Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

AI-Enabled Cryptocurrency Scams and Investor Vulnerability: Evidence on Scam Sophistication, Adoption, Financial Access and Losses

Sahila Chellamuthu

sahiladr@srmist.edu.in

SRM Institute of Science and Technology

Mcxin Tee

INTI International University

Research Article

Keywords: cryptocurrency scams, generative AI, investor protection, digital financial literacy, cybercrime, fraud losses

Posted Date: November 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11383134/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

AI-Enabled Cryptocurrency Scams and Investor Vulnerability: Evidence on Scam Sophistication, Adoption, Financial Access and Losses

Abstract

Cryptocurrency fraud is among the largest categories of financial crime, and generative AI may be raising the scale and persuasiveness of scams. We link AI-enabled scam sophistication, cryptocurrency adoption and financial access to investor losses, using public blockchain-analytics, survey and World Bank data. Scam operations linked to AI tooling earned about 4.5 times more than other operations, a descriptive comparison. Across eleven countries, loss prevalence ranges from 1% to 36%, and exposure alone does not explain the gap. Crypto adoption is unrelated to losses (Spearman $\rho = 0.21$; interval -0.46 to 0.74). Account ownership is negatively associated with losses ($\rho = -0.74$), but is confounded with income and region. We discuss implications for regulators, exchanges and AI providers.

Keywords: cryptocurrency scams; generative AI; investor protection; digital financial literacy; cybercrime; fraud losses

1. Introduction

Cryptocurrency is now a routine channel for financial crime against ordinary investors. Chainalysis estimates that crypto scams took at least US$14 billion on-chain in 2025, up from the US$9.9 billion first reported for 2024, and that the final tally could exceed US$17 billion as further data arrive (Chainalysis 2026, as reported by Cointelegraph, The Next Web and Forklog; see Appendix A). Over the same period the composition of scams changed. Chainalysis reports that impersonation scams grew by roughly 1,400% year on year, and that the average crypto payment across scam categories rose from US$782 to US$2,764. The

identity-verification firm Sumsub reports that deepfakes now account for about 11% of global fraudulent activity and that "sophisticated fraud" rose by 180% in 2025 (Sumsub 2026).

Survey evidence points the same way from the victim's side. The Global Anti-Scam Alliance (GASA) *Global State of Scams 2025* report finds that 57% of adults across 42 markets experienced a scam in the previous 12 months and that 23% lost money, rising to 41% in South America and Africa (GASA and Feedzai 2025). The same survey finds that 73% of adults say they are confident they can recognize a scam. High confidence alongside high loss is the puzzle that motivates this paper: individual awareness does not appear to be sufficient protection when deception is cheap, personalized and convincing.

Three strands of research bear on this problem, but they have largely developed separately. The first studies the economics of deception and the capability of generative models to produce convincing synthetic media (Akerlof and Shiller 2015; Chesney and Citron 2019; Mirsky and Lee 2021; Cross 2022). The second studies crypto-asset markets and crypto-enabled crime, including market manipulation and the flow of scam proceeds through exchanges (Gandal et al. 2018; Griffin and Mei 2024). The third studies the institutional environment: investor-protection law, financial literacy and the regulation of the crypto economy (La Porta et al. 2006; Lusardi and Mitchell 2014; Cumming et al. 2019; Ebekozien et al. 2024). To our knowledge, no peer-reviewed study connects these strands by asking how AI-enabled sophistication interacts with country-level adoption, literacy and regulation to shape investor losses. The closest work we found is practitioner research that documents AI-enabled typologies in crypto crime and quantifies the productivity of AI-linked scam operations, without testing variation across countries (Elliptic 2024; Chainalysis 2026). This statement rests on a targeted search of open web and preprint sources carried out on 3 October 2026, described in Section 3.5, and not on a systematic search of bibliographic databases; we return to this in Section 8.

This paper makes three contributions. First, it develops a framework and three hypotheses (H1 to H3) that can be assessed with public operation-level and cross-country data. Second, it assembles and documents, from public sources, a transparent evidence base that other researchers can reproduce, and it is explicit about what each source can and cannot measure. Third, it reports descriptive evidence on the operation-level premium associated with AI tooling and on the cross-country distribution of losses, and pre-specifies an estimation strategy appropriate for the small number of countries for which comparable survey data exist. Section 2 develops the hypotheses. Section 3 describes data and measurement. Section 4 sets out the empirical strategy. Section 5 presents descriptive evidence, Section 6 the estimation results, and Sections 7 to 9 discuss implications, limitations and conclusions.

2. Theoretical background and hypotheses

2.1 AI-enabled sophistication and losses

Akerlof and Shiller (2015) argue that in competitive markets profit-seeking sellers have an incentive to exploit the psychological weaknesses of buyers, so that deception can persist as an equilibrium rather than being competed away. Generative AI lowers the cost of the inputs to deception: fluent text in many languages, cloned voices, and synthetic faces and video (Chesney and Citron 2019; Mirsky and Lee 2021). Cross (2022) shows why this matters for victims: deepfake technology undermines standard checks, such as a reverse image search, that prevention messaging has advised people to use. We interpret AI-enabled sophistication as the combination of greater personalization, greater credibility signals and greater operational scale. Each raises the expected loss per contact and the number of contacts a given operator can sustain.

H1. Greater AI-enabled scam sophistication is associated with higher cryptocurrency scam revenue (assessed at the operation level).

2.2 Cryptocurrency adoption and exposure

Routine-activity theory holds that crime requires a motivated offender, a suitable target and the absence of a capable guardian (Cohen and Felson 1979). Higher crypto adoption enlarges the pool of suitable targets, because victims who hold or can acquire crypto can pay scammers in an irreversible, borderless form, and it may enlarge the pool of inexperienced first-time users. Griffin and Mei (2024) trace crypto flows into suspicious exchange accounts at around US$27.8 billion a year over 2021 to 2023, which illustrates how readily scam proceeds can use the same infrastructure as legitimate users. Adoption may also bring guardians, such as exchange compliance functions, so the sign of the net effect is an empirical question; we expect the exposure effect to dominate.

H2. Cryptocurrency adoption is positively associated with the prevalence of cryptocurrency scams.

2.3 Financial literacy and financial access

Financial literacy is associated with better financial decisions, and its absence has real costs (Lusardi and Mitchell 2014). The relevant skill for crypto scams, however, is not only numeracy but the ability to verify identity, assess unrealistic returns and resist social engineering in digital channels. GASA's finding that most adults are confident in their ability to spot scams while a large minority lose money suggests that perceived and actual competence diverge. We therefore expect measured capability, and the financial access that goes with it, rather than self-assessed confidence, to be protective.

H3. Greater digital financial capability and financial access are associated with lower scam losses.

2.4 Extensions that we do not test

Three further questions follow from the framework but cannot be tested with public data, and we do not state them as hypotheses. The first is whether the efficiency gain from AI is larger where crypto adoption is higher, so that sophistication and adoption interact. The second is whether regulatory strength weakens the link between sophistication and losses. La Porta et al. (2006) find that disclosure and liability rules are associated with better-developed stock markets, and Cumming et al. (2019) analyse the regulation of the crypto economy. For scams, the relevant levers include licensing of exchanges, know-your-customer and travel-rule enforcement, and powers to freeze funds, which act on the cash-out stage. The third is whether AI-enabled scams are more cross-border, since AI lowers the cost of translation, voice and culturally specific scripts, and Sumsub notes that payment details stolen in one country can be used in another where protections are weaker (Sumsub 2026). Each needs a country-level sophistication measure, a regulatory index or case-level operation data that are not publicly available.

3. Data and measurement

3.1 Sources

The evidence base combines public sources, summarized in Table 1 and Appendix A. We do not use proprietary data. A further source is the set of GASA 2026 country reports. GASA publishes a free version of each on its public report pages (GASA 2026a). We downloaded those for Brazil, Canada, Germany, Japan, Kenya, Nigeria, Portugal, South Africa, Spain, Taiwan and the United States, together with the Africa and Europe regional editions, on 3 October 2026. The full reports are for members and we did not use them. From each free version we took the headline results (share seeing a scam, interacting, losing money or data, reporting) and, where the report has one, the share of victims who say the scammer received

their money through a cryptocurrency exchange or digital currency (GASA question Q13). The Japan free version has no payment-channel page. GASA describes the 2026 programme as 58,900 respondents in 42 countries, so our eleven countries are the part of that wave whose free reports we could use. Where a figure comes from a primary report that we could not obtain directly, we cite the report and name the secondary outlet that quoted it (Appendix A), and the reader should treat such figures as reported rather than independently verified.

Table 1. Constructs, measures and sources

Construct	Measure	Source	Main limitation
Investor losses	Share of adults who lost money in last 12 months; average loss per victim (US$); estimated loss as % of GDP	GASA and Feedzai (2025), survey of 46,000 adults in 42 markets, fieldwork 7 to 20 March 2025	Online-panel survey; method differs from earlier GASA waves; top 2% to 5% of reported amounts trimmed; country values partly published only as charts
AI-enabled sophistication	Operation-level: on-chain links to AI tooling vendors. Region-level: deepfake and sophisticated-fraud rates	Chainalysis (2026), via press coverage; Sumsub (2026)	No country-level measure in hand; Sumsub is a vendor and measures identity-verification fraud, not crypto scams
Crypto adoption	Global Crypto Adoption Index: rank and 0 to 1 score from four value-received sub-indices	Chainalysis (2025)	Based on web-traffic attribution; VPNs can misattribute activity; 2025 method changed (retail DeFi removed, institutional sub-index added)
Digital financial capability (proxy)	Adults with an account; adults who made or received a digital payment; internet use (%, 2024)	Klapper et al. (2025), Global Findex 2025, Table B.1	Measures access and use, not knowledge; digital-payment cells blank for most high-income economies

3.2 Loss measures

Our primary outcome is the share of adults in a market who lost money to a scam in the previous 12 months, measured by GASA's question on total money lost, regardless of whether it was later recovered. The 2026 country pages report a slightly different quantity, the share of respondents who experienced financial or data loss, which is broader than money lost and is not specific to crypto. We use it as the outcome for the country analysis in Section 6 and label it "loss prevalence". Secondary outcomes are average loss per victim and estimated loss as a

share of GDP. GASA warns that its 2025 statistical approach differs from earlier reports and that historical comparisons should be treated with caution. The sensitivity is visible in the published numbers: GASA's 2023 report put Kenya's loss at nearly 4.5% of GDP, whereas the 2025 report lists 10.9%. We therefore do not splice waves into a single series. Reported-loss statistics in the 2021 and 2022 editions, which are compiled from national reports rather than survey responses, are treated as a separate robustness sample. The public report pages give a further sign of the sensitivity: the 2024 Japan page (GASA 2024) reports that 39% of participants lost money, whereas the 2026 Japan page reports that 1% experienced financial or data loss. We cannot tell from the public text how much of the difference reflects the change in question wording or sample, and this is another reason not to compare across waves.

3.3 Adoption and capability measures

The Chainalysis index ranks countries on four sub-indices of crypto value received: at centralized services, at centralized services in retail-sized transfers (under US$10,000), in DeFi protocols, and (new in 2025) institutional transfers above US$1 million. Each ranking is weighted for population and purchasing power, the geometric mean is taken and the result is normalized to a 0 to 1 score (Chainalysis 2025). The report describes 151 countries; the full-index table that we could extract lists countries through rank 131, and we use ranks as published. Because the index is purchasing-power weighted, it measures grassroots intensity rather than total volume. For digital financial capability we use Findex 2025 indicators for 2024 (Klapper et al. 2025). These are proxies for access and use, and we label them as such.

3.4 Sophistication

The cleanest sophistication evidence is operation-level: Chainalysis compares scam operations with observable on-chain links to AI tooling vendors against those without. For a country-level analysis we need a market-level indicator, and plan to use the deepfake or sophisticated-fraud

rates in Sumsub's identity-fraud reports and a text-based index built from public scam reports. Neither is yet integrated.

3.5 Literature search

We ran a targeted, non-systematic search on 3 October 2026 to check the claim of a gap made in Section 1. We searched the open web with two query sets, one combining generative AI, cryptocurrency scams, investor losses, cross-country comparison and regulation, and one combining deepfakes, AI-enabled crypto fraud, pig-butchering and victim losses. The searches returned industry and regulator material (including a horizon-scanning report by a blockchain-analytics firm, vendor and press coverage of the Chainalysis 2026 report, and securities-regulator warnings) and no peer-reviewed study that tests how AI-enabled sophistication, adoption, literacy and regulation jointly relate to losses across countries. Because the search was limited in scope and did not cover Scopus or Web of Science, we treat the absence of such a study as provisional. A systematic search with a PRISMA 2020 flow diagram (Page et al. 2021) is the appropriate way to confirm it, and we list it as a limitation.

4. Empirical strategy

4.1 Baseline specification

For market i in wave t, our baseline specification is

$$Loss_it = \alpha + \beta 1\ Adopt_it + \beta 2\ Access_it + \gamma' X_it + \tau_t + \varepsilon_it,$$

where Loss is the share of adults losing money (or its logit or log transform), Adopt the adoption index, Access the financial-access proxy, X a vector of controls (log GDP per capita, internet use, region), and τ_t wave effects when more than one comparable wave is available. H2 tests $\beta 1 > 0$ and H3 tests $\beta 2 < 0$. H1 concerns sophistication, for which there is no country-

level measure, and is assessed by comparing AI-linked and other scam operations on revenue and activity measures (Section 5.1).

Three features of the data shape the estimation. First, the number of markets is small. With 42 markets and four to six regressors, a conventional power calculation ($\alpha = 0.05$, power 0.80) implies that we can detect only effects with a partial R^2 of roughly 0.17 to 0.19 (Cohen's f^2 of about 0.20 to 0.24) for a single predictor. We therefore report confidence intervals and effect sizes, not just significance, and treat null results as uninformative about small effects. Second, there are few regional clusters, so we use heteroskedasticity-robust (HC3) errors, wild bootstrap inference with regional clusters, and leave-one-country-out influence checks. Third, we test a small number of directional hypotheses on one dataset and adjust for multiplicity with the Holm procedure (Section 4.2).

4.2 Estimation with the data in hand

The country data that we could assemble cover nine to eleven markets in the 2026 wave and 18 to 19 markets in the 2025 extreme-group sample, which is too few for the regression in Section 4.1 to be informative. We therefore estimate Spearman rank correlations and report three kinds of uncertainty. First, an exact permutation p-value (all 8! or 9! reorderings of one variable). Second, a percentile bootstrap interval from 5,000 resamples of countries. Third, the range of the coefficient when each country is dropped in turn. The confirmatory family contains three tests on the 2026 wave: H2 (adoption and loss prevalence), H3 (account ownership and loss prevalence), and a crypto-specific version of H2 (adoption and the approximate share of adults who lost money through a crypto exchange). The first two were specified before we obtained the payment-channel data. The third was added after we saw that the free reports contain them, and we flag it as post hoc. We adjust the three p-values with the Holm procedure. All other correlations are labelled exploratory and shown without adjustment. Where there are

more than nine countries (and for the 2025 sample) the p-values come from 400,000 or 200,000 random permutations. The 2025 sample is selected on the outcome and its results are descriptive. With nine to eleven countries, a rank correlation must exceed about 0.62 to 0.68 in absolute value to reach $p < 0.05$, so the analysis can only detect strong associations.

The design is observational. Even a significant $\beta 1$ would show association, not causation: operators that adopt AI tooling may differ in skill or targeting, countries with more sophisticated attackers may differ in unmeasured ways, and survey losses are self-reported. We use the language of association throughout, and the robustness section will report specifications that drop each control group in turn.

5. Descriptive evidence

5.1 Scale, composition and the AI premium

The aggregate numbers in Section 1 show a market in which both volume and average payment size grew. The operation-level comparison in Figure 1 is the most direct available evidence on sophistication. Chainalysis finds that operations with observable on-chain links to AI tooling vendors extracted on average US$3.2 million, against US$719,000 for those without, a ratio of about 4.5. They also received 35.1 transfers per day against 3.89, and earned US$4,838 per day against US$518 (Chainalysis 2026, as reported).

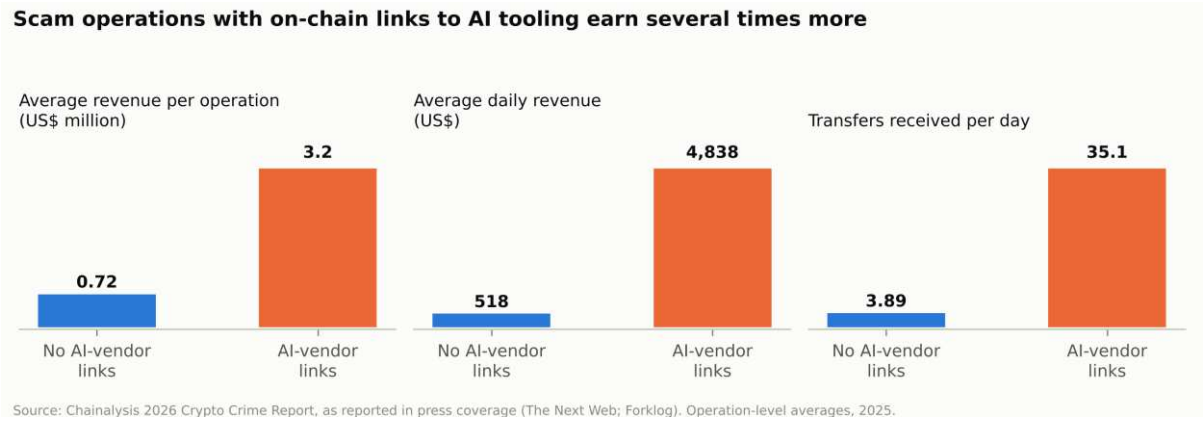


Figure 1. Scam operations with and without on-chain links to AI tooling vendors, 2025. Values as reported in press coverage of the Chainalysis 2026 Crypto Crime Report. The comparison is descriptive: operators that use AI tooling may differ in skill or scale for other reasons.

Two readings are consistent with these numbers. On the first, AI tools raise revenue by making each approach more persuasive and by letting an operation contact many more people, which is the mechanism behind H1. On the second, more capable and better-resourced operators are the ones that buy AI tooling, so the premium partly reflects selection. The data in hand cannot separate the two, but the reported averages point towards volume. AI-linked operations receive about nine times as many transfers per day (35.1 against 3.89), and the implied revenue per transfer is similar in the two groups (about US$138 against US$133, computed from the reported daily revenue and transfer counts), so the gain appears to come mainly from reach rather than from larger payments. Any claim that AI *causes* the premium would need within-operation data before and after adoption.

5.2 Losses across regions and countries

Figure 2 shows the two regional series that GASA publishes. The share of adults losing money is 41% in both South America and Africa, against 19% in Asia and 20% in Europe. Crypto exchanges are the payment channel in 11% of victim payments worldwide, but the regional shares do not track loss prevalence: South America has the lowest crypto share (4%), while

Africa (13%) and the Middle East (14%) are at the top. With seven regions we draw no statistical conclusion, but the pattern warns against reading regional loss differences as a crypto effect. Wire or bank transfer (29%) and credit-card payment (18%) remain the most common channels globally (GASA and Feedzai 2025).

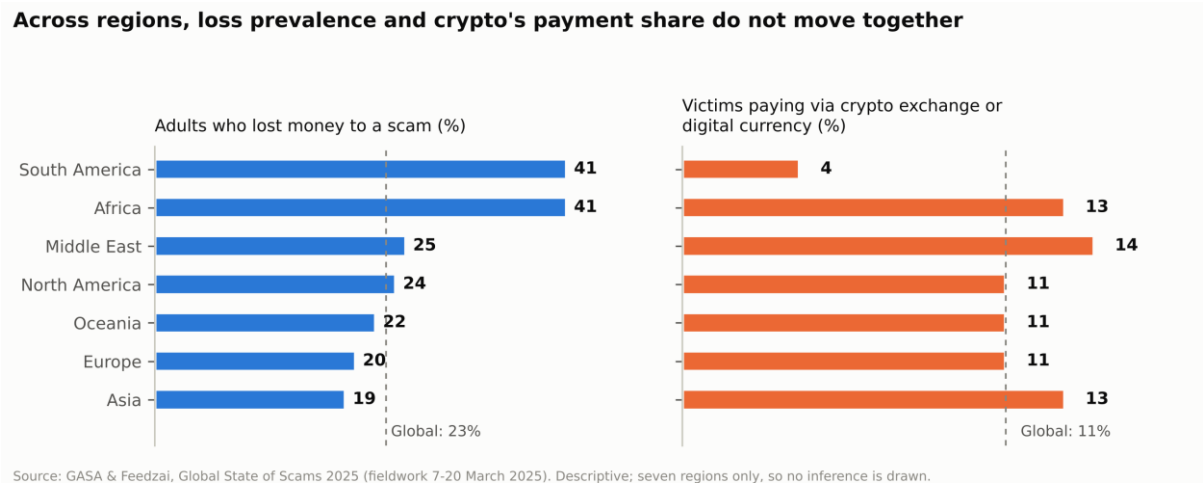


Figure 2. Loss prevalence and crypto payment share by region, GASA Global State of Scams 2025. Regional bases are as reported by GASA; the global values are 23% and 11%.

Figure 3 and Table 2 turn to countries. GASA publishes the ten highest and nine lowest loss burdens as a share of GDP, and we attach each country's Chainalysis adoption rank and Findex indicators. Estimated losses range from 11.3% of GDP in Nigeria and 10.9% in Kenya to 0.02% in Sweden. High adoption does not sort countries cleanly: India (adoption rank 1), Pakistan (3), Brazil (5) and Nigeria (6) are among the highest-loss markets, while the United States (2) and Indonesia (7) are among the lowest-loss ones. The median adoption rank is 14.5 in the high-loss group and 48.5 in the low-loss group (Taiwan excluded because it is not in the Chainalysis table), a difference in the direction of H2. This comparison cannot be treated as a test, because the countries were selected for being at the extremes of the outcome, which inflates differences, and the groups are small.

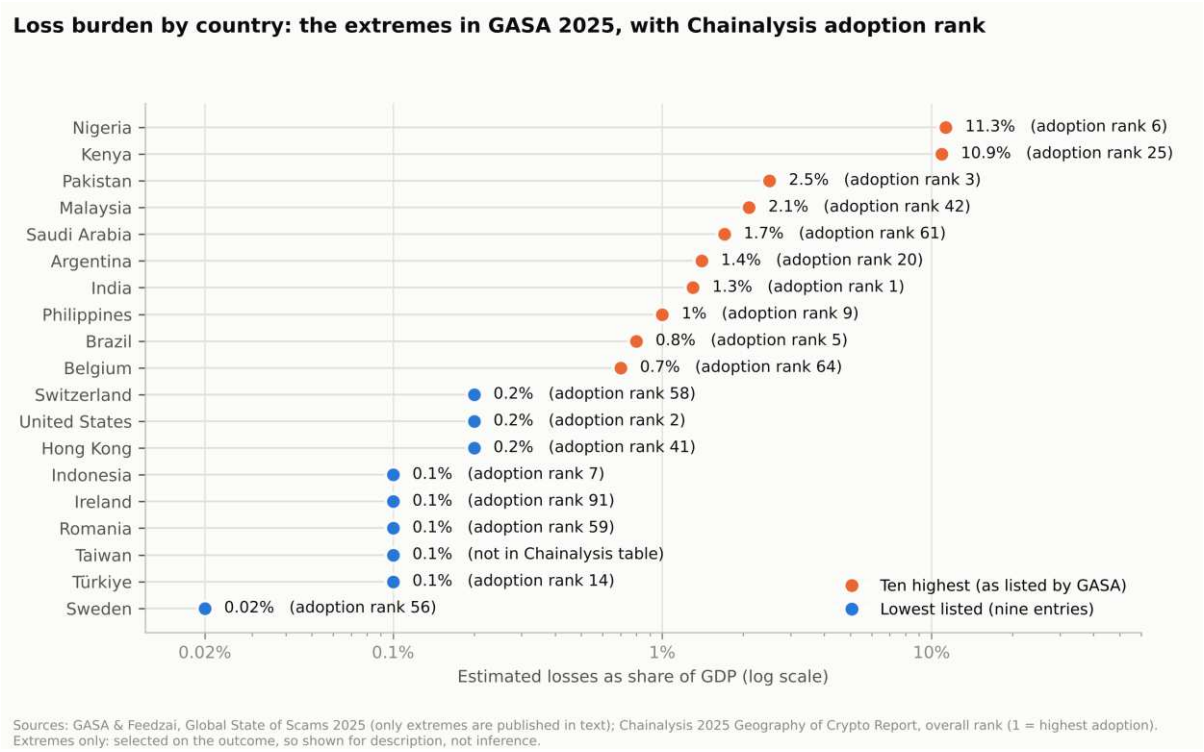


Figure 3. Estimated loss as a share of GDP for the ten highest and nine lowest markets listed by GASA (2025), with Chainalysis overall adoption rank (1 = highest). Log scale. Taiwan does not appear in the extracted Chainalysis table.

Table 2. Markets at the extremes of GASA's 2025 loss-to-GDP ranking, with adoption and capability indicators

Market	Loss, % of GDP (GASA 2025)	GASA group	Chainalysis adoption rank	Findex: has account (%)	Findex: digital payment (%)
Nigeria	11.3	Highest 10	6	63	54
Kenya	10.9	Highest 10	25	90	89
Pakistan	2.5	Highest 10	3	27	25
Malaysia	2.1	Highest 10	42	89	77
Saudi Arabia	1.7	Highest 10	61	79	76
Argentina	1.4	Highest 10	20	82	72
India	1.3	Highest 10	1	89	48
Philippines	1.0	Highest 10	9	50	40
Brazil	0.8	Highest 10	5	86	77
Belgium	0.7	Highest 10	64	98	n/a
Switzerland	0.2	Lowest listed	58	98	n/a
United States	0.2	Lowest listed	2	97	n/a
Hong Kong SAR	0.2	Lowest listed	41	97	n/a

Market	Loss, % of GDP (GASA 2025)	GASA group	Chainalysis adoption rank	Findex: has account (%)	Findex: digital payment (%)
Indonesia	0.1	Lowest listed	7	56	43
Ireland	0.1	Lowest listed	91	98	n/a
Romania	0.1	Lowest listed	59	71	64
Taiwan	0.1	Lowest listed	not listed	96	n/a
Türkiye	0.1	Lowest listed	14	82	71
Sweden	0.02	Lowest listed	56	99	n/a

Notes: GASA 2025 lists nine entries in its lowest group. Findex digital-payment values are blank for most high-income economies because those data were not collected for 2024 (n/a). Sources: GASA and Feedzai (2025); Chainalysis (2025); Klapper et al. (2025), Table B.1.

Average losses per victim tell a complementary story (Figure 4). They are highest in Austria (US$2,885), Saudi Arabia (US$2,511) and Australia (US$2,384), and lowest in Romania (US$113), Indonesia (US$115) and Türkiye (US$116). Wealthier markets therefore tend to have large losses per victim, while lower-income markets have a larger share of the population losing smaller amounts, so that the loss burden relative to GDP is highest in some lower-income markets. Loss as a share of GDP is a ratio of two quantities that differ across countries for reasons unrelated to scams, and we report it alongside prevalence for that reason.



Figure 4. Average loss among victims, ten highest and ten lowest markets, GASA Global State of Scams 2025 (victims who lost money, n = 10,732). Markets with insufficient bases are omitted by GASA.

5.3 What the descriptive evidence does and does not show

The evidence is consistent with H1 at the operation level and weakly in the direction of H2 at the country level. It gives no country-level information on H1, because no country-level sophistication measure is available. On H3, Table 2 shows that account ownership varies widely at both ends of the loss ranking (from 27% in Pakistan to 90% in Kenya in the high-loss group, and from 56% in Indonesia to 99% in Sweden in the low-loss group), and that digital-payment use is high in several high-loss markets (Kenya 89%, Brazil 77%, Malaysia 77%). The indicators do not line up with the loss ranking in any obvious way. High digital-payment use in some high-loss markets is compatible with the view that digital access also raises exposure, but it mainly underlines that access and use are poor proxies for the verification skills that H3 concerns.

6. Estimation results

6.1 The 2026 country dataset

Table 3 collects the headline results of eleven GASA 2026 country reports, together with the Chainalysis adoption rank and the Findex account-ownership rate. Loss prevalence runs from 1% in Japan to 36% in Kenya, and the three African markets occupy the top three places. The ordering is not the ordering of exposure. Kenya has the highest exposure (92%), but Canada and Nigeria report the same exposure of 84% while 6% of Canadians and 22% of Nigerians lost money or data. The difference arises after the first contact. Of those who saw a scam, 80% of Nigerians interacted with it, against 45% of Canadians. Dividing loss prevalence by exposure gives a rough conversion rate from exposure to loss, which is 39% in Kenya, 26% in Nigeria and 22% in South Africa, against 11% in the United States, 7% in Canada and 5% in Japan. Reporting to the authorities is low everywhere (3% to 7%), apart from Japan, where GASA's figure is the share of scam victims who reported to the police and is not comparable.

The payment-channel question lets us look at crypto directly. Among victims who lost money, the share who say the scammer received their money through a cryptocurrency exchange or digital currency ranges from 4% in Brazil to 21% in Spain, and is 11% to 13% in the three African markets. Multiplying this share by loss prevalence gives an approximate crypto-channel loss rate, the share of adults who lost money through a crypto exchange. It is 4.0% in Kenya, 2.6% in Nigeria and 2.5% in South Africa, and between 0.4% and 1.9% in the other countries with data. The figure is an approximation and probably an upper bound, because the payment question allows several answers and its base is people who lost money, whereas the loss-prevalence figure also counts people who lost only data. The bases are also small, from 63 to 444 victims. For Canada (17%, n = 70) a 95% Wilson interval runs from about 10% to 28%. We therefore treat the crypto-channel figures as coarse.

Table 3. Headline results of the GASA 2026 country reports, with crypto-channel, adoption and account-ownership indicators

Country	n	Saw scam (%)	Confid. (%)	Interact. of exposed (%)	Lost money/ data (%)	Report. (%)	Crypto share of payments, % (base)	Crypto-channel loss (%)	Chain. rank	Account (%)
Kenya	1,230	92	54	83	36	7	11 (444)	4.0	25	90
Nigeria	1,180	84	49	80	22	4	12 (254)	2.6	6	63
South Africa	1,280	85	85	71	19	6	13 (243)	2.5	28	81
Brazil	1,170	81	42	56	10	4	4 (113)	0.4	5	86
United States	3,110	82	86	56	9	4	13 (280)	1.2	2	97
Spain	1,360	77	-	53	9	4	21 (125)	1.9	31	98
Germany	2,550	66	57	44	7	3	17 (167)	1.2	21	98
Taiwan	1,240	75	66	52	7	3	12 (85)	0.8	n/l	96
Canada	1,160	84	88	45	6	3	17 (70)	1.0	27	98
Portugal	1,360	77	-	41	5	3	11 (63)	0.6	45	91
Japan	1,200	22	57	53	1	18*	n/a	n/a	19	99

Notes: Sorted by loss prevalence. Crypto share = percentage of victims who lost money whose scammer received it through a cryptocurrency exchange or digital currency (GASA Q13; base in brackets). Crypto-channel loss = loss prevalence × crypto share (approximate). Chain. rank = Chainalysis 2025 overall adoption rank (1 = highest); n/l = not listed in the extracted table. Account = adults with an account, Findex 2024. *Japan reports the share of scam victims who reported to the police. The Portugal landing page gives 9% for loss prevalence, but the report itself gives 5%, which we use. Sources: GASA (2026a) country reports; Chainalysis (2025); Klapper et al. (2025).

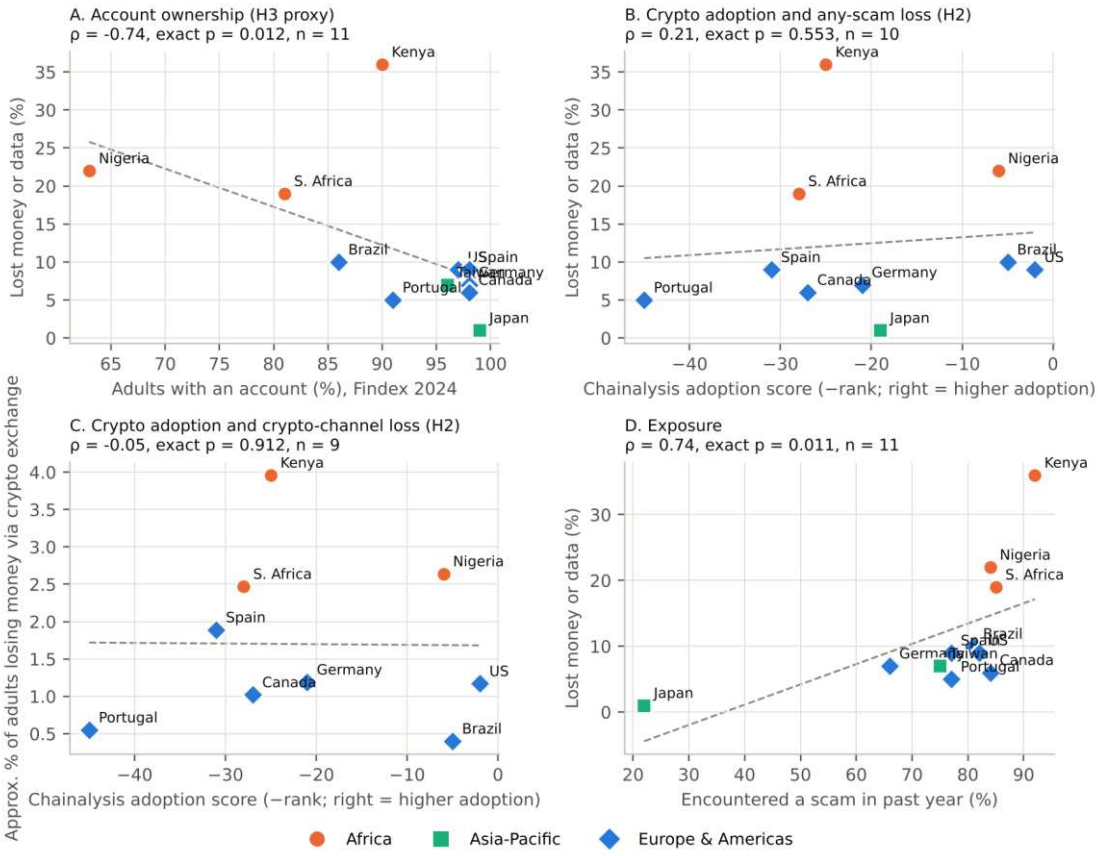


Figure 5. Loss prevalence and crypto-channel loss against country characteristics, GASA 2026 wave. Dashed lines are least-squares lines for orientation only. Panels B and C omit Taiwan, which is not in the Chainalysis table; panel C also omits Japan, which has no payment-channel data.

6.2 Tests of H2 and H3

Table 4 and Figure 6 summarize the rank correlations. **H2 receives no support, but is not rejected.** The correlation between adoption and loss prevalence is small and positive (ρ = 0.21, bootstrap interval -0.46 to 0.74). With the crypto-specific outcome it is zero (ρ = -0.05, interval -0.72 to 0.79), and adoption is not related to the crypto share of victim payments either (ρ = -0.25). The intervals are wide enough to contain strong effects of either sign. In the 2025 extreme-group sample the sign is positive (ρ = 0.31, p = 0.22), which fits the direction of H2, but that sample is selected on the outcome and the interval, -0.16 to 0.65, includes zero. The

median adoption rank is 14.5 in the ten highest-loss markets and 48.5 in the eight lowest-loss markets with a rank, and a permutation test of the difference gives p = 0.20 (Mann-Whitney p = 0.36). The sign of the 2026 coefficient depends on the sample, since it is 0.41 outside Africa (n = 7) and 0.21 overall. We conclude that the public data are uninformative about H2 and are not evidence against a modest positive effect. The pattern in the country data is easy to see: Brazil and the United States are among the highest-adoption markets but have low crypto-channel losses (0.4% and 1.2%), while Kenya, which ranks 25th, has the highest (4.0%).

Table 4. Rank correlations between loss measures and country characteristics

Association	n	Spearman ρ	Permutation p	Bootstrap 95% interval	Leave-one-out range	Holm p
H2 Adoption and any-scam loss prevalence (2026)	10	0.21	0.553	-0.46 to 0.74	0.02 to 0.36	1.00
H2 Adoption and crypto-channel loss (2026)	9	-0.05	0.912	-0.72 to 0.79	-0.33 to 0.29	1.00
H3 Account ownership and any-scam loss prevalence (2026)	11	-0.74	0.012	-0.96 to -0.20	-0.84 to -0.65	0.035
Account ownership and crypto-channel loss (2026)	10	-0.23	0.531	-0.76 to 0.75	-0.44 to 0.02	exploratory
Adoption and crypto share of victim payments (2026)	9	-0.25	0.509	-0.89 to 0.57	-0.58 to -0.04	exploratory
Adoption and exposure (2026)	10	-0.01	0.979	-0.65 to 0.62	-0.20 to 0.11	exploratory
Scam-recognition confidence and loss (2026)	9	-0.47	0.205	-0.86 to 0.29	-0.62 to -0.33	exploratory
Exposure and any-scam loss prevalence (2026)	11	0.74	0.012	0.17 to 0.98	0.65 to 0.91	exploratory
H2 Adoption and loss % of GDP (2025 extremes)	18	0.31	0.215	-0.16 to 0.65	0.22 to 0.41	descriptive
H3 Account ownership and loss % of GDP (2025 extremes)	19	-0.41	0.084	-0.77 to 0.09	-0.54 to -0.30	descriptive

Notes: Permutation p-values are exact for n of nine or fewer and based on 400,000 random permutations for the 2026 rows with ten or eleven countries and 200,000 for the 2025 sample. Holm p-values apply to the three confirmatory tests (the crypto-specific version of H2 was added after the payment-channel data were obtained). The 2025 rows use log loss as a percentage of GDP and are descriptive because the markets were selected for extreme values. Adoption is entered as minus the Chainalysis rank, so a positive coefficient means that higher adoption goes with higher loss.

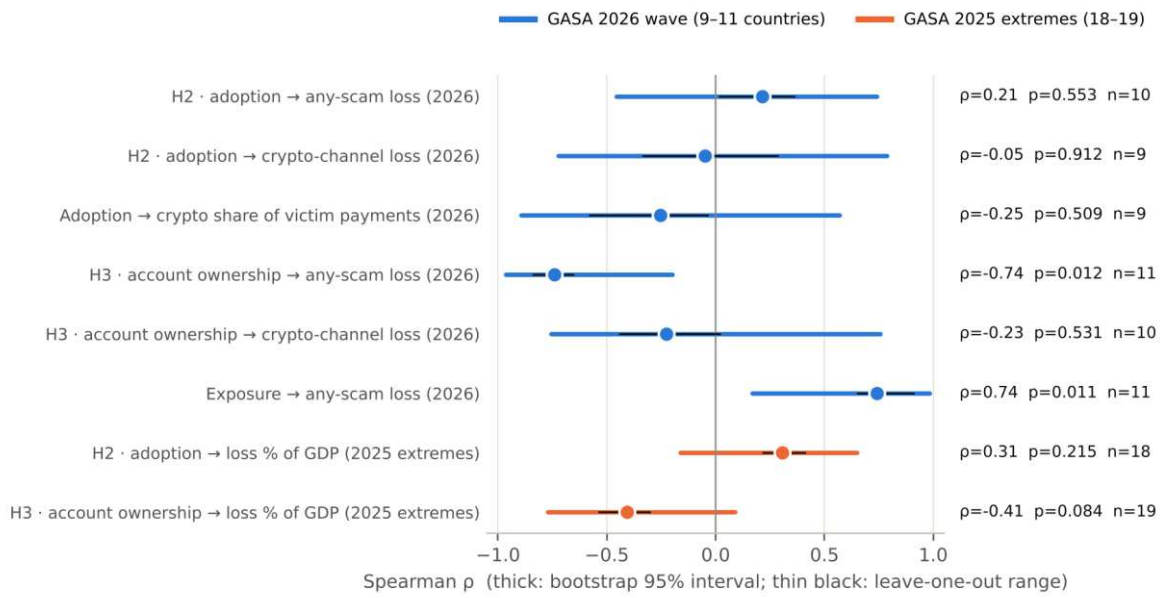


Figure 6. Rank correlations with bootstrap and leave-one-out uncertainty. Blue: GASA 2026 wave. Orange: GASA 2025 extreme-group sample.

The direction of H3 is supported, with strong caveats. Across eleven countries, account ownership and loss prevalence are negatively related (ρ = -0.74, permutation p = 0.012, Holm-adjusted p = 0.035). The coefficient stays between -0.84 and -0.65 when any one country is dropped. Outside Africa, however, it is smaller (ρ = -0.46, n = 8), and within the three African countries alone the relation is positive (Kenya has the highest account ownership of the three and the highest loss), so a good part of the association is a contrast between the African markets and the rest. Account ownership is also only weakly related to the crypto-channel loss rate (ρ = -0.23, p = 0.53), so the result is about scam losses in general and not about crypto. Account ownership measures access to the financial system, not the ability to judge a scam. It also tends to rise with income and institutional quality, so the coefficient probably bundles several channels. Self-rated confidence in recognizing scams, which is closer to a capability measure, is negatively but weakly related to loss (ρ = -0.47, p = 0.21, n = 9). Confidence is highest in Canada (88%), the United States (86%) and South Africa (85%), and South Africa's loss prevalence is 19%, which echoes GASA's own remark that high confidence does not protect

people from engaging with scams. We read the H3 result as consistent with the hypothesis in direction and as evidence about financial access and development, not literacy.

Among the exploratory results, exposure is positively related to loss prevalence ($\rho = 0.74$, $p = 0.012$), but the coefficient is 0.42 outside Africa. The share of exposed people who interacted with the scam is almost mechanically related to loss ($\rho = 0.86$), since loss requires interaction, and we treat it as a description of the funnel and not as a test. Figure 7 shows the pairwise correlations. Account ownership is negatively related to exposure, interaction, loss and the share reporting mental-health effects, which suggests that the African markets differ from the others on several dimensions at the same time. One further pattern concerns crypto. The crypto share of victim payments is higher where self-rated confidence is high ($\rho = 0.75$, $n = 9$) and where account ownership is high ($\rho = 0.70$, $n = 10$). In richer markets fewer people lose money, but a larger share of those who do pay through exchanges. That clustering is the main reason we do not read any of these coefficients causally.

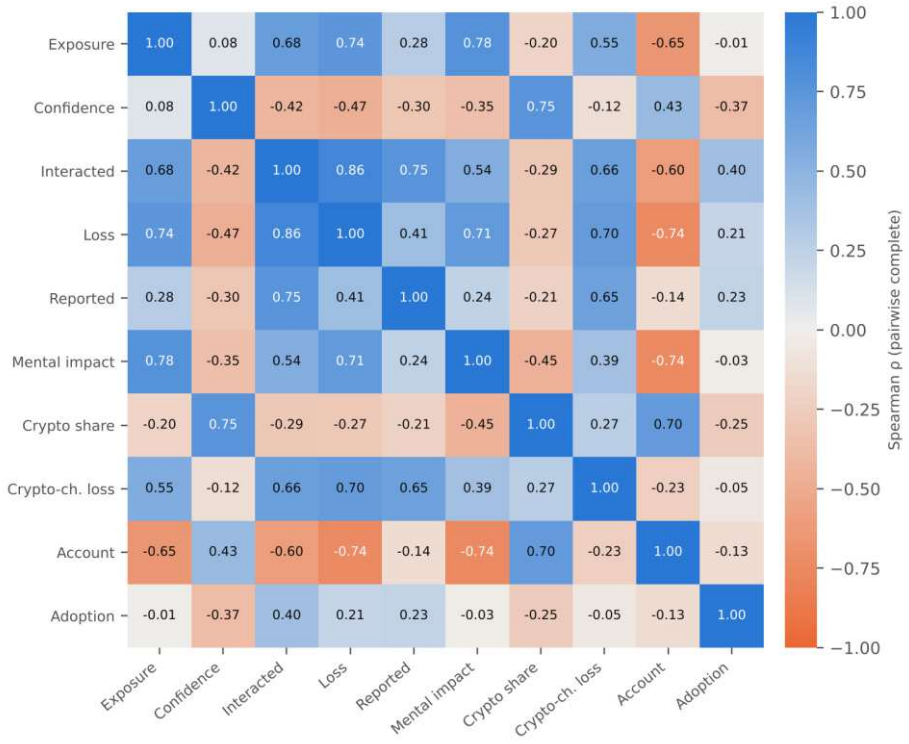


Figure 7. Spearman correlations among the 2026 country indicators (8 to 11 countries per cell). Blue = positive, orange = negative.

6.3 Robustness and consistency

Three checks bear on how far these results can be trusted. First, dropping each country in turn leaves the sign of the H3 and exposure coefficients unchanged, and it does not move the H2 coefficients far from zero (any-scam loss: 0.02 to 0.36; crypto-channel loss: -0.33 to 0.29). No single market therefore produces the results that are not near zero. Second, the 2026 loss prevalence agrees with the 2025 loss-to-GDP ranking in the five markets that appear in both (Kenya, Nigeria, Brazil, the United States and Taiwan): the rank correlation is 0.90, with an exact two-sided p-value of 0.083. Five points are few, but they suggest that the 2026 outcome orders countries similarly to the 2025 outcome, even though the two are defined differently. Third, Figure 8 repeats the main comparison for the 2025 extreme-group sample, where account ownership is again negatively related to loss as a share of GDP (ρ = -0.41, p = 0.08), and adoption is positively but weakly related. Because that sample keeps only the countries at

the top and bottom of the loss ranking, correlations in it overstate the spread of the outcome and understate the middle of the distribution.

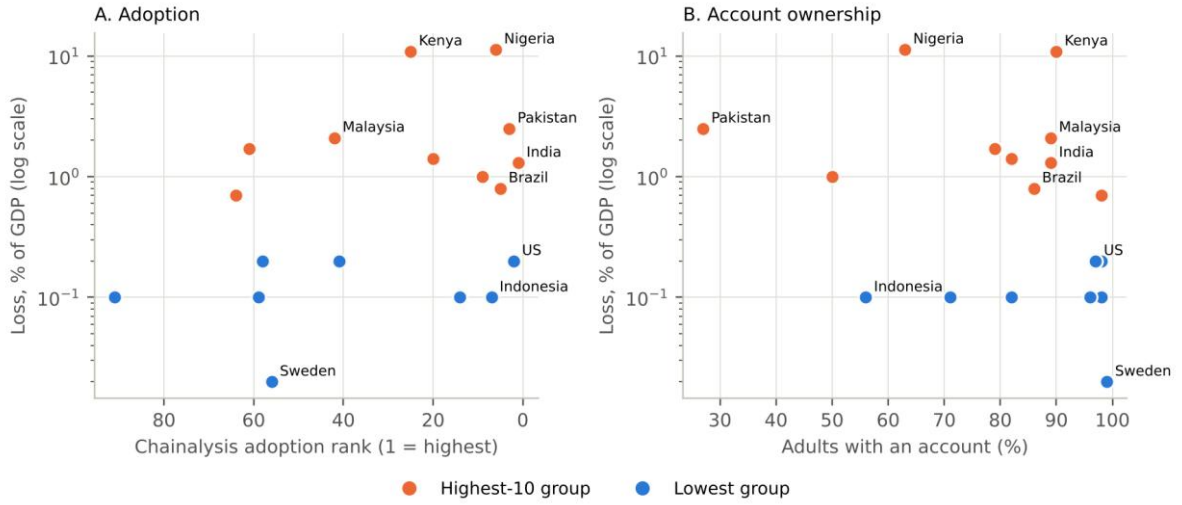


Figure 8. Estimated loss as a share of GDP against adoption rank (A) and account ownership (B), GASA 2025 extreme groups. Log scale; Taiwan is omitted from panel A.

6.4 What cannot be tested with the public data

H1 at the country level, and the extensions in Section 2.4. These require a country-level measure of AI-enabled scam sophistication, a regulatory-strength index or case-level operation data. None exists in the public data that we could reach. The GASA reports contain a question about whether people use AI tools to check suspicious messages, but it measures consumer behaviour and not how sophisticated the scams are. The only direct evidence on sophistication is the operation-level comparison in Section 5.1, which cannot be matched to countries. We therefore make no claim about the interaction of sophistication with adoption or regulation, or about cross-border reach, beyond the descriptive evidence in Section 5. A regulatory index could be constructed from official sources, but it would have to be coded under rules fixed in advance, and we have not done that here.

7. Discussion and implications

The evidence assembled so far supports a cautious reading, and the new country data change its emphasis. At the operation level, AI-linked scam operations are far more productive on every measure Chainalysis reports. At the country level, the burden of losses is not a simple function of adoption, regional crypto payment shares do not track loss prevalence, and average losses per victim are highest in wealthier markets while the share of adults losing money is highest in South America and Africa. The crypto-specific data add a caution. Crypto exchanges are named as the payment route by roughly one victim in ten in the three African markets and by up to one in five in Spain, so in every country with data most victim payments still go through banks, cards and wallets. Crypto adoption does not predict where crypto-channel losses are highest. A further result from the 2026 wave is that exposure does not predict loss as closely as the later stages of the funnel do: markets with the same exposure differ by a factor of three or four in loss prevalence, and the difference lies in how many exposed people engage with the scam and how many of those lose. This is the pattern that the term investor vulnerability is meant to capture, and it implies that measures acting after first contact, such as friction at payment, warnings at the moment of transfer and fast reporting, may matter as much as measures that reduce the number of approaches. The policy lesson is that interventions aimed only at crypto users will miss much of the problem, since bank transfers and card payments remain the most common victim payment channels and the scam ecosystem blends channels. Chainalysis also notes that many pig-butchering and investment scams combine impersonation, social engineering and wallet-based techniques (as reported by Cointelegraph).

For regulators, the findings point to the cash-out and payment stages as the places where AI-driven scale can be constrained. Licensing and travel-rule enforcement, faster freezing of suspected scam funds and clearer liability for platforms that carry scam advertising all act at

those stages. Whether such measures weaken the link between sophistication and losses is a question for future work that has a country-level sophistication measure and a regulatory index. For exchanges and AI providers, the descriptive premium for operations linked to AI tooling suggests that vendor-level monitoring and abuse reporting could be a high-leverage control. For investors, the confidence gap documented by GASA suggests that awareness campaigns should teach verification routines, such as calling back on an independently sourced number, rather than relying on cues such as poor grammar that generative AI removes (GASA finds that 27% of adults check spelling and grammar, one of the less effective verification methods it identifies).

8. Limitations

- **Outcome measurement.** Survey losses are self-reported, collected through an online panel, and trimmed for outliers. GASA states that its method changed in 2025, so wave-to-wave comparisons are unreliable.
- **Novelty claim.** The statement that no study connects the three strands rests on a targeted web search, not on a systematic search of Scopus and Web of Science. It should be confirmed with a registered search protocol and a PRISMA flow diagram.
- **Small samples.** The 2026 analysis rests on nine to eleven countries, three of them African. Intervals are wide, p-values are exploratory except for the two confirmatory tests, and clustering by region and income cannot be separated from the effects of interest.
- **Outcome definition.** The 2026 outcome is financial or data loss from any scam. It is not specific to crypto, and the public summaries give no payment-channel breakdown by country. The crypto-channel loss rate is a product of two survey proportions with small bases, and the tests of H3 concern scam losses in general.
- **Incomplete country data.** Country values for 2025 are published mainly in charts or an interactive map, and the PDF text gives only extreme groups. Full 2026 reports are restricted to

GASA members and were not used, and the free version of the Japan report has no payment-channel page.

- **Proxy validity.** Findex indicators measure access and use, not literacy. Sumsub rates measure identity-verification fraud and come from a commercial vendor. The Chainalysis index is based on web-traffic attribution.
- **Secondary sourcing.** The Chainalysis 2026 figures are quoted from press coverage of the report. They should be checked against the primary report.
- **Selection and ecology.** The extremes in Table 2 are selected on the outcome, and country-level associations need not hold for individuals.
- **Power.** With about 42 markets, only moderate to large effects can be detected.

9. Conclusion

Industry data show that scam operations linked to AI tooling are several times more productive than other operations, and that impersonation has become the fastest-growing scam type. Public country data from GASA's 2026 reports show that loss prevalence varies by a factor of more than thirty across eleven countries, and that the variation arises mostly after exposure, since markets with the same share of people seeing a scam differ widely in the share who lose money. In these data, crypto adoption is not associated with loss prevalence or with the approximate crypto-channel loss rate (H2 is unsupported but not rejected), and account ownership is negatively associated with loss prevalence (H3 supported in direction). Both results rest on very few countries, and the second is confounded with income and region and is weaker outside Africa. The AI-sophistication hypothesis (H1) is supported only at the operation level and descriptively. Whether sophistication interacts with adoption and regulation, and whether AI-enabled scams are more cross-border, are left for future work, because the required country-

level measures do not exist in public sources. The next step is to build those measures, obtain the full 2026 country files, and estimate the models on the 42-country sample.

References

Akerlof, G. A., and R. J. Shiller. 2015. *Phishing for Phools: The Economics of Manipulation and Deception*. Princeton, NJ: Princeton University Press.

Chainalysis. 2025. *The 2025 Geography of Crypto Report: What Regional Trends Reveal About What's Next in Crypto*. Chainalysis, Inc.

Chainalysis. 2026. *2026 Crypto Crime Report* (crypto scams chapter, January 2026). Chainalysis, Inc. Figures cited as reported by Cointelegraph, The Next Web and Forklog (see Appendix A).

Chesney, R., and D. K. Citron. 2019. "Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security." *California Law Review* 107 (6): 1753–1819. https://doi.org/10.15779/Z38RV0D15J.

Cohen, L. E., and M. Felson. 1979. "Social Change and Crime Rate Trends: A Routine Activity Approach." *American Sociological Review* 44 (4): 588–608. https://doi.org/10.2307/2094589.

Cross, C. 2022. "Using Artificial Intelligence (AI) and Deepfakes to Deceive Victims: The Need to Rethink Current Romance Fraud Prevention Messaging." *Crime Prevention and Community Safety* 24 (1): 30–41. https://doi.org/10.1057/s41300-021-00134-w.

Cumming, D. J., S. Johan, and A. Pant. 2019. "Regulation of the Crypto-Economy: Managing Risks, Challenges, and Regulatory Uncertainty." *Journal of Risk and Financial Management* 12 (3): 126. https://doi.org/10.3390/jrfm12030126.

Ebekozien, A., C. Aigbavboa, W. D. Thwala, M. S. Samsurijan, M. A. H. Ahmed, J. Aliu, and S. A. Adekunle. 2024. "Appraising the Application of Cryptocurrency Technologies in the Nigerian Built Environment: Stakeholders' Perspective." *International Journal of Building Pathology and Adaptation* 42 (7): 93–112. https://doi.org/10.1108/IJBPA-04-2024-0070.

Elliptic. 2024. *Horizon Scanning Report 2024: Exploring Emerging Risks and Trends in Crypto and Artificial Intelligence*. Elliptic. https://dcfintechweek.org/wp-content/uploads/2024/09/Elliptic-AI-enabled_crime_in_the_cryptoasset_ecosystem.pdf.

Gandal, N., J. T. Hamrick, T. Moore, and T. Oberman. 2018. "Price Manipulation in the Bitcoin Ecosystem." *Journal of Monetary Economics* 95: 86–96. https://doi.org/10.1016/j.jmoneco.2017.12.004.

Global Anti-Scam Alliance (GASA). 2024. *State of Scams in Japan Report 2024*. Report landing page. https://gasa.org/knowledge-base/reports/state-of-scams-in-japan-report-2024.

Global Anti-Scam Alliance (GASA). 2026a. *State of Scams Country Reports 2026* (Brazil, Canada, Germany, Japan, Kenya, Nigeria, Portugal, South Africa, Spain, Taiwan, United States; Africa and Europe regional editions). Free versions of the reports. https://gasa.org/knowledge-base. Accessed October 3, 2026.

Global Anti-Scam Alliance (GASA), and Feedzai. 2025. *Global State of Scams 2025*. Report; fieldwork by Opinium, March 7–20, 2025. Global Anti-Scam Alliance.

Griffin, J. M., and K. Mei. 2024. "How Do Crypto Flows Finance Slavery? The Economics of Pig Butchering." SSRN Working Paper 4742235. https://doi.org/10.2139/ssrn.4742235.

Klapper, L., D. Singer, L. Starita, and A. Norris. 2025. *The Global Findex Database 2025: Connectivity and Financial Inclusion in the Digital Economy*. Washington, DC: World Bank. https://doi.org/10.1596/978-1-4648-2204-9.

La Porta, R., F. Lopez-de-Silanes, and A. Shleifer. 2006. "What Works in Securities Laws?" *The Journal of Finance* 61 (1): 1–32. https://doi.org/10.1111/j.1540-6261.2006.00828.x.

Lusardi, A., and O. S. Mitchell. 2014. “The Economic Importance of Financial Literacy: Theory and Evidence.” *Journal of Economic Literature* 52 (1): 5–44. https://doi.org/10.1257/jel.52.1.5.

Mirsky, Y., and W. Lee. 2021. “The Creation and Detection of Deepfakes: A Survey.” *ACM Computing Surveys* 54 (1): Article 7. https://doi.org/10.1145/3425780.

Page, M. J., J. E. McKenzie, P. M. Bossuyt, et al. 2021. “The PRISMA 2020 Statement: An Updated Guideline for Reporting Systematic Reviews.” *BMJ* 372: n71. https://doi.org/10.1136/bmj.n71.

Sumsub. 2026. “Fraud Trends 2026: AI Scams, Deepfakes, and Emerging Threats.” Web article, April 15, 2026. https://sumsub.com/blog/fraud-trends/.