and resilience: Investments in cyber capacity, incident response, and resilience may make it harder or less attractive for states to rely on widespread exploitation of known vulnerabilities, even as those vulnerabilities continue to be discovered and catalogued.
- Strategic substitution: State actors may increasingly rely on Odays, Denial of Service attacks, social engineering, password spraying or living-off-the-land techniques without the use of software vulnerabilities to achieve political objectives, decoupling their activity from the public CVE stream. Another explanation could be that China decoupled from the global CVE system and built its own vulnerability ecosystem. Since 2021 Chinese security researchers do no longer report CVE's to the US-led system but have to report them to Chinese vulnerability databases and ministries (Cary & Del Rosso 2023).
- Selective enforcement and signalling: Stronger attribution regimes, such as the EU Cyber-Diplomacy Toolbox (since 2017) and more or offensive-denial strategies such as through US persistent engagement (since 2018) could deter some forms of state behaviour without changing the underlying technical attack surface.

The concordance analysis refines this picture by focusing on dynamic alignment. The modestly elevated concordance with a one-year lag, particularly for non-state actors in EuRepoC, is compatible with a view in which at least some attackers adjust activity with delay to changes in the vulnerability landscape, which is a necessary condition for deterrence-by-denial mechanisms to operate. At the same time, concordance values clustered only slightly above 0.5 underline that dynamic alignment is moderate, not strong, and that the short effective sample size cautions against firm claims about the presence or effectiveness of denial strategies. The analysis can therefore indicate compatibility with such mechanisms but cannot demonstrate that denial measures have actually reduced attack frequency.

The differences between CFR and EuRepoC correlations and concordance scores underscore the importance of actor heterogeneity. EuRepoC's broader coverage of nonstate and mixed-motivation incidents appears more tightly coupled to CVE volumes in the most recent period, while the state-focused CFR series decouples both in levels and growth rates. This pattern fits a view in which: Nonstate actors and opportunistic groups continue to benefit directly from a growing pool of known vulnerabilities, especially when they can rapidly integrate published CVEs into their tools off-the-shelf tools (such as the Metasploit Framework). State actors may be constrained by strategic, legal, or normative considerations that make reliance on easily attributable, widely known vulnerabilities less attractive over time. For instance, Chinese operations often favour stealth and hidden long-term access, often based on 0-days. In contrast, Russian operations are relatively noisy and also utilize known CVEs (Steiger 2022).

In theoretical terms, the ease-of-attack logic may apply most straightforwardly to actors for whom marginal cost and tool availability are primary constraints, whereas more sophisticated state actors operate under a more complex mix of constraints and incentives such as long-term planning, or techniques that are less directly tied to vulnerabilities. The results therefore suggest that offense dominance may be segmented: strong and persistent for certain classes of actors, attenuated or reshaped for other.

6. Conclusion

Overall, the findings support a nuanced version of the main hypothesis: publicly disclosed vulnerabilities and politically significant cyber operations generally move together over the long run, but this relationship is contingent on time, actor type, and institutional context. However, correlations cannot establish causality or uncover attacker motivations. Future work should move beyond correlation by explicitly modelling time trends, unit roots, and lag structures. Simple models that correlate incidents with lagged CVEs (and vice versa), or that correlate detrended residuals rather than raw levels, would help distinguish shared growth from more substantive temporal linkages. With richer data, vector autoregression or panel timeseries designs could test whether changes in vulnerability exposure systematically precede changes in incident rates. This study focuses on publicly disclosed vulnerabilities. Incorporating information on Odays, exploit markets, and stockpiling

practices—where available—would help clarify when state behaviour decouples from public CVE streams and relies instead on private vulnerability arsenals. This would speak directly to debates about strategic substitution and the limits of attack surface measures based solely on disclosed vulnerabilities.

Ethics Declaration: Ethical approval was not required for this research.

AI Declaration: Perplexity AI was used to assist in streamlining drafts of the manuscript. The final text was streamlined by the authors.

References

- Aldasoro, I., Fender, I., Hardy, B. and Tarashev, N. (2022) 'The drivers of cyber risk', *Journal of Financial Stability*, 60, DOI: 10.1016/j.jfs.2022.100989.
- Bendiek, A. and Metzger, T. (2015) Deterrence theory in the cyber-century: Lessons from a state-of-the-art literature review. *Stiftung Wissenschaft und Politik (SWP) Working Paper*, (02), https://www.swp-berlin.org/publications/products/arbeitspapiere/Bendiek-Metzger_WP-Cyberdeterrence.pdf, (Accessed 22 January 2026).
- Borghard, E.D. and Loneragan, S.W. (2017) The Logic of Coercion in Cyberspace. *Security Studies*, 26(3), pp. 452-481.
- Borghard, E.D. and Loneragan, S.W. (2023) Deterrence by denial in cyberspace. *Journal of Strategic Studies*, 46(3), pp. 534-569.
- Brantly, A.F. (2018) 'The Cyber Deterrence Problem', in *Proceedings of the 10th International Conference on Cyber Conflict (CyCon)*. Tallinn: NATO CCD COE Publications, pp. 31-54.
- Campbell, P. and Donahoo, M.J. (2024) 'Harnessing the Power of Cyber Defence', in *Survival: April-May 2024*. London: Routledge, pp. 127-142.
- Cary, D. and Del Rosso, K. (2023) Sleight of hand: How China weaponizes software vulnerabilities. Atlantic Council, <https://www.atlanticcouncil.org/in-depth-research-reports/report/sleight-of-hand-how-china-weaponizes-software-vulnerability/> (Accessed 20 January 2026).
- Council on Foreign Relations (n.d.) *Cyber operations tracker* [online]. Council on Foreign Relations, <https://www.cfr.org/cyber-operations/> (Accessed 13 January 2026).
- Fischer, M. (2019) The Concept of Deterrence and Its Applicability in the Cyber Domain. *Connections: The Quarterly Journal*, 18(1/2), pp. 69-92.
- Fischerkeller, M.P. and Harknett, R.J. (2017) Deterrence is Not a Credible Strategy for Cyberspace. *Orbis*, 61(3), pp. 381-393.
- Fischerkeller, M.P., Goldman, E.O. and Harknett, R.J. (2022) *Cyber persistence theory: redefining national security in cyberspace*. New York: Oxford University Press.
- Garfinkel, B. and Dafoe, A. (2019) How does the offense-defense balance scale? *Journal of Strategic Studies*, 42(6), pp. 736-763.
- Healey, J. (2019) The implications of persistent (and permanent) engagement in cyberspace. *Journal of Cybersecurity*, 5(1), pp. 1-15.
- Huntley, W. and Shives, T. (2024) 'The Offense-Defense Balance in Cyberspace', in *23rd European Conference on Cyber Warfare and Security*, Chester, UK: Academic Conferences and Publishing International Limited, pp. 214-221. Available at: <https://papers.academic-conferences.org/index.php/eccws/article/view/2500> (Accessed: 22 January 2026).
- Jervis, R. (1978) Cooperation under the Security Dilemma. *World Politics*, 30(2), pp. 167-214.
- Kello, L. (2017) *The virtual weapon and international order*. New Haven: Yale University Press.
- Libicki, M.C. (2009) *Cyberdeterrence and Cyberwar*. Santa Monica, CA: RAND Corporation.
- Loneragan, E. and Montgomery, M. (2021) What is the Future of Cyber Deterrence? *SAIS Review of International Affairs*, 41(2), pp. 61-73.
- Lim, J.W.Z. and Thing, V.L.L. (2022) 'Towards Effective Cybercrime Intervention', arXiv. Available at: <https://arxiv.org/abs/2211.09524> (Accessed: 22 January 2026).
- Malone, P.J. (2012) *Offense-defense balance in cyberspace: a proposed model*. Master's Thesis. Monterey, CA: Naval Postgraduate School.
- Saltzman, I.Z. (2013) 'Cyber posturing and the offense-defense balance', *Contemporary Security Policy*, 34(1), pp. 40-63. doi:10.1080/13523260.2013.771031.
- Nye, J.S. (2017) Deterrence and Dissuasion in Cyberspace. *International Security*, 41(3), pp. 44-71.
- Slayton, R. (2017) What Is the Cyber Offense-Defense Balance? Conceptions, Causes, and Assessment. *International Security*, 41(3), pp. 72-109.
- Steiger, S. (2022) 'Auf leisen Pfoten oder brüllend laut: Chinesische und russische Cyberangriffe im Vergleich', in Harnisch, S., Zettl, K. and Hansel, M. (eds.) *Asymmetrien in Cyberkonflikten*. Baden-Baden: Nomos, pp. 45-68.
- Soesanto, S. (2022) Cyber Deterrence Revisited. Maxwell Air Force Base, AL: Air University Press, CPP-8.
- Taddeo, M. (2018) The Limits of Deterrence Theory in Cyberspace. *Philosophy & Technology*, 31(3), pp. 339-355.
- Valeriano, B. (2022) 'The failure of offense/defense balance in cyber security', *Cyber Defense Review*, 7(3), pp. 91-103. Available at:

https://cyberdefensereview.army.mil/Portals/6/Documents/2022_summer_cdr/08_Valeriano_CDR_V7N3_Summer_2022.pdf (Accessed: 22 January 2026).

Van de Velde, J. (2023) Cyber Deterrence Is Dead! Long Live “Integrated Deterrence”!. *Joint Force Quarterly*, (109), pp. 41-50.

Van Evera, S. (1998) Offense, Defense, and the Causes of War. *International Security*, 22(4), pp. 5-43.

Zettl-Schabath, K., Bund, J., Müller, M., Borrett, C., Hemmelskamp, J., Alibegovic, A., Bajra, E., Jazxhi, A., Kellenter, E., Sachs, A. and Shelley, C. (2025) Global dataset of cyber incidents (version 1.3.2) [data set]. European Repository of Cyber Incidents. Available at: <https://doi.org/10.5281/zenodo.14965395>.

Appendix 1

	2000 2004	to	2005 to 2009	2010 to 2014	2015 to 2019	2020 to 2024	total ^b
Pearson							
CVEs x CFR	-		0.157	0.910	0.899	-0.770	0.624** (n=20)
CVEs x EuRepoC	-0.321		-0.370	0.083	-0.757	0.878	0.854*** (n=25)
Pearson(log)							
CVEs x CFR	-		-0.073	0.827	0.914	-0.769	0.665** (n=20)
CVEs x EuRepoC	-0.364		-0.340	0.227	-0.718	0.900	0.830*** (n=25)
Spearman's							
CVEs x CFR	-		0.100	0.500	0.900	-0.200	0.677** (n=20)
CVEs x EuRepoC	-0.300		-0.100	-0.200	-0.900	0.900	0.664*** (n=25)
Kendall's tau							
CVEs x CFR	-		0.000	0.400	0.800	-0.200	0.495** (n=20)
CVEs x EuRepoC	-0.200		0.000	-0.200	-0.800	0.800	0.491*** (n=25)
CVEs	7,287		32,320	26,897	61,416	132,633	260,553
CFR	0		25	86	266	520	897
EuRepoC	16		108	547	643	2,008	3,322

^b significance only given for total values
significance level: * = 0.05 ** = 0.01 *** = 0.001

Year-on-year (YoY) values are log-transformed:

