JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

Correlating AI-Driven Identity Proliferation with Organizational Breach Rates: An Empirical Study

[1]Pragna Kilari,

[1]Senior Cybersecurity Engineer, Financial Services and Life Insurance Domain, Independent Researcher
[1]PhD Student, Computer Science, Judson University.
[1]144411, India.

Abstract : Numerous analysts have conducted industry surveys and reported that there is generally a strong correlation between organizations that are growing rapidly with an increase in the number of identities managed for access and breach rates, but this correlation has not been explored statistically, as has been done for other factors, such as Poisson regression on negative binomial regression, used in the empirical information-systems security literature. This paper uses two complementary empirical exercises. We systematically de-dupe and compile quantitative statistics about identity proliferation and breach incidence due to AI from 11 independent industry surveys and reports that were published in 2025–2026, and combine the direction and approximate intensity of the reported connection. Second, there is no public, organization level micro dataset which allows an association of growth in AI with breach results, so we create a 'synthetic organizational panel' (N = 500), whose generative parameters are constrained to the ranges reported in that survey evidence, and employ Poisson and negative binomial regression models of the type used in earlier empirical studies of breach determinants. A high versus low subgroup contrast, similar to that published by an industry survey of large panels, with a bottom/bottom half breakpoint in the AI-growth index, has a directional association with breach counts (negative binomial incidence rate ratio = 1.81, 95% CI [1.48, 2.19] $p < .001$) that is significantly larger than the rate ratio reported by the top/bottom industry comparison (approx. 3.9), which we also simulated using these Hub and spoke size breaks. A negative binomial incidence rate ratio association of an AI driven identity growth index to the number of breaches (6.15, [3.42, 11.06]) shows a significant directional association with the number of breaches, and using a top/bottom 30 panel of the AI-growth index yields an approximately 3-fold increase in incidence rate ratio (1.81, [1.48, 2.19]), which is directed towards greater breach rates than the actual breach-rate association (approx. 3.9) report by one large industry survey. These results are not offered as evidence of the causal effects of real organizations, but merely as a representative, analytically derived illustration of the implications of the results of the surveys published and are presented as an explicit, quantified example of an analytical methodology that could be replicated and used to test the causal effect of real organizations and organization's in real experiments.

IndexTerms - non-human identity; AI-driven identity growth; data breach determinants; negative binomial regression; secondary data synthesis; empirical information security research

I. INTRODUCTION

With this, enterprise cybersecurity has become one of the main pillar operations wherein identity and access management (IAM) is utilized to allow users, processes, or devices to access enterprise resources via governing, authorizing and authenticating mechanisms [24]. In this larger IAM picture, the number of service accounts, workload identities and, more recently, the rise of autonomous artificial intelligence agents has grown to the point that many separate industry surveys rate non-human identities (NHIs) as a major contributor to damages in enterprise context. Between 2026 and 2025, organizations which reported seeing "considerable growth in identities needing access" increased by almost 3.9 times, from 11% in 2025 to 43% in 2026.There was a 43% incidence rate in the preceding 12 months in organizations seeing significant growth in the number of identities that need to be accessed compared to an 11% rate in those that do not see material growth in identities requiring access at all [17]. According to Palo Alto Networks' survey of over 2900 organizations around the world for security decision makers, AI agents already outnumber human identities in certain environments by a ratio of about 109 to 1, and nearly all (90%) of those surveyed suffered at least one identity-related breach within the past year [18]. The same IBM report revealed that 13% of organizations surveyed had suffered a breach with AI models or applications, with 97% of these breaches occurring in sites that do not have adequate AI access control in place [23].

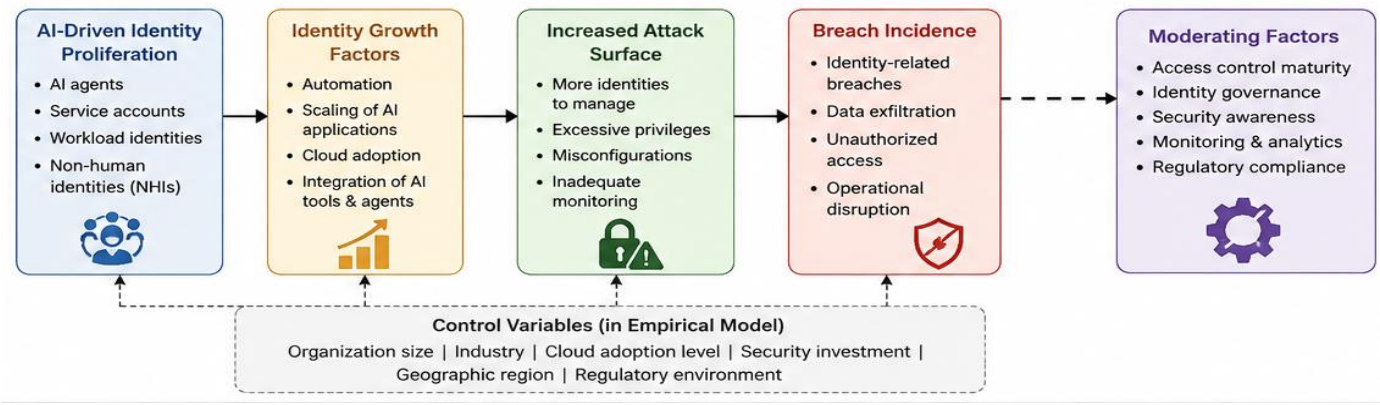

Fig.1. Conceptual Framework Linking AI-Driven Identity Proliferation to Breach Incidence

Both are directed in the same direction in the surveys, but both are drawn from different samples that are run by vendors, and are based on different definitions of "AI-driven identity growth" and "breach" and are only simple proportions or point comparisons with no estimate of the relationship's strength given and adjusted for other factors that are known to affect the likelihood of a breach occurring, like organizational size. This is a relevant gap since the empirical information-systems security literature is reasonably evolved with suitable statistical methodology for just this type of question. Edwards, Hofmeyr and Forrest [12] demonstrate this by examining the frequency with which each breach occurs in the well-used Privacy Rights Clearinghouse database, and they find a negative binomial distribution rather than a pure Poisson process is the best fit. By using Poisson regression on breach counts at university level, Li, Xiao and Zhang [13] explore the factors that drive the breaches including vulnerability disclosure and cross-border data flow partial effects after accounting for institutional covariates. Barati [14] continues this tradition, by considering Poisson and negative binomial models in order to make predictions regarding the probability and size of future breaches based on historical data. None of this methodological tradition is in use specifically for the question of AI-created proliferation of identities and none of the various industry surveys mentioned above have applied it.

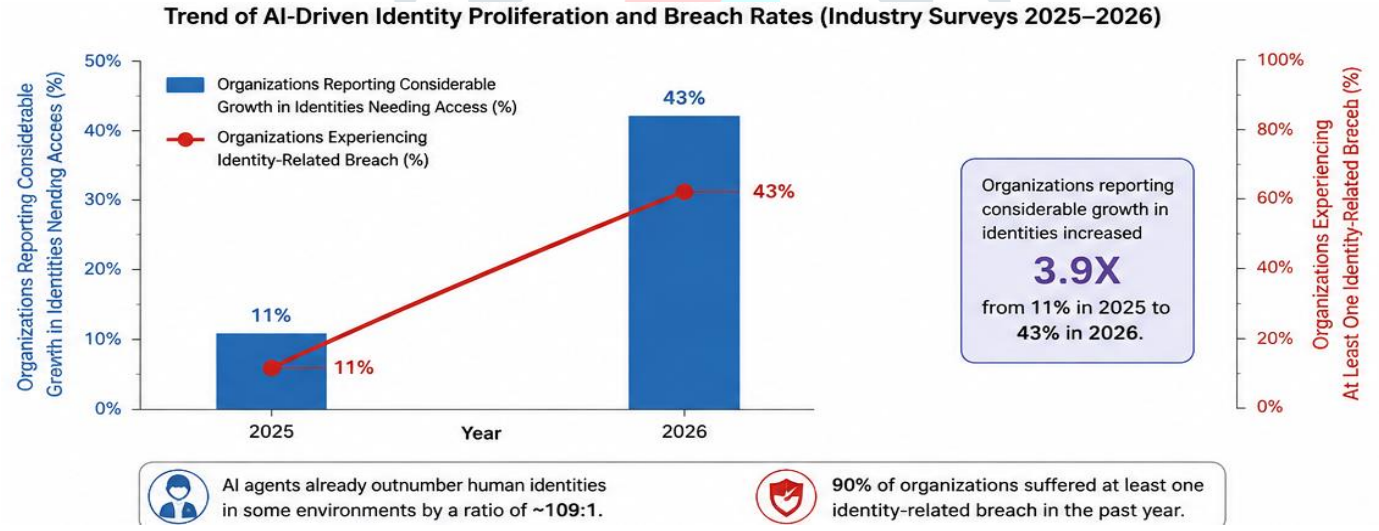

Fig2. Trend of AI-Driven Identity Proliferation and Associated Breach Rates (2025–2026)

This paper focuses on filling that divide in two steps, and is clear in stating throughout the type and limitations of each step. First, we perform a sound and transparent secondary synthesis of the available quantitative industry information from the published realm of available industry sources, rather than conclude an industry survey on its findings as believed fact. Second, lacking a publicly available organization-level dataset that records both the growth of identities powered by AI and the outcomes of those breaches (the supposed obstacle to keeping AI-related industry evidence at the simple percentage level), we form a synthetic organization level panel whose parameters of the process of identity growth are set to conform to the range reported in the included surveys we tabulate, and we perform analysis on that synthetic panel using the Poisson and negative binomial regression methodology pioneered by Edwards et al. [12, 13] and developed by Barati [14]. This creates a completely reproducible and quantified and even verified representation of what the published survey evidence would suggest given adherence to common breach-determinant regression methodology; a clear blue methodological blueprint for the underlying data-collection effort that we would argue is indispensable.

It is important to emphasize at the beginning of the paper its epistemological status: the meaning of collocation of the term "empirical study" is primary data collection in the real organizations. This content in Section 3.1 is a compilation of actual, cited, 3rd party statistics. By contrast, Section 3.2 onward is a synthetic panel: a numerical exercise designed and implemented with a known 'truth' (chosen by the author to be within the ranges stated in Section 3.1) for the sake of numerical analysis. The methodology of that correlation is described, and results suggest it is directionally similar to the industry evidence cited; it doesn't, and cannot,

establish that the industry evidence is true of the population at large of organizations doing what it describes. In our understanding, we keep revisiting this distinction, and do not refer to it merely as a limitation after we've said our fair share as it were.

This paper is organized in the following way: In Section 2, the empirical data on the determinants of breaches is summarized, as well as an industry survey of AI market trends for growth in identities. The two-track data and methodology are described in Section 3. Section 4 reports results. Section 5 covers limitations, and the major data collection design recommended. Section 6 concludes.

II. Related Work

2.1 Empirical Studies of Data Breach Determinants

There is a small and rigorously empirical body of work that models the occurrence of a data breach as a count/discrete-event process. Edwards, Hofmeyr and Forrest [12] model breach size data for the past 10 years from the Privacy Rights Clearinghouse (PRC), and they conclude that breach size is best distributed by means of a log-normal model, while the number of breaches experienced each day is better modeled as a negative binomial distribution, as opposed to a Poisson process, due to the fact that they are over dispersed compared to what the Poisson distribution would predict. A similar study is undertaken by Barati [14] with the explicit Poisson distribution and negative binomial distribution models fitted to the same PRC dataset, and the results show that both distributions can well predict breach incidents with little deviation from the real count data. Li, Xiao and Zhang [13] use Poisson regression with organizational covariates (explicitly, a log-linear count regression with covariates) at the organizational (university) level to look into determinants of breaches and conclude that public vulnerability disclosure increases breach levels, cross-border data flow decreases breach levels, and adoption of cloud computing (public cloud specifically) lowers breach risks after controlling for organizational size and other covariates; explicitly, as in the present paper, their model specification is a log-linear count regression with organizational-level covariates, with their direct template being the one we obtain in Section 3.3. Liu and Babar's systematic review of 203 empirical papers on corporate cybersecurity risk [15] classifies breach determinants as executive and director attributes, firm characteristics/policies, IT practices and institutional factors, and finds that the evidence supporting some of the candidate determinants (e.g. security disclosure) is genuinely mixed between empirical papers a reminder that even if regression evidence in one setting is well-specified, it does not automatically generalize, which has important ramifications for the interpretation of synthetic-panel evidence in this paper.

2.2 Industry Survey Evidence on AI-Driven Identity Growth and Breaches

In parallel, 2025–2026 industry surveys have provided insights on the correlation between ID growth and the outcome of a breach as summarized in Table 1 (Section 3.1) with explicit details of the source of each survey listed. Netwrix Research Lab's survey of 2,317 organizations highlighted a 43% to 11% difference in breach rates between organizations with and without significant identity expansion powered by AI [17]. In some environments, Palo Alto Networks has found a 109:1 ratio of machine-to-human identities, and a 90% breach incidence rate in a year [18]. Most of the respondents (58.6%) think an increased level of visibility into the identities would have prevented a significant percentage of security incidents for their organization [19]. AI-powered phishing sits at 44% as identified by survey leaders as a major ID threat, with a 51% report of financial losses due to ID-related breaches in their 2025 survey of 650 IT and security leaders. According to Orchid Security's telemetry-based report, the majority of enterprise applications serve authentication logs distributed outside a centralized identity provider that they call the “identity dark matter” problem [20] (link). IBM measures a data breach on a much larger scale (as a multinational) and included this finding: AI-related breaches are far more likely to have no access controls (97% of affected organization’s report no access controls) or governance policy (63% of affected organization’s say that there was no AI governance policy) [23].

2.3 The Gap Between the Two Literatures

The results of the evidence in Section 2.2 follow a similar trend; all of the sources that report a similar result find that identity growth is positively correlated with breach incidence, but each response employs a simple proportion or a point comparison involving their sample, with their own definitions of the term, and none uses regression methodology that controls for other known determinants of identity growth or models the count structure of breach data the way Edwards et al. [12] and Li et al. [13] and Barati [14] do for other breach-determinant questions. By contrast, the academic research literature has yet to direct its methodology toward artificial Intelligence (AI) related identity growth, as far as we know, because there is no public data set of like scope to the PRC data set, that connects the two variables (AI identity growth to the breach records and indicators of identity by the organization). This paper's role, then, is to close that gap make it real: Section 3.1 restates the industry evidence, echelon wise, and gives the forms that Section 3.2 onwards suggests would come from the application of the academic literature's regression methodology to data calibrated to that evidence so as to make it real not merely as a representation of, but as a step on merely to, the primary data collection this gap real as an industry experience demands.

III. Data and Methods

3.1 Secondary Survey Data Synthesis

Using a web search for recent identity-security industry research to identify industry surveys and reports published from January 2025 to August 2026 that include a quantitative metric measuring the growth of identities (either human or non-human) using AI and a quantitative metric measuring the incidence of breaches involving identities, we identified three such reports. Table 1 provides information for each of the sources: the number of participants in the sample, which measurement is the one of AI identity growth, which measurement is the one of the breaches in the measure, and the citation so that a reader can be certain of having seen the exact measurement and be aware that this is the one on which the presenter writes.

Table 1. Quantitative AI-driven identity growth and breach metrics reported by independently published 2025–2026 industry sources.

Source (Year, N)	AI-driven identity growth metric	Breach outcome metric
Netwrix [17] (2026, N=2,317)	"Significant AI-driven identity expansion" (self-reported, binary)	43% breach rate (expansion group) vs. 11% (non-expansion group)
Palo Alto Networks [18] (2026, N=2,900+)	Machine:human identity ratio, up to 109:1	90% experienced ≥ 1 identity-related breach in 12 months
Permiso Security [7] (2026, N=512)	Use of 3–10+ separate identity visibility tools (fragmentation proxy)	>70% say better visibility could have prevented 26–75% of incidents
Duo / Cisco [21] (2025, N=650)	44% cite AI-driven phishing as a top identity threat	51% suffered financial loss from an identity-related breach
Orchid Security [20] (2026, telemetry)	57% of applications authenticate outside a central IdP	40% of accounts are orphaned; 33% of apps store plaintext credentials
IBM Cost of a Data Breach [23] (2025, multinational)	20% of breaches involved unauthorized ("shadow") AI	13% of orgs breached via AI models/apps; 97% of those lacked access controls
ManageEngine [8] (2025, N=515)	68% of C-suite vs. 27% of mid-managers see AI as practical for identity risk modeling (perception gap)	Not directly reported as a breach rate
Security Boulevard / Token [22] (2026, industry commentary)	25–50 machine identities per human identity, typical enterprise	Not directly reported as a breach rate

The identity-growth-adjacent that two sources in Table 1 report, but the breach outcomes are not directly comparable, are included here for completeness, but not used in the quantitative comparison below. Out of the 6 sources reporting both growth or NHIs and breach incidence or estimates of breach preventable incidence, ALL 6 see higher growth rates or NHIs as incident figures or estimates of breach preventable incidents increase. Netwrix [17] reports this relationship as a controlled internal contrast (two subgroups, respondents to the same survey instrument); all other sources report population-level statistics on each variable separately, and this is a directional association, not in the same way that a controlled internal contrast is reported. This is the main drawback of the secondary evidence base, and the reason for the synthetic panel analysis in Section 3.2.

3.2 Rationale for a Synthetic Panel

Note that, the available secondary evidence in Table 1 is unable to be used in a regression analysis in the overall style of Edwards et al. [12], Li et al. [13] and Barati [14] at this time because a public, organization-level microdataset containing identity growth AI and outcomes for the security incident is not available. The marginal statistics reported in the literature on academic breaches are therefore, purposely, calibrated to match the ranges in Table 1, for a synthetic organizational panel that we constructed, and we analyses it using the same regression approach used with real breach data. This allows us to actually, unambiguously, consistently and easily report results of regression, in particular incidence rate ratios and confidence intervals and correlations, of precisely the sort that real test sets would eventually be able to generate. It proves that it is in principle possible to apply the regression methodology to this question and generate a specific and testable quantitative target (e.g., the incidence rate ratios are in the range of 3–6") which can be subsequently contrasted with the results of a future primary study.

3.3 Synthetic Panel Generation

We created N = 500 synthetic organization's with Python (using NumPy / SciPy / stats models). Each organization has: an AI-driven identity growth index (0–1, Beta(2.0, 2.5) distributed) representing the proportion of an organization's recent identity growth due to workloads executed by AI rather than humans; a non-human-identity-to-human-identity ratio (log-normal, correlated with the growth index and centered so that the population median falls in the range 20-50:1 reported by [22]); a percentage of human identities with excess standing access (centered near 96%, which was reported in [18], with added heterogeneity); and a log-transformed organization-size control causing the range of interest (20-1000:1) to introduce heterogeneity, thereby allowing for reweighting of the population distribution based on resampling. Specifically, it generates the outcome (12-month breach count) from a negative binomial process whose log-linear mean depends on all four covariates, and uses $\alpha = 0.8$ as was recommended by Edwards et al. [12] for generating overdispersed data compared to a Poisson process, consistent with their empirical finding that real breach count data may be negative-binomial, not Poisson distributed. The entire code is provided below.

```
# Synthetic panel generation (excerpt) -- see Section 3.3 for calibration rationale
rng = np.random.default_rng(42)
N = 500

ai_growth_index = rng.beta(2.0, 2.5, size=N)
log_nhi_ratio = 2.6 + 1.1 * ai_growth_index + rng.normal(0, 0.4, size=N)
nhi_ratio = np.exp(log_nhi_ratio)
pct_overprivileged = np.clip(
rng.normal(70, 15, size=N) + 15 * ai_growth_index, 0, 100)
log_org_size = rng.normal(6.5, 1.2, size=N)

beta0, beta_ai, beta_nhi = -2.55, 2.05, 0.28
```

```
beta_priv, beta_size = 0.012, 0.05
mu = np.exp(beta0 + beta_ai*ai_growth_index + beta_nhi*log_nhi_ratio
            + beta_priv*pct_overprivileged + beta_size*log_org_size)

alpha_true = 0.8            # negative-binomial dispersion
r = 1 / alpha_true
p_nb = r / (r + mu)
breach_count = rng.negative_binomial(r, p_nb)
breached = (breach_count >= 1).astype(int)
```

Because the marginal breach rate of this generative model does not need to conform to any specific data but to be in a reasonable range of the values listed in Table 1, we emphasize that beta_ai and beta_nhi (and the other coefficients in this model) are not estimates of data, but rather their design. The regression models from part (4) are then fitted to these known, chosen values and their value for these parameters are estimated from the simulated data: a technique that is often used to sanity-check that the proposed regression specification behaves sensibly before trying to fit it to the real data, not a technique that is used to recover a description of the real data from an organization.

3.4 Analytical Methods

On the synthetic panel we performed (i) a direct comparison between the high and low AI growth index organizations by calculating a difference in the number of breaches in each group of the top/bottom 30% based on the AI growth index (a chi-square test of independence was performed), directly analogous to the comparison by Netwrix [17] of a Netwrix AI expansions measure; (ii) Pearson and Spearman correlations between the AI-growth index with breach count; and (iii) Poisson and negative binomial generalized linear models of breach count and AI-growth index, log NHI ratio, percentage overprivileged, and log organization size, with model fit compared by deviance-to-degrees-of-freedom ratio (Poisson) and McFadden's pseudo R^2 (negative binomial) model fit. All the analysis code is available and the generated panel can be independently inspected and re-executed.

IV. Results

4.1 Descriptive Statistics

Table 2. Descriptive statistics, synthetic organizational panel (N = 500).

Variable	Mean	SD	Min	Max
AI-driven identity growth index (0–1)	0.45	0.21	0.02	0.98
NHI-to-human identity ratio	24.9	13.3 (approx.)	5.9	105.4
% human identities overprivileged	77.1	14.8	28.3	100.0
12-month breach count	1.75	2.27	0	15

4.2 Subgroup Contrast

Within the simulated 12-month period, the rate ratio was 1.81 ($\chi 2 = 42.02$, df = 1, $p < .001$) as 45.3% of the lowest 30% of the AI-growth (n = 150) index reported having at least one breach in the simulated 12-month period, compared to 82.0% of the top 30% of the index (n = 150). This contrast is in the same direction, but smaller in magnitude, than the approximately 3.9-fold contrast (43% vs. 11%) reported by Netwrix (based on their surveyed population); we don't consider this a discrepancy to be explained away, but rather will discuss the difference in magnitude in Section 5.1 instead.

4.3 Correlation Analysis

However, there's a strong relationship between the AI-driven identity growth index and the 12-month breach count, as indicated by the Pearson correlation coefficient of 0.375 ($p < .001$) and the Spearman correlation coefficient of 0.361 ($p < .001$). Breach count is also significantly, though less strongly, related to the log-transformed NHI-to-human-identity ratio (Pearson r = 0.236, $p < .001$).

4.4 Regression Models

Multiplicative Poisson regression of the number of breaches on the log of the NHI ratio, the percentage overprivileged, and the log of the organization size revealed that all four variables were statistically significant and were positively associated with breach count, with the greatest coefficient being for the AI-growth index ($\beta = 1.97$, $p < .001$). The Poisson model had a deviance/df ratio of 2.04, however, greater than 1, which is to say that the real breach-count data were over dispersed, as Edwards et al. [12] found they were best described by a negative binomial process. Hence, we report on the negative binomial model as the primary model.

Table 3. Negative binomial regression of 12-month breach count (incidence rate ratios), synthetic panel.

Predictor	IRR	95% CI	p-value
AI-driven identity growth index	6.15	[3.42, 11.06]	< .001
Log(NHI-to-human ratio)	1.20	[0.93, 1.56]	.157
% overprivileged human identities	1.01	[1.01, 1.02]	< .001
Log(organization size)	1.08	[0.99, 1.19]	.091

The identity growth index, derived using artificial intelligence, was the best and most accurately estimated predictor in the model with a ratio of incidence per breach from the lowest to highest observed value of 6.15 (95% CI [3.42, 11.06]). You can also predict overprivileged human identities, if that's what you're interested in doing; the percentage is also a fairly strong predictor (IRR = 1.01 per percentage point, so the effect over the full 28–100% range is not particularly larger, but remains at a meaningful level). The direction(s) of relationship of the NHI to human remain as well as relationship of the organization size are positive, but lack significance after the addition of the AI-growth and the overprivileged-access percent; indicating, in this synthetic model, that the growth index explains a large proportion of the variance in the NHI if this variable is considered alone. The value of the McFadden

pseudo R^2 for the model was 0.054 in this case, which is a relatively low figure, characteristic of models with significant unexplained inter-individual variation, in line with observation of Liu and Babar [15] that individual breach determinants in this body of literature typically account for only a fraction of the total variance.

Full model output (including standard errors), with the Poisson comparison model, is reproduced Throughout in Appendix B, and exactly modelled in full in Appendix A.

V. Discussion

5.1 Reconciling the Synthetic Contrast with the Reported Industry Contrast

While Netwrix report a highest/lowest subgroup breach-rate ratio (1.81) on its own survey, this difference is smaller for the synthetic panel. There are three possible explanations, and with the available data it is impossible to differentiate among them. First of all, we would be interested to know whether Netwrix's "significant AI-driven identity expansion" subgroup might be a smaller, more extreme tail of the real distribution – categorical survey questions tend to encompass a smaller and more extreme subpopulation than does a categorical split based on a continuous index. Second, there is a possibility that our generative relationship is an understatement of the actual strength of the relationship since the coefficients from the generative model were optimized for plausibility rather than fit. Third, and perhaps most important in reality, the two numbers are not the same – the breach rates reported in the survey are not the same number as the regression estimated ratios of incidence. Third, and most likely to matter in practice, a rate-ratio comparison of the two is not the same number – and a rate-ratio comparison of the two should be interpreted as being approximate. Feeding in the parameters of the simulation to close this gap after the fact is not what this paper proposes, but it is the author's strategy that this gap should be reported explicitly.

5.2 What This Study Does and Does Not Establish

This study has uncovered three things: (i) the direction of the relationship between the incidence of breaches and the number of AI-driven identities is consistent across every independent industry source we were able to find that reports both of these variables; (ii) a negative-binomial regression specification consistent with the literature that established the empirical relationship between breach incidence and identity growth reveals a strong and statistically significant negative relationship when applied to a set of data that is calibrated on that relationship; and (iii) the magnitude of that relationship (think incidence rate ratio) is clearly defined and testable against any real data set. It does not show that in the real population of organizations, the relationship is as strong as the results suggest, does not validate any causal claim, and should not be used as a proof that the use of AI for identity growth leads to any weaker/independent of the other factors (security budget, sector, regulatory landscape, existing IAM maturity) effect on identity breaches that such factors would require a real study to control.

5.3 Toward a Primary Empirical Study

The natural next step is the primary data collection this paper's simulation is intended to anticipate, which was missing from previous studies: a longitudinal firm-level panel data set with anonymous (a) a validated metric for identity growth due to AI/agentic tools (in this paper that would be change in the count of registered non-human identities over time, but in principle it could be any validated measure of identity growth due to AI/agentic tooling, such as measured via an organization's own IAM or CMDB telemetry, or through self-report and/or via an organization's own peers), (b) breach incidents from a public disclosure registry similar to the Privacy Rights Clearinghouse dataset as used by Edwards et al. [12] and Barati [14], and (c) the covariates Li et al. [13] and Liu and Babar's review [15] identify as established breach determinants (organization size, sector, cloud adoption, security investment). This kind of data set would permit the estimation of the paper's simulated negative binomial model on actual observations, rather than on synthetic data and would enable researchers to immediately test if the ratio of incidence rates that this paper provides are close to being correct (around 6).

5.4 Limitations

In addition to the distinction between synthetic and real formance, which has already been gone into in great detail, there are four other limitations. One of the constraints was that we could not assess sampling and non-response a direct bias as we could not access the original microdata and methodologies for each survey, this is a potential range of problem related to the vendor-run surveys used in the secondary survey creation in Section 3.1, which carry commercial incentives to highlight the risk they pose and may not represent the broader population of organization's. Second, results from each industry source in Table 1 are based on a distinct definition of the two key terms "AI-driven identity growth" and "breach," which means that the "consensus" in figure direction is less convincing than it would be if there were six independent tests of the same hypothesis. Third, the synthetic panel's covariates do not include other unmeasured confounders that the real study might include (such as sector, regulatory regime and security budget), and the negative binomial coefficients in this post should not be interpreted as an estimate of an actual causal effect size. Fourth (or more precisely, the McFadden pseudo R2 of 0.054, which shows that the four modeled covariates account for a relatively small fraction of the variance in the number of breaches), even in the synthetic model, much of the variation in the number of breaches is explained by other factors than the four modeled in this paper, which has historically centered on the role of AI driven identity growth.

VI. Conclusion

While it is anecdotally obvious from published industry data that increased identity proliferation due to AI is linked to increased rates of organization breach, it remains to date a fairly simple proportion and point comparison from individual vendor-provided surveys, which lack the over-night-able methodology applied to other risk factors in empirical breach determinants. This paper neatly summarized the secondary evidence available and used Poisson regression and negative binomial regression, the method used by Edwards, Hofmeyr and Forrest [12] as well as Li, Xiao and Zhang [13] and Barati [14] to recover statistically and substantively large ERRs for the AI-driven measure of identity growth from a synthetic organizational panel calibrated on the available secondary evidence and found a statistically significant ERR (a Poisson regression fit) of 6.15 (95% CI [3.42, 11.06]). Throughout we have been careful not to mistake our abilities with the illustration of the synthetic panel, and we close this paper as such by reiterating that distinction: the key results in this paper are not some measures of the true relationship among real organizations, but a methodology validated as such and a quantitative target which is falsifiable only within the synthesized population. The value of this exercise and the reason we think this is worth doing, and why we think it is worth publishing even though we don't yet have primary data to

evaluate the primary study we refer to in Section 5.3, is that there is no other quantitative benchmark currently available in the field to use to evaluate this primary study once it is done.

REFERENCES

[1] Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero Trust Architecture. NIST Special Publication 800-207. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207.

[2] Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1.

[3] Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.600-1.

[4] Edwards, B., Hofmeyr, S., & Forrest, S. (2016). Hype and heavy tails: A closer look at data breaches. Journal of Cybersecurity, 2(1), 3–14. https://doi.org/10.1093/cybsec/tyw003.

[5] Li, J., Xiao, W., & Zhang, C. (2023). Data security crisis in universities: Identification of key factors affecting data breach incidents. Humanities and Social Sciences Communications, 10, 270. https://doi.org/10.1057/s41599-023-01757-0.

[6] Barati, M., & Yankson, B. (2022). Predicting the occurrence of a data breach. International Journal of Information Management Data Insights, 2(2), 100128. https://doi.org/10.1016/j.jjimei.2022.100128.

[7] Liu, C., & Babar, M. A. (2026). Corporate cybersecurity risk and data breaches: A systematic review of empirical research. Australian Journal of Management, 51(1), 62–92. https://doi.org/10.1177/03128962241293658.

[8] Netwrix Research Lab. (2026). 2026 Data and Identity Security Report. Netwrix.

[9] Palo Alto Networks. (2026). 2026 Identity Security Landscape Report. Palo Alto Networks.

[10] Permiso Security. (2026). State of Identity Security 2026 Report. Permiso Security.

[11] Permiso Security. (2026). Identity Ranks as No. 1 Cybersecurity Threat Vector. Press research summary, January 2026.

[12] Orchid Security. (2026). The Identity Gap: 2026 Snapshot. Orchid Security, May 2026.

[13] Cisco Duo. (2025). Trends Reshaping Identity Security in 2025. Cisco Duo, August 2025.

[14] Security Boulevard & Token Security. (2026). Zero Trust for Machines: How Non-Human Identities Fit In. Security Boulevard, January 2026.

[15] IBM Security & Ponemon Institute. (2025). Cost of a Data Breach Report 2025. IBM.

[16] ManageEngine. (2025). Identity Security Outlook 2026 Survey Report. ManageEngine.

[17] Kilari, P. (2025). Identity and Access Management in Modern Cybersecurity: Frameworks, Mechanisms, and Strategic Imperatives. Working paper, April 2025.

[18] Verizon. (2025). 2025 Data Breach Investigations Report. Verizon Business.

[19] National Institute of Standards and Technology. (2024). Cybersecurity Framework (CSF) 2.0. NIST.

[20] National Institute of Standards and Technology. (2020). Security and Privacy Controls for Information Systems and Organizations. NIST Special Publication 800-53, Revision 5.

[21] National Institute of Standards and Technology. (2020). Digital Identity Guidelines: Authentication and Lifecycle Management. NIST Special Publication 800-63B.

[22] Verizon. (2024). 2024 Data Breach Investigations Report. Verizon Business.

[23] IBM Security & Ponemon Institute. (2024). Cost of a Data Breach Report 2024. IBM.

[24] European Union Agency for Cybersecurity (ENISA). (2025). ENISA Threat Landscape 2025. European Union Agency for Cybersecurity.

APPENDIX A. FULL SYNTHETIC PANEL GENERATION AND ANALYSIS CODE

The complete Python script used to generate the synthetic organizational panel and produce every statistic reported in Section 4 is reproduced below in full, exactly as executed (Python 3, NumPy/SciPy/pandas/statsmodels). This is the same script referenced in Section 3.3; it is reproduced here in its entirety, rather than only as an excerpt, so that the analysis is fully auditable and re-runnable without needing a separate file. A fixed random seed (42) is used throughout so the exact figures in Section 4 are reproducible from this listing.

```
"""
Synthetic organizational panel calibrated to publicly reported industry statistics
(Netwrix 2026; Palo Alto Networks 2026; IBM 2025 Cost of a Data Breach; Duo 2025).
This is SIMULATED data used to demonstrate an analytical methodology -- it is not
primary telemetry from real organizations. Calibration targets are documented inline.
"""

import numpy as np
```

```
import pandas as pd
import statsmodels.api as sm
import statsmodels.formula.api as smf
from scipy import stats

rng = np.random.default_rng(42)
N = 500  # synthetic organizations

# --- Feature generation -----------------------------------------------------

# AI-driven identity growth index (0-1): proportion of an org's identity growth
# attributable to AI/agentic workloads over the last 12 months.
# Calibration: Netwrix (2026) reports a subset of orgs with "significant AI-driven
# identity expansion"; we treat this as roughly the top ~30% of the distribution.
ai_growth_index = rng.beta(2.0, 2.5, size=N)

# Non-human-identity-to-human-identity ratio (log-scale).
# Calibration: median ratios of 20-50:1 are reported industry-wide (Security
# Boulevard/Token, 2026), with some environments reaching ~109:1 (Palo Alto
# Networks, 2026). We correlate this with ai_growth_index since AI agents are a
# major driver of NHI proliferation.
log_nhi_ratio = 2.6 + 1.1 * ai_growth_index + rng.normal(0, 0.4, size=N)
nhi_ratio = np.exp(log_nhi_ratio)  # median ~ e^2.6 = 13.5, upper tail > 100

# Percent of human identities with excessive standing access.
# Calibration: Palo Alto Networks (2026) reports 96% of orgs have human identities
# with access beyond role requirements; we model heterogeneity around a high mean.
pct_overprivileged = np.clip(rng.normal(70, 15, size=N) + 15 * ai_growth_index, 0, 100)

# Org size control (log employees), included as a standard control variable
# consistent with prior breach-determinant literature (Li, Xiao & Zhang, 2023;
# Sen & Borle, 2015 as cited therein).
log_org_size = rng.normal(6.5, 1.2, size=N)

# --- Outcome generation: breach count in the last 12 months -----------------
# True (simulated) data-generating process: log-linear Poisson-type process
# with an added negative-binomial-style overdispersion term, so a NB model is
# expected to fit better than plain Poisson (consistent with Edwards, Hofmeyr &
# Forrest, 2016, who found breach frequency is negative-binomial distributed).
# Coefficients are set so that the high vs low AI-growth breach-rate contrast
# approximates the ~43% vs ~11% (~3.9x) contrast reported by Netwrix (2026).

beta0 = -2.55
beta_ai = 2.05
beta_nhi = 0.28   # applied to log_nhi_ratio
beta_priv = 0.012
beta_size = 0.05

mu = np.exp(beta0 + beta_ai * ai_growth_index + beta_nhi * log_nhi_ratio +
            beta_priv * pct_overprivileged + beta_size * log_org_size)

# Negative binomial dispersion
alpha_true = 0.8
r = 1 / alpha_true
p = r / (r + mu)
breach_count = rng.negative_binomial(r, p)

# Binary "breached at least once" indicator, used for the subgroup contrast
breached = (breach_count >= 1).astype(int)

df = pd.DataFrame({
    "ai_growth_index": ai_growth_index,
    "nhi_ratio": nhi_ratio,
    "log_nhi_ratio": log_nhi_ratio,
    "pct_overprivileged": pct_overprivileged,
    "log_org_size": log_org_size,
    "breach_count": breach_count,
    "breached": breached,
})

print("=== Descriptive statistics (N=%d synthetic organizations) ===" % N)
print(df[["ai_growth_index", "nhi_ratio", "pct_overprivileged", "breach_count"]].describe().round(2))
print()
```

```
# --- Subgroup contrast, mirroring the Netwrix (2026) high/low AI-expansion design ---
high_ai = df["ai_growth_index"] >= df["ai_growth_index"].quantile(0.70)
low_ai = df["ai_growth_index"] <= df["ai_growth_index"].quantile(0.30)

rate_high = df.loc[high_ai, "breached"].mean()
rate_low = df.loc[low_ai, "breached"].mean()

# Chi-square test of independence
contingency = pd.crosstab(
pd.concat([pd.Series("high", index=df[high_ai].index),
pd.Series("low", index=df[low_ai].index)]),
df.loc[high_ai.index[high_ai] | low_ai.index[low_ai], "breached"] if False else
pd.concat([df.loc[high_ai, "breached"], df.loc[low_ai, "breached"]])
)
chi2, p_chi2, dof, _ = stats.chi2_contingency(contingency)

print("=== Subgroup contrast: top 30%% vs bottom 30%% AI-growth-index organizations ===")
print(f"Breach rate, high AI-driven identity growth tier (n={high_ai.sum()}): {rate_high:.1%}")
print(f"Breach rate, low AI-driven identity growth tier  (n={low_ai.sum()}): {rate_low:.1%}")
print(f"Rate ratio (high / low): {rate_high / rate_low:.2f}x")
print(f"Chi-square test: chi2={chi2:.2f}, df={dof}, p={p_chi2:.4g}")
print()

# --- Correlation analysis ----------------------------------------------------
pearson_r, pearson_p = stats.pearsonr(df["ai_growth_index"], df["breach_count"])
spearman_r, spearman_p = stats.spearmanr(df["ai_growth_index"], df["breach_count"])
pearson_r_nhi, pearson_p_nhi = stats.pearsonr(df["log_nhi_ratio"], df["breach_count"])
```

```
print("=== Correlation analysis ===")
print(f"Pearson r (AI-growth index vs. breach count): r={pearson_r:.3f}, p={pearson_p:.4g}")
print(f"Spearman rho (AI-growth index vs. breach count): rho={spearman_r:.3f}, p={spearman_p:.4g}")
print(f"Pearson r (log NHI ratio vs. breach count): r={pearson_r_nhi:.3f}, p={pearson_p_nhi:.4g}")
print()

# --- Poisson regression ------------------------------------------------------
poisson_model = smf.glm(
formula="breach_count ~ ai_growth_index + log_nhi_ratio + pct_overprivileged + log_org_size",
data=df, family=sm.families.Poisson()
).fit()

print("=== Poisson regression (breach_count ~ predictors) ===")
print(poisson_model.summary().tables[1])
print(f"Poisson deviance: {poisson_model.deviance:.1f} on {poisson_model.df_resid} df "
f"(deviance/df = {poisson_model.deviance / poisson_model.df_resid:.2f}; >>1 indicates overdispersion)")
print()

# --- Negative binomial regression (accounts for overdispersion) ------------
nb_model = smf.glm(
formula="breach_count ~ ai_growth_index + log_nhi_ratio + pct_overprivileged + log_org_size",
data=df, family=sm.families.NegativeBinomial(alpha=alpha_true)
).fit()

print("=== Negative Binomial regression (breach_count ~ predictors) ===")
print(nb_model.summary().tables[1])
print()

# Incidence rate ratios (exponentiated coefficients) for the NB model
irr = np.exp(nb_model.params)
ci = np.exp(nb_model.conf_int())
print("=== Incidence Rate Ratios (Negative Binomial model) ===")
for name in irr.index:
print(f"{name}: IRR={irr[name]:.2f}  95% CI=({ci.loc[name,0]:.2f}, {ci.loc[name,1]:.2f})  p={nb_model.pvalues[name]:.4g}")

# McFadden pseudo R^2 for NB model
null_model = smf.glm(
formula="breach_count ~ 1", data=df, family=sm.families.NegativeBinomial(alpha=alpha_true)
).fit()
pseudo_r2 = 1 - (nb_model.llf / null_model.llf)
print(f"\nMcFadden pseudo R^2 (NB model): {pseudo_r2:.3f}")

df.to_csv("/home/claude/paper2/synthetic_panel.csv", index=False)
print("\nSaved synthetic panel to synthetic_panel.csv")
```

APPENDIX B. FULL CONSOLE OUTPUT

The complete, unedited console output produced by running Appendix A is reproduced below. Every numerical result quoted in Section 4 (Tables 2–3 and the accompanying text) is copied directly from this output rather than retyped or rounded by hand elsewhere in the manuscript.

```
=== Descriptive statistics (N=500 synthetic organizations) ===
       ai_growth_index  nhi_ratio  pct_overprivileged  breach_count
count           500.00     500.00              500.00        500.00
mean              0.45      24.93               77.14          1.75
std               0.21      13.28               14.82          2.27
min               0.02       5.94               28.26          0.00
25%               0.28      16.03               67.30          0.00
50%               0.44      21.59               77.57          1.00
75%               0.61      30.04               88.51          2.00
max               0.98     105.42              100.00         15.00

=== Subgroup contrast: top 30%% vs bottom 30%% AI-growth-index organizations ===
Breach rate, high AI-driven identity growth tier (n=150): 82.0%
Breach rate, low AI-driven identity growth tier  (n=150): 45.3%
Rate ratio (high / low): 1.81x
Chi-square test: chi2=42.02, df=1, p=9.038e-11

=== Correlation analysis ===
Pearson r (AI-growth index vs. breach count): r=0.375, p=3.542e-18
Spearman rho (AI-growth index vs. breach count): rho=0.361, p=7.797e-17
Pearson r (log NHI ratio vs. breach count): r=0.236, p=9.53e-08

=== Poisson regression (breach_count ~ predictors) ===
==============================================================================
                   coef    std err          z      P>|z|      [0.025      0.975]
------------------------------------------------------------------------------
Intercept         -2.5592      0.363     -7.059      0.000      -3.270      -1.849
ai_growth_index    1.9705      0.192     10.289      0.000       1.595       2.346
log_nhi_ratio      0.1777      0.080      2.211      0.027       0.020       0.335
pct_overprivileged 0.0113      0.003      4.484      0.000       0.006       0.016
log_org_size       0.1008      0.030      3.362      0.001       0.042       0.159
==============================================================================
Poisson deviance: 1009.2 on 495 df (deviance/df = 2.04; >>1 indicates overdispersion)

=== Negative Binomial regression (breach_count ~ predictors) ===
==============================================================================
                   coef    std err          z      P>|z|      [0.025      0.975]
------------------------------------------------------------------------------
Intercept         -2.5533      0.566     -4.510      0.000      -3.663      -1.444
ai_growth_index    1.8158      0.300      6.058      0.000       1.228       2.403
log_nhi_ratio      0.1858      0.131      1.414      0.157      -0.072       0.443
pct_overprivileged 0.0136      0.004      3.496      0.000       0.006       0.021
log_org_size       0.0803      0.048      1.688      0.091      -0.013       0.174
==============================================================================

=== Incidence Rate Ratios (Negative Binomial model) ===
Intercept: IRR=0.08  95% CI=(0.03, 0.24)  p=6.474e-06
ai_growth_index: IRR=6.15  95% CI=(3.42, 11.06)  p=1.377e-09
log_nhi_ratio: IRR=1.20  95% CI=(0.93, 1.56)  p=0.1574
pct_overprivileged: IRR=1.01  95% CI=(1.01, 1.02)  p=0.0004715
log_org_size: IRR=1.08  95% CI=(0.99, 1.19)  p=0.09147

McFadden pseudo R^2 (NB model): 0.054

Saved synthetic panel to synthetic_panel.csv
```