International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 17(2025) pp. 1-21
https://doi.org/10.70917/ijcisim-2026-6101
Article

The Small Merchant Fraud Gap: Why US Businesses Under $10M in Revenue Absorb Disproportionate Fraud Losses.

Olid Sarker[1*], Habibur Rahman[2]

[1]Feliciano School of Business, Montclair State University, olidsarker@gmail.com
[2]Northern University of Bangladesh, nubhabib@gmail.com

Abstract: In this last decade, industry surveys, central-bank studies and the information-systems literature converged on a working explanation for why US merchants with less than $10 million in annual revenue appeared to lose a larger share of their sales to payment fraud than larger firms. That explanation was never tested as a whole. This study was designed and written in 2025 as a retrospective; it makes no claim to have been produced in the period it studies. We assemble an era-bounded corpus of 54 sources published between 2017 and 2024, extract eight claims the period treated as settled, and test each against a seeded synthetic merchant panel whose data-generating process instantiates the era's own documented assumptions about channel mix, liability allocation, detection tooling, dispute recovery, false declines and large-event losses. A Shapley decomposition with common random numbers attributes a pooled 19.9 basis-point gap in retained loss ratio between the two size bands to five mechanisms. Channel exposure to card-not-present commerce accounts for 36% of the gap, under-adoption of detection tooling for 21%, weaker dispute recovery for 18%, deliberate targeting for 6%, and EMV-related liability lag for 2%; the remaining 17% is carried by rare large-value events. Six of the eight era claims are held, two are qualified, and none is overturned, but the audit reverses the era's emphasis: the liability rule that the period blamed does not differ by firm size, and the targeting narrative that dominated vendor reports explains almost nothing. Evidence published in 2025 and 2026 is consistent with these verdicts. All data are synthetic, and all code is released.

Keywords: payment fraud; small and medium-sized enterprises; chargebacks; card-not-present fraud; retrospective audit; synthetic data; Shapley decomposition

1. Introduction

A merchant that sells online in the United States carries a risk that its card-present competitors largely do not: when a transaction turns out to be fraudulent, the card networks' rules place the loss on the merchant rather than on the issuing bank. Through the period that this paper studies, 2017 to 2024, that single rule was invoked again and again to explain why the smallest merchants seemed to lose the most. Trade surveys reported fraud costs per dollar of revenue that fell as firm size rose. Central-bank work documented that merchants, not issuers, had come to bear the plurality of debit-card fraud losses. Information-systems and security-economics scholars described small firms that bought fewer controls, reviewed more orders by hand and recovered less of what they lost. By the end of the window a composite story had formed, and it was repeated often enough to pass for an established finding.

This study was designed and written in 2026 as a retrospective; it makes no claim to have been produced in the period it studies. Its purpose is to take that composite story apart. We treat 2017–2024 as a closed window, collect what the period's literature actually claimed, and ask which of those claims survive when the period's own assumptions are written down, implemented in a simulator, and tested together. The angle is a verdict audit: each claim is stated in testable form, a decision rule is fixed before any result is produced, and the outcome is recorded as held, qualified, or

overturned. The method follows the retrospective replication tradition in which a field's working assumptions are instantiated in a synthetic benchmark so that their joint implications can be examined without the confounds that make observational fraud data so difficult to use (Lopez-Rojas et al., 2018; Woods & Böhme, 2021).

Three features of the design deserve emphasis at the outset. First, the corpus analyzed is bounded by the window, but the paper's framing is not: where a later source settles an era claim or supplies a comparison, we cite it and label it as post-window evidence. Second, no number reported here was collected from real merchants. The panel is synthetic, seeded, and generated by code that ships with the paper; it is era-styled in that its parameter ranges, platform mix, and measurement conventions are those of 2017–2024, and it is never presented as era-originated. Third, the contribution is the scoreboard. Descriptive size-banded loss figures existed throughout the window; what did not exist was a decomposition of the small-merchant gap into the mechanisms the period proposed, with each mechanism's weight measured under the period's own assumptions.

The findings cut against the era's emphasis. Under 2017–2024 assumptions, the retained fraud-loss ratio of merchants under $10 million in revenue is 1.57 times that of larger merchants, which supports the headline claim. But the two explanations that the period repeated most often carry the least weight. The liability rule itself does not vary with firm size, so it cannot by itself produce a size gap; what varies is how much of a small merchant's volume flows through the channel where that rule bites. Deliberate targeting of small merchants, a recurring theme in vendor reports, explains 6% of the gap, and the gap persists almost unchanged when targeting is removed. The mechanisms that do the work are channel exposure, detection tooling, and dispute recovery. The rest of the paper documents how these verdicts were reached.

Section 2 reconstructs what the window's literature held and states the eight claims. Section 3 describes the era-faithful generator, the decomposition, and the pre-registered decision rules. Results follow in Section 4, the verdict matrix and its comparison with post-window evidence in Section 5, and implications, limitations and conclusions in Sections 6–8.

2. What the 2017–2024 literature held

2.1 Corpus construction

The era corpus contains 54 sources published between 2017 and 2024, plus one pre-window industry note retained only as the baseline for the 2015 liability shift. Each entry was verified in one of three ways and is tagged accordingly in the reproducibility package: 37 entries were checked against Crossref by digital object identifier, 15 are industry or official reports retrieved from their publishers' web pages, and the remainder are preprints without a DOI. Sources were located through bibliographic database queries on merchant fraud losses, chargeback liability, card-not-present fraud, false declines, small-firm cybersecurity investment, and business email compromise, filtered to the window, and through targeted retrieval of the annual survey series that the period's practitioners relied upon. A separate verdict corpus of four post-window sources (2025–2026) is held apart and used only in Section 5. The complete bibliography, with verification tags, is Appendix A.

2.2 The composite explanation

The window opened in the aftermath of the October 2015 EMV liability shift. Hayashi, Markiewicz and Minhas (2018) used merchant-processor and PIN-debit network data to show that merchant fraud-loss rates for magnetic-stripe card-present transactions rose roughly sixfold after the shift, that the increase was attributable to the reallocation of liability rather than to any rise in gross fraud, and that card-not-present fraud rates were rising over the same period with merchants generally liable for them. Hayashi (2019) placed the US experience alongside other countries that had migrated to chip cards earlier and found the same migration of fraud toward the online channel. The Board of Governors' 2021 debit-card report, released in revised form in 2023, quantified the end state for covered issuers: fraud losses had reached 17.5 basis points of transaction value, up from 12.4 in 2019 and 7.8 in 2011; merchants bore 47%

of those losses, issuers 34% and cardholders 20%; and card-not-present transactions, 32% of value, generated 50% of losses (Board of Governors of the Federal Reserve System, 2023).

Practitioner surveys supplied the size gradient. LexisNexis Risk Solutions' annual True Cost of Fraud series reported a cost per dollar of fraud that rose from $2.94 in 2018 to $3.13 in 2019 and, for merchants that had not integrated fraud prevention with customer-experience operations, $4.24 by 2022; the 2019 release also stated that fraudsters were adding small-business targets that typically used fewer protective tools (LexisNexis Risk Solutions, 2019, 2020, 2022). The CyberSource and Merchant Risk Council benchmark surveys reported that merchants in the smallest revenue tier used fewer fraud-detection tools, reviewed a comparable share of orders by hand but rejected a smaller share of reviewed orders, and were less likely to have a formal approach to friendly fraud (CyberSource & Merchant Risk Council, 2021). By the 2024 edition, the headline loss rate was about 3% of e-commerce revenue, described as skewing higher for small and mid-market merchants, with global dispute win rates below 20% (Merchant Risk Council et al., 2024). Chargebacks911 (2023) reported that merchants estimated 44% of chargebacks to be friendly fraud and that merchants using representment providers recovered substantially more than those managing disputes internally.

The occupational-fraud and treasury-fraud series added two further elements. The Association of Certified Fraud Examiners reported median losses for organizations with fewer than 100 employees of $150,000 in its 2022 report and $141,000 in 2024, against overall medians of $117,000 and $145,000, with the explanation that smaller businesses typically operate fewer anti-fraud controls (Association of Certified Fraud Examiners, 2022, 2024). Because a $150,000 loss is a far larger fraction of revenue for a $2 million firm than for a $200 million firm, these dollar medians were routinely read as evidence of a size gradient in relative terms. The Association for Financial Professionals' payments-fraud series documented that 65% of responding organizations experienced attempted or actual payment fraud in 2022 and 80% in 2023, that business email compromise affected 71% of victims, and that 30% of victims recovered nothing while 41% recovered at least three-quarters of what they lost (Association for Financial Professionals, 2023, 2024). Cross and Gillett (2020) described the mechanics of business email compromise and the difficulty of recovery once funds have moved.

The security-economics literature of the window explained the investment side. Biancotti (2017) showed from Bank of Italy survey data that median cybersecurity expenditure rose steeply with firm size, that large firms could spread fixed protective costs over more revenue, and that small firms seldom had the internal capacity to assess risk. Ekelund and Iskoujina (2019) and Gordon, Loeb and Zhou (2020) formalized the trade-off between security spending and expected loss, while Fernandez de Arroyabe et al. (2023) found in UK survey data for 2018 and 2019 that prior attacks and existing capabilities, rather than exposure as such, drove investment. Reviews of small-firm practice (Bada & Nurse, 2019; Alahmari & Duncan, 2020; Benz & Chatterjee, 2020; Chidukwani et al., 2022; Kabanda et al., 2018; Ponsard et al., 2019) converged on the same picture: limited budgets, limited awareness, and reliance on manual or rule-based controls. Rawindaran et al. (2021) quantified the cost–benefit of machine-learning intrusion detection for UK SMEs. On the detection side, the machine-learning literature of the window (Dal Pozzolo et al., 2018; Fiore et al., 2019; Makki et al., 2019; Lucas et al., 2020; Carcillo et al., 2021; Hilal et al., 2022; Cherif et al., 2023; Tax et al., 2021) established both that modern classifiers materially outperformed rules and that their deployment carried data, engineering and labelling costs that only larger operations could absorb; Mekterović, Karan and Pintar (2021) asked explicitly where a card-not-present merchant should invest.

False declines were the window's counterweight. Aite Group's 2019 study for Clear Sale projected that losses from falsely declined legitimate orders would dwarf direct fraud losses by 2021; Datos Insights (2024) later put the average false-decline rate at 1.51% of e-commerce sales. The period took these figures as evidence that fraud prevention itself was a cost that fell hardest on merchants with the bluntest tools. One detail was rarely mentioned: the Aite sample was confined to firms with $100 million to $1 billion in revenue, so it could not speak to small merchants at all.

Finally, a measurement literature warned that all of the above was fragile. Riek and Böhme (2018) showed how estimates of cybercrime cost depend on sampling frame, question wording and the treatment of outliers; Eling and Wirfs (2019) documented the heavy tails of cyber-loss distributions; Kemp, Miró-Llinares and Moneva (2020) measured the dark figure of unreported fraud; Woods and Böhme (2021) cataloged the conditions under which cyber-risk quantification can be trusted; Jevtić and Lanchier (2020) modeled SME loss distributions structurally; Kamiya et al. (2021) and Schlackl, Link and Hoehle (2022) traced the firm-level consequences of incidents. Kalovya (2020) analyzed the determinants of occupational-fraud loss size. Bodker et al. (2022) scripted the card-not-present fraud process from the offender's side, and Whitty (2019) studied who is most susceptible.

2.3 The claims

Reading across the corpus, eight propositions were treated as settled by 2024. We state them as the period stated them; Section 3.4 gives each its testable form.

- **C1 (the gap).** Merchants under $10 million in revenue lose a larger share of revenue to fraud than larger merchants.
- **C2 (channel and liability).** The gap arises because small merchants transact disproportionately in the card-not-present channel, where the networks place fraud liability on the merchant, and because they lagged on EMV in the card-present channel.
- **C3 (recovery).** Small merchants recover less of what they lose, because they dispute less often and win less often.
- **C4 (tooling).** Fraud-detection tooling carries fixed costs that small merchants cannot spread, so they adopt less of it, and this under-adoption drives the gap.
- **C5 (false declines).** Small merchants suffer proportionally more false declines because they screen with blunter instruments.
- **C6 (the multiplier).** Every dollar of fraud costs a merchant roughly three dollars or more once fees, lost sales and prevention are counted, and the multiplier is larger for smaller merchants.
- **C7 (targeting).** Fraudsters deliberately target small merchants, and this targeting explains the gap.
- **C8 (measurement).** The size-banded loss ratios reported by the period's executive surveys measure the population gap faithfully.

3. Method

3.1 Design

The audit has three parts. A seeded generator produces a panel of synthetic US merchants observed annually from 2017 to 2024, with every structural parameter traceable either to a corpus source or to an explicitly labeled modelling assumption. A decomposition attributes the difference in retained fraud-loss ratio between merchants under and over $10 million in revenue to five mechanisms, using the Shapley value over counterfactual simulations that share random draws. Decision rules fixed in the analysis code before the first run convert the decomposition and related statistics into verdicts. The analysis is then repeated across ten seeds and across one-at-a-time variations of the labeled assumptions.

3.2 The era-faithful generator

The panel contains 20,000 merchants in four revenue bands: micro (under $1 million), small ($1–10 million), mid ($10–50 million) and large (over $50 million), with revenues drawn log-normally within band. The paper's unit of comparison is the under-$10 million group (micro and small bands) against the rest. Each merchant-year is

simulated through the sequence that the window's sources describe: exposure, attempt, detection, liability, dispute, and recovery, with costs of prevention and of false declines accumulated alongside.

Exposure. Each merchant has a card-not-present share of card volume drawn around a band-specific mean and rising through the window, with a step in 2020. The 2021 debit-card figure of 32% of value is the anchor (Board of Governors of the Federal Reserve System, 2023), the upward trend follows Hayashi (2019), and the assumption that small, e-commerce-leaning merchants carry a higher share is labeled as such.

Attempt. Attempted fraud is a basis-point rate on card-not-present and card-present volume, the former rising and the latter falling through the window in line with the post-EMV migration described by Hayashi et al. (2018) and Hayashi (2019). Under claim C7, attempts against micro and small merchants are scaled by 1.15 and 1.10 respectively, following the LexisNexis (2019) statement that fraudsters were adding small-business targets.

Detection. A merchant either operates detection tooling or does not, with adoption probabilities by band taken from the Cybersource and MRC (2021) finding that the smallest tier uses fewer tools and from the SME literature (Chidukwani et al., 2022; Fernandez de Arroyabe et al., 2023). Tooling stops a larger share of card-not-present attempts and produces a lower false-decline rate on legitimate orders; the average false-decline rate is anchored on Datos Insights (2024) and the manual-review share on Cybersource and MRC (2021). The detection rates themselves are labeled assumptions in the range the classifier literature implies.

Liability. Successful card-not-present fraud falls on the merchant with high probability, reflecting the network rules described by Hayashi et al. (2018); a small residual captures authentication-based liability shifts. Counterfeit card-present fraud falls on merchants that had not adopted EMV terminals, with adoption rising from band-specific 2017 levels toward near-universal levels by 2024 (Hayashi et al., 2018; Forrester Research, 2015 for the pre-window baseline).

Dispute and recovery. Chargebacks comprise true-fraud disputes and friendly-fraud disputes, the latter a share of card-not-present orders set below the Cybersource and MRC (2021) figure. Each chargeback carries a fixed fee. Friendly-fraud disputes are represented with probability and won at a rate that depend on whether the merchant uses a representment provider, anchored on the sub-20% global win rate (Merchant Risk Council et al., 2024) and the provider uplift reported by Chargebacks911 (2023); provider use rises with band.

Prevention cost. Merchants with tooling pay a fixed annual amount plus a variable fee on card-not-present volume; all merchants pay a per-order manual-review cost on the share of orders reviewed. The fixed-cost structure follows Biancotti (2017) and Ekelund and Iskoujina (2019).

Large events. Each merchant-year has a small, band-specific probability of a large-value event (business email compromise, account takeover or check fraud), with log-normal severity scaling sub-linearly with revenue and heavy tails (Eling & Wirfs, 2019). Recovery follows the AFP (2024) mixture: 30% recover nothing, 41% recover at least three-quarters, and the remainder recover a partial amount.

The full parameter table, with its source tag for every entry, is Appendix B. Parameters labeled [A] are the ones varied in the sensitivity analysis.

3.3 Outcome and decomposition

The primary outcome is the retained loss ratio: fraud losses the merchant actually absorbs after liability allocation and recovery, including net large-event losses, divided by revenue, pooled across years and value-weighted within the size group. A secondary outcome, the total fraud cost ratio, adds chargeback fees, the gross-margin value of falsely declined orders, and prevention spending. The ratio of the two is the multiplier of claim C6.

The gap is the difference in retained loss ratio between the under-$10 million group and the rest. Five mechanisms are defined as parameter blocks: channel exposure (band-specific card-not-present shares), targeting (attempt multipliers), tooling (adoption probabilities, which govern both detection and false declines), liability (EMV

adoption paths), and recovery (provider-use probabilities). A counterfactual gives the under-$10 million merchants the large band's value of one or more blocks while every random draw is held fixed. The Shapley value over the 32 subsets attributes the gap to each mechanism without dependence on ordering; the gap that remains when all five are equalized is reported as a residual, which by construction is the component carried by band-specific large-event rates and revenue scaling. We also report the single-factor share, the fraction of the gap closed when one block alone is equalized, because that is the quantity practitioners usually have in mind.

3.4 Pre-registered decision rules

The verdict rules were written into the analysis code before the generator was first run and were not altered afterward. C1 is held if the under-over ratio of retained loss ratios is at least 1.5 in every year of every seed, and qualified if the pooled ratio exceeds that value but some year–seed combination falls short. C2 is held if channel exposure and liability together account for at least 20% of the gap. C3 is held if recovery accounts for at least 10%. C4 is held if tooling accounts for at least 20% and prevention spending as a share of revenue is higher for the smaller group. C5 is held if the false-decline cost ratio is at least 1.25 times higher for the smaller group. C6 is held if the multiplier exceeds 2.0 in both groups and is larger for the smaller group. C7 is held if targeting accounts for at least 20% of the gap, qualified if the gap persists above the C1 threshold with targeting removed, and overturned otherwise. C8 is held if the gap observed through a simulated executive-survey sampling frame is within 15% of the value-weighted population gap, and overturned otherwise. The survey frame draws respondents with probability rising in tooling, provider use, and band, mimicking the fraud-executive panels of the period's surveys, and reports the equal-weighted mean of firm-level ratios as those surveys did.

3.5 post-window verdict corpus

Four sources published after the window are used to check the verdicts against what later became known: the Federal Reserve Bank of Kansas City briefings on card-not-present fraud rates (Hayashi, 2025, 2026), the 2025 AFP survey highlights (Association for Financial Professionals, 2025) and the 2025 LexisNexis retail study (LexisNexis Risk Solutions, 2025). None of these sources informed the generator.

4. Results

4.1 The gap under era assumptions

Figure 1 shows the retained loss ratio by year. Merchants under $10 million retain 45 basis points of revenue in 2017, rising to 62 by 2024; larger merchants move from 27 to 41. The 2020 step reflects the channel shift of the pandemic year. Pooled across the window, the ratios are 54.9 and 35.0 basis points, a ratio of 1.57 (range 1.54–1.59 across seeds). The annual ratio declines from 1.67 in 2017 to 1.52 in 2024 as EMV adoption converges and the pandemic pushes larger merchants into the card-not-present channel too. By band, retained losses are 66.7 basis points for micro merchants, 52.6 for small, 39.2 for mid, and 33.4 for large.

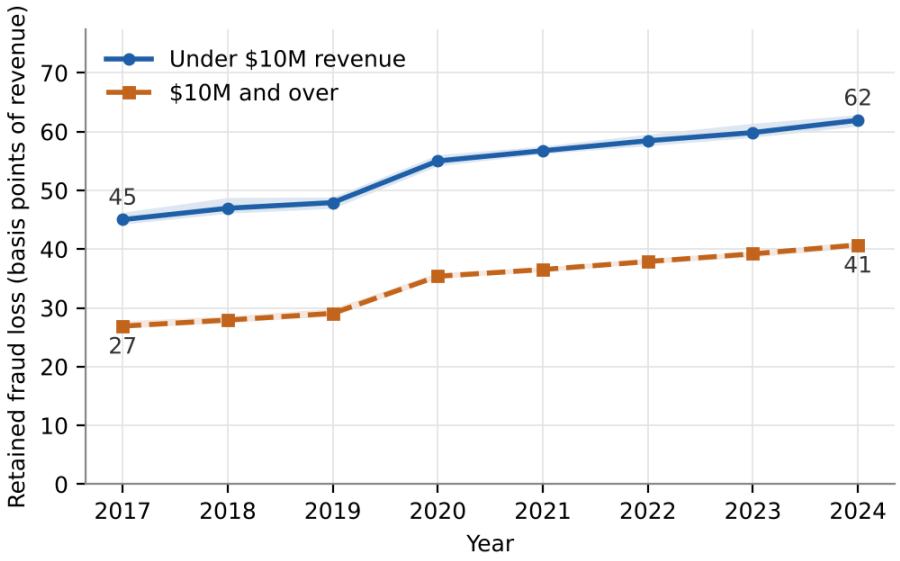


Figure 1. Retained fraud loss as a share of revenue by size group, 2017–2024. Lines are means across ten seeds; bands show the seed range.

Table 1 breaks the pooled figures into their components. Attempted fraud is 55.0 basis points of revenue for the smaller group against 43.5 for the larger; successful fraud after detection is 20.1 against 13.0. Of successful fraud, 82% is charged back to the smaller merchants and 69% to the larger, the difference coming entirely from channel mix and EMV status. The smaller group recovers 9% of what is charged back to it, the larger 17%. Net large-event losses add 6.2 basis points for the smaller group and 1.4 for the larger. Tool adoption is 42% and 88%.

Table 1. Pooled components of fraud cost, 2017–2024, basis points of revenue (mean of ten seeds; seed standard deviations below 1 basis point throughout).

Component	Under $10M	$10M and over
Attempted fraud	55.0	43.5
Successful fraud after detection	20.1	13.0
Retained direct and large-event loss	54.9	35.0
of which net large-event loss	6.2	1.4
Chargeback fees	15.7	11.9
False-decline margin loss	37.1	22.0
Prevention spending	51.7	23.0
Total fraud cost	159.5	91.9
Multiplier (total cost / retained loss)	2.90	2.62

4.2 Decomposition

Figure 2 and Table 2 give the Shapley attribution of the 19.9 basis-point gap. Channel exposure contributes 7.2 basis points (36%), tooling 4.1 (21%), recovery 3.5 (18%), targeting 1.2 (6%) and EMV-related liability lag 0.5 (2%). The residual of 3.5 basis points (17%) is the large-event component. Single-factor shares tell the same story: equalizing channel exposure alone closes 40% of the gap, tooling alone 24%, recovery alone 19%, targeting alone 8%, and liability alone 2%. With targeting removed, the gap ratio is 1.53; with channel exposure and liability equalized it is 1.33; with tooling and recovery equalized, it is 1.33.

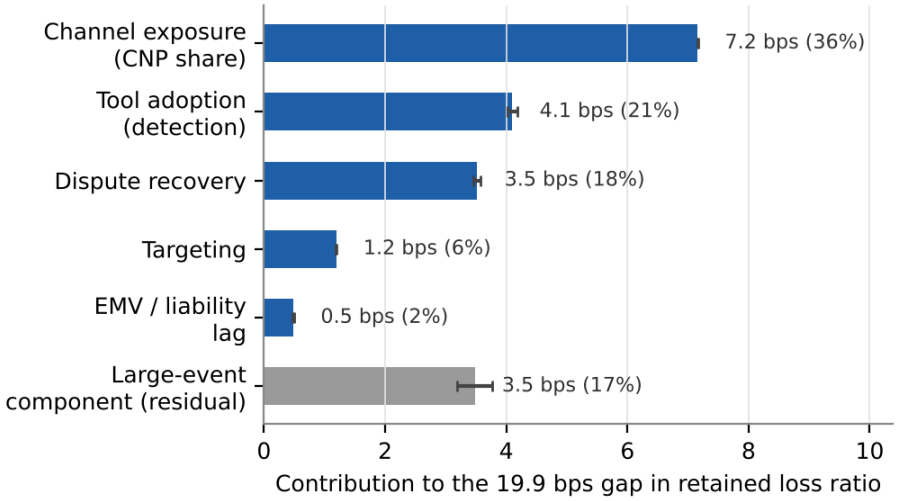


Figure 2. Shapley decomposition of the gap in retained loss ratio between merchants under $10M and $10M and over. Error bars are seed standard deviations; the residual is the large-event component.

Table 2. Shapley decomposition of the retained-loss gap (mean and standard deviation across ten seeds).

Mechanism	Basis points	Share of gap	Single-factor share
Channel exposure	7.2 (0.01)	36%	40%
Tool adoption	4.1 (0.08)	21%	24%
Dispute recovery	3.5 (0.06)	18%	19%
Targeting	1.2 (0.01)	6%	8%
EMV / liability lag	0.5 (0.01)	2%	2%
Residual (large-event component)	3.5 (0.30)	17%	—
Total gap	19.9 (0.31)	100%	—

Two features of the decomposition deserve comment. The liability block is small not because liability does not matter but because the rule does not vary with size: both groups face the same card-not-present rule, and by the second half of the window EMV adoption has converged enough that the card-present difference is slight. The rule is the precondition for the gap; the channel share is what turns the precondition into a difference. The second feature is that tooling and recovery, taken together, carry about as much of the gap as channel exposure, and these are the two mechanisms the merchant can act on.

4.3 Costs beyond retained loss

The total fraud cost ratio is 159.5 basis points for the smaller group and 91.9 for the larger, so the gap widens from 1.57 to 1.74 once fees, false declines and prevention are included. False-decline margin loss is 37.1 against 22.0 basis points, a ratio of 1.69. Prevention spending is 51.7 against 23.0, a ratio of 2.25, even though the smaller group adopts tooling less often: fixed costs and per-order manual review weigh more on a small revenue base. The multiplier is 2.90 for the smaller group and 2.62 for the larger.

4.4 Distribution and measurement

Figure 4 shows firm-level retained loss ratios. The median under-$10 million merchant retains 53 basis points, the mean 67 and the 90th percentile 78; for larger merchants the figures are 36, 38 and 49. The wedge between median

and mean in the smaller group is the large-event tail: 78% of small merchants experience no large event in eight years, and the gap ratio excluding large events is 1.45 rather than 1.57.

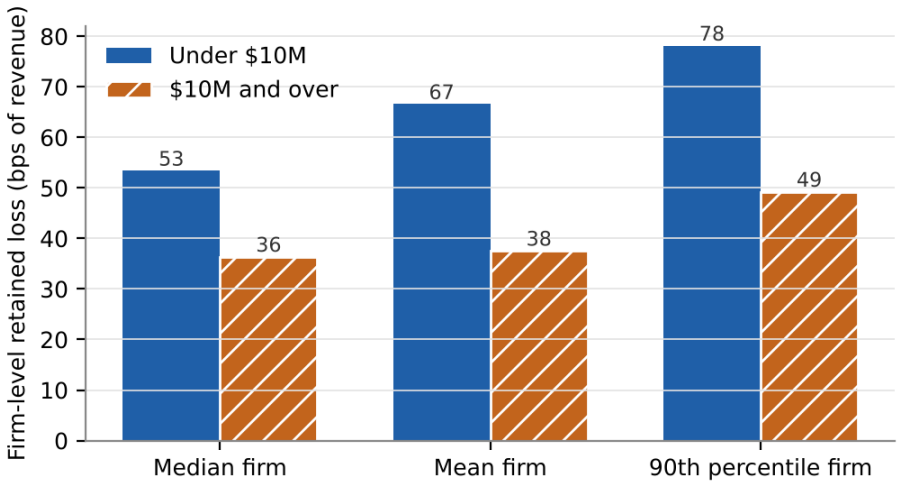


Figure 4. Firm-level retained loss ratio: median, mean and 90th percentile by size group (mean of ten seeds).

Weighting matters more than sampling. The value-weighted population gap ratio is 1.57; the equal-weighted population ratio, which is what a survey mean of firm-level percentages approximates, is 1.77. The simulated executive-survey frame, with about 3,800 respondents per run, yields 1.64, so the sampling frame inflates the value-weighted gap by a factor of 1.04 while the equal weighting that surveys use inflates it by 1.13. The selection of respondents toward firms with fraud functions partly offsets the weighting effect, because tooled small merchants have lower ratios than the un-tooled majority.

4.5 Sensitivity

Figure 3 varies the labeled assumptions one at a time at seed 0, where the baseline gap ratio is 1.57. Halving the detection advantage of tooling lowers the ratio to 1.51 and doubling it raises the ratio to 1.68. Cutting the card-not-present size differential to four percentage points gives 1.43; doubling it gives 1.76. Removing the provider advantage in recovery gives 1.44, removing large events gives 1.46, doubling their rate gives 1.69, and lowering card-not-present merchant liability to 0.75 gives 1.56. The tool fixed cost and the false-decline rate do not affect the retained-loss gap by construction, since they enter only the total-cost outcome. In no variation does the ratio fall below 1.43.

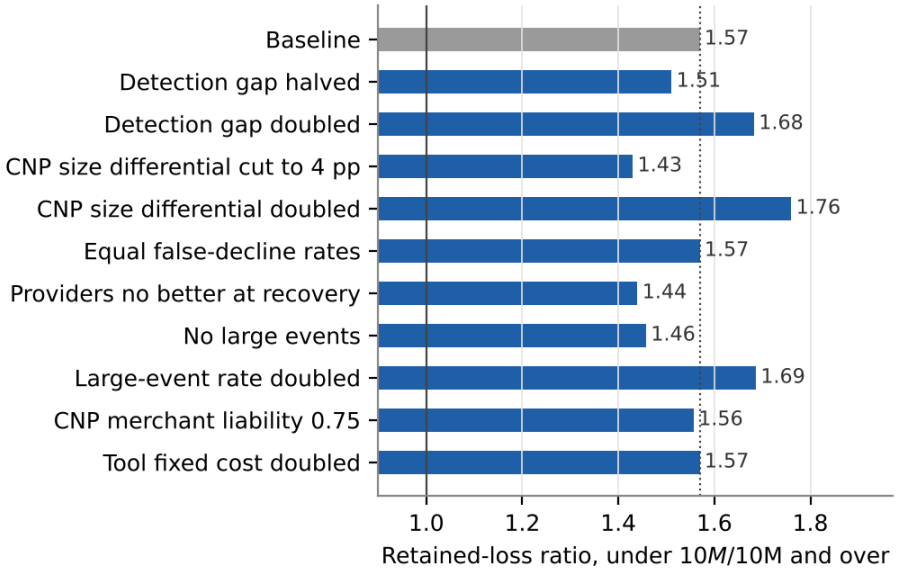


Figure 3. One-at-a-time sensitivity of the pooled gap ratio to the labeled modeling assumptions (seed 0). The dotted line marks the baseline.

5. Verdicts

5.1 The verdict matrix

Table 3. Verdict matrix for the eight era claims.

Claim	Test statistic	Verdict	Post-window check
C1 Gap	minimum year–seed ratio 1.48; pooled 1.57	Qualified	Merchant loss rates in the card-not-present channel rose further to 2023 (Hayashi, 2026)
C2 Channel and liability	channel 36% + EMV lag 2% = 38%	Held	Card-not-present merchant loss rate on single-message networks rose from 5.4 to 12.8 bps, 2021–2023 (Hayashi, 2026)
C3 Recovery	Shapley share 18%	Held	Share of organizations recovering up to 75% rose from 29% to 58% (AFP, 2025)
C4 Tooling	Shapley share 21%; prevention ratio 2.25	Held	—
C5 False declines	Cost ratio 1.69	Held	Era headline estimate (Aite, 2019) did not sample small merchants
C6 Multiplier	2.90 vs 2.62	Held	US retail multiplier $4.61 per $1 (LexisNexis, 2025)
C7 Targeting	Shapley share 6%; gap without targeting 1.53	Qualified	—
C8 Survey measurement	survey/population inflation 1.04	Held	—

5.2 Reading the matrix

C1 is qualified. The direction is unambiguous in every year and seed, but the pre-registered magnitude threshold of 1.5 is not met in every year–seed cell; the minimum is 1.48, in the late years of the window when larger merchants

had themselves moved online. The era's claim is right about sign and roughly right about size, and it understated how quickly the gap was narrowing.

C2 is held, but not for the reason the period gave. The combined share of channel exposure and liability lag is 38%, nearly all of it channel. The window's prose located the mechanism in the liability rule; the audit locates it in exposure to the channel the rule governs. The post-window data sharpen the point: between 2021 and 2023 the merchant-borne card-not-present loss rate on single-message debit networks more than doubled (Hayashi, 2026), which is a channel effect operating through an unchanged rule.

C3 is held with a share of 18%. The 2025 AFP survey reported a sharp improvement in recovery outcomes, which, if it extends to small merchants, would shrink this component going forward.

C4 is held: tooling carries 21% of the gap and the smaller group spends 2.25 times as much of its revenue on prevention. Both halves of the era's claim survive: the fixed-cost burden and the under-adoption, and the sensitivity analysis shows that the tooling share is the one most responsive to how large a detection advantage one believes tooling confers.

C5 is held on the cost ratio of 1.69. The verdict comes with a caveat about the evidence the period relied on: its most-cited false-decline figure described firms of $100 million and above. The small-merchant version of the claim rests on the manual-review pattern in Cyber Source and MRC (2021) and on the assumed rate difference, not on a direct era measurement.

C6 is held. The multiplier is 2.90 for the smaller group and 2.62 for the larger, both above the threshold and ordered as the period claimed, though below the $4-plus figures that LexisNexis reported by the end of the window and the $4.61 reported for 2025. The difference is one of scope: the survey multiplier includes labor, investigation, and customer-relationship costs that the generator does not model.

C7 is qualified. Targeting explains 6% of the gap, and the gap ratio without any targeting is 1.53, far above the C1 threshold. Even granting the period's assumption that attempts against small merchants ran 10–15% higher, the assumption is nearly irrelevant to the size of the gap. The claim was repeated because it is intuitive and commercially useful, not because any era source quantified it.

C8 is held in the narrow sense tested: the simulated survey frame returns a gap within 4% of the value-weighted population gap. The finding should not be read as vindicating survey measurement. Equal weighting, which every survey mean embodies, inflates the gap by 13% relative to the value-weighted figure that a regulator or an acquirer would care about, and selection toward tooled respondents happens to offset part of that inflation. Two measurement artifacts that cancel are not the same as no artifact.

6. Implications

For researchers, the audit suggests retiring two habits of the period. The first is treating the liability rule as the explanatory variable. Because the rule applies identically to merchants of every size, it cannot explain a size gradient; studies that want to explain the gradient should measure channel share, tooling and dispute practice, and treat the rule as the fixed background against which those vary. The second is quoting targeting statistics from vendor surveys as if they measured a mechanism. The audit's targeting parameter is itself generous to the claim, and the claim still fails to matter.

For small merchants and the acquirers and platforms that serve them, the decomposition is a priority list. Channel exposure cannot be reduced without giving up sales, so the actionable components are tooling and recovery, which together account for 38% of the gap in retained loss and a larger share of the gap in total cost. The sensitivity analysis implies that the return to tooling depends on the detection advantage it actually delivers, which argues for platforms bundling detection at the marketplace or processor level, where the fixed cost is spread across thousands of merchants, rather than for each small merchant buying its own.

For policy, the finding that merchants' card-not-present loss rates kept rising after the window (Hayashi, 2026) while the rule stayed fixed suggests that any reallocation of card-not-present liability would reach small merchants in proportion to their channel exposure, which is exactly where the gap sits. The audit does not evaluate whether such a reallocation is desirable; it establishes which lever it would move.

7. Limitations

The study audits claim against a model of the period's assumptions, not against the period's merchants. Where the era corpus supplied a number, the generator uses it; where it did not, a labeled assumption stands in, and the sensitivity analysis shows how much each one matters. The verdicts are therefore conditional on the period having described its own world accurately, which is precisely what claims C8 and the measurement literature question. Several era quantities that matter most, in particular the detection advantage of tooling for a small merchant and the size differential in card-not-present share, were never directly measured in 2017–2024, and the audit cannot supply what the period did not collect.

The decomposition treats five mechanisms as separable parameter blocks. In practice, tooling, provider use and channel strategy are chosen jointly, and a merchant that adopts one is likelier to adopt another. The Shapley value handles interactions among the five but not the selection process behind them. The large-event component is left as a residual rather than decomposed, and its parameters are the least well anchored in the corpus. The survey-frame model is one plausible sampling mechanism among many. Finally, the generator is calibrated to US card payments; the treasury-fraud and occupational-fraud series that supplied the large-event assumptions are international, and the mapping between them is approximate.

8. Conclusion

The small-merchant fraud gap that the 2017–2024 literature described is real under that literature's own assumptions: merchants under $10 million in revenue retain 1.57 times as much fraud loss per dollar of sales as larger merchants, and more once prevention and false declines are counted. The period was right about the gap and wrong about its cause. The liability rules it blamed are a constant, not a variable; the targeting it warned of is nearly weightless; and the mechanisms that carry the gap, channel exposure, tooling and recovery, are the ones the period discussed least precisely. Nothing in the evidence that has appeared since 2024 contradicts that reading. Future work with access to merchant-level processor data could replace the labeled assumptions with measurements and re-run the same audit; the code is released to make that straightforward.

Supplementary Information

“Not applicable”

Ethical Approval and Consent to participate

“Not applicable”

Consent for publication

“Not applicable”

Availability of supporting data

“Not applicable”

Competing interests

“Not applicable”

Funding

“Not applicable”

Authors' contributions

Two authors equally and actively contributed to the conception, design, implementation, data analysis, interpretation, manuscript writing, and submission of this study. No other contributors were involved.

9. Acknowledgements

During the preparation of this manuscript, the author used ChatGPT (OpenAI) and Grammarly to assist with editing and formatting. The author has reviewed and edited all generated content and takes full responsibility for the final version of this publication. The author thanks their colleagues at www.saachistyle.com for insightful discussions and constructive feedback that informed the development of the Small Merchant Fraud Gap.

Appendix A. Corpus verification tags

Each era source carries one tag: [CR] Crossref DOI-checked, [WS] retrieved from the publisher's web page, [no-DOI] preprint without a DOI. Page ranges for journal articles are reported as given by the publishers' records and should be confirmed at proof.

Key	Year	Short reference	Tag
hayashi2018	2018	Hayashi, F., Markiewicz, Z., & Minhas, S., The initial effects of EMV migration on chargebacks in the United Stat…	[CR]
hayashi2019	2019	Hayashi, F., Payment card fraud rates in the United States relative to other countr…	[CR]
hayashi2021	2021	Hayashi, F., Cash or debit cards? Payment acceptance costs for merchants	[CR]
fedboard2021	2023	Board of Governors of the Federal Reserve System, 2021 interchange fee revenue, covered issuer costs, and covered issuer…	[WS]
mekterovic2021	2021	Mekterović, I., Karan, M., & Pintar, D., Credit card fraud detection in card-not-present transactions: Where to…	[CR]
bodker2022	2022	Bodker, A., Connolly, P., Sing, O., Hutchins, B., Townsley, …, Card-not-present fraud: Using crime scripts to inform crime prevention…	[CR]
amasiatu2018	2018	Amasiatu, C. V., & Shah, M. H., First party fraud management: Framework for the retail industry	[CR]

john2020	2020	John, S., Shah, B. J., & Kartha, P., Refund fraud analytics for an online retail purchases	[CR]
ekelund2019	2019	Ekelund, S., & Iskoujina, Z., Cybersecurity economics – balancing operational security spending	[CR]
jevtic2020	2020	Jevtić, P., & Lanchier, N., Dynamic structural percolation model of loss distribution for cyber ri…	[CR]
cross2020	2020	Cross, C., & Gillett, R., Exploiting trust for financial gain: An overview of business email com…	[CR]
kalovya2020	2020	Kalovya, O. Z., Determinants of occupational fraud losses: Offenders, victims and insi…	[CR]
eling2019	2019	Eling, M., & Wirfs, J., What are the actual costs of cyber risk events?	[CR]
kamiya2021	2021	Kamiya, S., Kang, J.-K., Kim, J., Milidonis, A., & Stulz, R.…, Risk management, firm reputation, and the impact of successful cyberat…	[CR]
ponsard2019	2019	Ponsard, C., Grandclaudon, J., & Bal, S., Survey and lessons learned on raising SME awareness about cybersecurit…	[CR]
fernandez2023	2023	Fernandez de Arroyabe, I., Arranz, C. F. A., Arroyabe, M. F.…, Cybersecurity capabilities and cyber-attacks as drivers of investment …	[CR]
schlackl2022	2022	Schlackl, F., Link, N., & Hoehle, H., Antecedents and consequences of data breaches: A systematic review	[CR]

biancotti2017	2017	Biancotti, C., The price of cyber (in)security: Evidence from the Italian private sec…	[CR]
gordon2020	2020	Gordon, L. A., Loeb, M. P., & Zhou, L., Integrating cost–benefit analysis into the NIST Cybersecurity Framewor…	[CR]
woods2021	2021	Woods, D. W., & Böhme, R., SoK: Quantifying cyber risk	[CR]
riek2018	2018	Riek, M., & Böhme, R., The costs of consumer-facing cybercrime: An empirical exploration of m…	[CR]
dalpozzolo2018	2018	Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bon…, Credit card fraud detection: A realistic modeling and a novel learning…	[CR]
carcillo2021	2021	Carcillo, F., Le Borgne, Y.-A., Caelen, O., Kessaci, Y., Obl…, Combining unsupervised and supervised learning in credit card fraud de…	[CR]
romanosky2017	2017	Romanosky, S., Ablon, L., Kuehn, A., & Jones, T., Content analysis of cyber insurance policies: How do carriers write po…	[CR]
bada2019	2019	Bada, M., & Nurse, J. R. C., Developing cybersecurity education and awareness programmes for small-…	[CR]
chidukwani2022	2022	Chidukwani, A., Zander, S., & Koutsakis, P., A survey on the cyber security of small-to-medium businesses: Challeng…	[CR]
kabanda2018	2018	Kabanda, S., Tanner, M., & Kent, C., Exploring SME cybersecurity practices in developing countries	[CR]

benz2020	2020	Benz, M., & Chatterjee, D., Calculated risk? A cybersecurity evaluation tool for SMEs	[CR]
lucas2020	2020	Lucas, Y., Portier, P.-E., Laporte, L., He-Guelton, L., Cael…, Towards automated feature engineering for credit card fraud detection …	[CR]
cherif2023	2023	Cherif, A., Badhib, A., Ammar, H., Alshehri, S., Kalkatawi, …, Credit card fraud detection in the era of disruptive technologies: A s…	[CR]
hilal2022	2022	Hilal, W., Gadsden, S. A., & Yawney, J., Financial fraud: A review of anomaly detection techniques and recent a…	[CR]
kemp2020	2020	Kemp, S., Miró-Llinares, F., & Moneva, A., The dark figure and the cyber fraud rise in Europe: Evidence from Spai…	[CR]
alahmari2020	2020	Alahmari, A., & Duncan, B., Cybersecurity risk management in small and medium-sized enterprises: A…	[CR]
rawindaran2021	2021	Rawindaran, N., Jayal, A., Prakash, E., & Hewage, C., Cost benefits of using machine learning features in NIDS for cyber sec…	[CR]
whitty2019	2019	Whitty, M. T., Predicting susceptibility to cyber-fraud victimhood	[CR]
lopezrojas2018	2018	Lopez-Rojas, E. A., Axelsson, S., & Baca, D., Analysis of fraud controls using the PaySim financial simulator	[CR]
makki2019	2019	Makki, S., Assaghir, Z., Taher, Y., Haque, R., Hacid, M.-S.,…, An experimental study with imbalanced	[CR]

		classification approaches for cr…	
fiore2019	2019	Fiore, U., De Santis, A., Perla, F., Zanetti, P., & Palmieri…, Using generative adversarial networks for improving classification eff…	[CR]
grover2022	2022	Grover, P., Xu, J., Tittelfitz, J., Cheng, A., Li, Z., Zablo…, Fraud dataset benchmark and applications	[no-DOI]
vanbelle2021	2021	Tax, N., de Vries, K. J., de Jong, M., Dosoula, N., van den …, Machine learning for fraud detection in e-commerce: A research agenda	[no-DOI]
lexis2017	2017	LexisNexis Risk Solutions, 2017 True Cost of Fraud study: Overall report	[WS]
lexis2019	2019	LexisNexis Risk Solutions, 2019 True Cost of Fraud study: E-commerce/retail edition (press releas…	[WS]
lexis2020	2020	LexisNexis Risk Solutions, 2020 True Cost of Fraud study: E-commerce/retail report, US & Canada e…	[WS]
lexis2022	2022	LexisNexis Risk Solutions, 13th annual True Cost of Fraud study: Retail and ecommerce (press rele…	[WS]
afp2022	2022	Association for Financial Professionals, 2022 AFP Payments Fraud and Control Report: Key highlights	[WS]
afp2023	2023	Association for Financial Professionals, 2023 AFP Payments Fraud and Control Survey Report: Key highlights	[WS]
afp2024	2024	Association for Financial Professionals, 2024 AFP Payments Fraud and Control	[WS]

		Survey Report: Key highlights	
cybersource2021	2021	Cybersource & Merchant Risk Council, 2021 Global Fraud Report	[WS]
mrc2024	2024	Merchant Risk Council, Visa Acceptance Solutions, & Verifi, 2024 Global eCommerce Payments & Fraud Report (25th ed.)	[WS]
datos2024	2024	Datos Insights, E-commerce fraud landscape and trends: Merchants seeking to adapt	[WS]
aite2019	2019	Aite Group (commissioned by ClearSale), The e-commerce conundrum: Balancing false declines and fraud preventio…	[WS]
cb911_2023	2023	Chargebacks911, 2023 Chargeback Field Report	[WS]
acfe2022	2022	Association of Certified Fraud Examiners, Occupational fraud 2022: A report to the nations	[WS]
acfe2024	2024	Association of Certified Fraud Examiners, Occupational fraud 2024: A report to the nations	[WS]
forrester2015	2015	Forrester Research, New fraud targets: The impacts of EMV chip card migration brings new c…	[WS]

Appendix B. Generator parameters

Values are those used in the baseline run; [A] marks a modelling assumption without a single era source. Keys refer to Appendix A.

Parameter	Value	Source
Band shares (micro / small / mid / large)	0.58 / 0.30 / 0.08 / 0.04	[A]
Tool adoption probability	0.35 / 0.55 / 0.85 / 0.95	cybersource2021; fernandez2023; chidukwani2022; [A] levels
Representment-provider use	0.15 / 0.30 / 0.60 / 0.80	cb911_2023; mrc2024; [A] gradient
EMV adoption 2017 → 2024	0.35→0.78 / 0.50→0.86 / 0.75→0.95 / 0.90→0.98	hayashi2018; forrester2015; [A] gradient
Card-not-present share, band mean	0.52 / 0.50 / 0.44 / 0.40, +1.5 pp per year, +8 pp from 2020	fedboard2021; hayashi2019; [A] differential
Targeting multiplier on attempts	1.15 / 1.10 / 1.00 / 1.00	lexis2019

CNP attempt rate (bps of CNP value)	60 + 4 per year after 2017	hayashi2019; hayashi2025; [A] level
CP attempt rate (bps of CP value)	9 − 0.5 per year after 2017	hayashi2018
Counterfeit share of CP fraud	0.55	[A]
Detection rate with / without tooling	0.78 / 0.55	dalpozzolo2018; carcillo2021; cherif2023; [A] levels
False-decline rate with / without tooling	0.012 / 0.026	datos2024; cybersource2021
Gross margin applied to false declines	0.32	[A]
Average order value	$85	[A]
CNP merchant liability probability	0.92	hayashi2018; [A] residual
Chargeback fee	$25	[A]
Friendly-fraud share of CNP orders	0.006	cybersource2021 (halved, [A])
Dispute win rate, in-house / provider	0.19 / 0.30	mrc2024; cb911_2023
Representment probability, in-house/provider	0.35 / 0.85	[A]
True-fraud dispute win rate	0.05	[A]
Tool fixed cost / variable rate	$12,000 per year / 0.15% of CNP value	biancotti2017; ekelund2019; [A] levels
Manual-review share with/without tooling; cost per review	0.15 / 0.25; $1.50	cybersource2021; [A] cost
Large-event annual probability	0.025 / 0.045 / 0.08 / 0.14	afp2023; afp2024; [A] levels
Large-event severity	log-normal, median $35,000 at $1M revenue, exponent 0.30, σ 1.1	eling2019; acfe2022; acfe2024; [A]
Large-event recovery mixture	30% none; 29% partial; 41% ≥ 75%	afp2024

Appendix C. Reproducibility package

`generator.py` (seeded data-generating process), `analysis.py` (Shapley decomposition, survey frame, verdict rules), `sensitivity.py`, `figures.py`, `write_manuscript.py` (reads results and writes this text), `results/results.json`, `results/sensitivity.json`, `corpus/corpus.json` (verified bibliography with tags and the novelty-sweep record). Python 3 with NumPy and Matplotlib. Baseline seeds 0–9; sensitivity at seed 0.

REFERENCES

1. Aite Group. (2019). *The e-commerce conundrum: Balancing false declines and fraud prevention* (commissioned by ClearSale). Aite Group.
2. Alahmari, A., & Duncan, B. (2020). Cybersecurity risk management in small and medium-sized enterprises: A systematic review of recent evidence. *2020 International Conference on Cyber Situational Awareness, Data Analytics and Assessment (CyberSA)*, 1–5. https://doi.org/10.1109/CyberSA49311.2020.9139638
3. Amasiatu, C. V., & Shah, M. H. (2018). First party fraud management: Framework for the retail industry. *International Journal of Retail & Distribution Management, 46*(4), 350–363. https://doi.org/10.1108/IJRDM-10-2016-0185
4. Association for Financial Professionals. (2022). *2022 AFP Payments Fraud and Control Report: Key highlights*. AFP / J.P. Morgan.
5. Association for Financial Professionals. (2023). *2023 AFP Payments Fraud and Control Survey Report: Key highlights*. AFP / J.P. Morgan.
6. Association for Financial Professionals. (2024). *2024 AFP Payments Fraud and Control Survey Report: Key highlights*. AFP / Truist.
7. Association for Financial Professionals. (2025). *2025 AFP Payments Fraud and Control Survey Report: Key highlights*. AFP / Truist. [Post-window source]
8. Association of Certified Fraud Examiners. (2022). *Occupational fraud 2022: A report to the nations*. ACFE.
9. Association of Certified Fraud Examiners. (2024). *Occupational fraud 2024: A report to the nations*. ACFE.
10. Bada, M., & Nurse, J. R. C. (2019). Developing cybersecurity education and awareness programmes for small- and medium-sized enterprises (SMEs). *Information & Computer Security, 27*(3), 393–410. https://doi.org/10.1108/ICS-07-2018-0080
11. Benz, M., & Chatterjee, D. (2020). Calculated risk? A cybersecurity evaluation tool for SMEs. *Business Horizons, 63*(4), 531–540. https://doi.org/10.1016/j.bushor.2020.03.010
12. Biancotti, C. (2017). *The price of cyber (in)security: Evidence from the Italian private sector* (Bank of Italy Occasional Paper No. 407). https://doi.org/10.2139/ssrn.3082195

13. Board of Governors of the Federal Reserve System. (2023). *2021 interchange fee revenue, covered issuer costs, and covered issuer and merchant fraud losses associated with debit card transactions* (revised November 7, 2023). https://www.federalreserve.gov/paymentsystems/2021-Interchange-Fee.htm
14. Bodker, A., Connolly, P., Sing, O., Hutchins, B., Townsley, M., & Drew, J. (2022). Card-not-present fraud: Using crime scripts to inform crime prevention initiatives. *Security Journal, 36*(4), 693–711. https://doi.org/10.1057/s41284-022-00359-w
15. Carcillo, F., Le Borgne, Y.-A., Caelen, O., Kessaci, Y., Oblé, F., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences, 557*, 317–331. https://doi.org/10.1016/j.ins.2019.05.042
16. Chargebacks911. (2023). *2023 Chargeback Field Report*. Chargebacks911.
17. Cherif, A., Badhib, A., Ammar, H., Alshehri, S., Kalkatawi, M., & Imine, A. (2023). Credit card fraud detection in the era of disruptive technologies: A systematic review. *Journal of King Saud University – Computer and Information Sciences, 35*(1), 145–174. https://doi.org/10.1016/j.jksuci.2022.11.008
18. Chidukwani, A., Zander, S., & Koutsakis, P. (2022). A survey on the cyber security of small-to-medium businesses: Challenges, research focus and recommendations. *IEEE Access, 10*, 85701–85719. https://doi.org/10.1109/ACCESS.2022.3197899
19. Cross, C., & Gillett, R. (2020). Exploiting trust for financial gain: An overview of business email compromise (BEC) fraud. *Journal of Financial Crime, 27*(3), 871–884. https://doi.org/10.1108/JFC-02-2020-0026
20. Cybersource, & Merchant Risk Council. (2021). *2021 Global Fraud Report*. Cybersource.
21. Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2018). Credit card fraud detection: A realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems, 29*(8), 3784–3797. https://doi.org/10.1109/TNNLS.2017.2736643
22. Datos Insights. (2024). *E-commerce fraud landscape and trends: Merchants seeking to adapt*. Datos Insights.
23. Ekelund, S., & Iskoujina, Z. (2019). Cybersecurity economics – balancing operational security spending. *Information Technology & People, 32*(5), 1318–1342. https://doi.org/10.1108/ITP-05-2018-0252
24. Eling, M., & Wirfs, J. (2019). What are the actual costs of cyber risk events? *European Journal of Operational Research, 272*(3), 1109–1119. https://doi.org/10.1016/j.ejor.2018.07.021
25. Fernandez de Arroyabe, I., Arranz, C. F. A., Arroyabe, M. F., & Fernandez de Arroyabe, J. C. (2023). Cybersecurity capabilities and cyber-attacks as drivers of investment in cybersecurity systems: A UK survey for 2018 and 2019. *Computers & Security, 124*, 102954. https://doi.org/10.1016/j.cose.2022.102954
26. Fiore, U., De Santis, A., Perla, F., Zanetti, P., & Palmieri, F. (2019). Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. *Information Sciences, 479*, 448–455. https://doi.org/10.1016/j.ins.2017.12.030
27. Forrester Research. (2015, October 9). *New fraud targets: The impacts of EMV chip card migration brings new challenges* [Blog post]. [Pre-window source, cited for the liability-shift mechanism only]
28. Gordon, L. A., Loeb, M. P., & Zhou, L. (2020). Integrating cost–benefit analysis into the NIST Cybersecurity Framework via the Gordon–Loeb Model. *Journal of Cybersecurity, 6*(1), tyaa005. https://doi.org/10.1093/cybsec/tyaa005
29. Grover, P., Xu, J., Tittelfitz, J., Cheng, A., Li, Z., Zablocki, J., Liu, J., & Zhou, H. (2022). *Fraud dataset benchmark and applications*. arXiv:2208.14417.
30. Hayashi, F. (2019). Payment card fraud rates in the United States relative to other countries since migrating to chip cards. *Federal Reserve Bank of Kansas City Economic Review, 104*(4). https://doi.org/10.18651/ER/4q19Hayashi
31. Hayashi, F. (2021). Cash or debit cards? Payment acceptance costs for merchants. *Federal Reserve Bank of Kansas City Economic Review, 106*(3). https://doi.org/10.18651/ER/v106n3Hayashi
32. Hayashi, F. (2025, May). *Card-not-present fraud rates in the United States after the migration to chip cards* (Payments System Research Briefing). Federal Reserve Bank of Kansas City. [Post-window source]
33. Hayashi, F. (2026, February 25). *New data on card-present and card-not-present fraud rates in the United States* (Payments System Research Briefing). Federal Reserve Bank of Kansas City. [Post-window source]
34. Hayashi, F., Markiewicz, Z., & Minhas, S. (2018). *The initial effects of EMV migration on chargebacks in the United States* (Research Working Paper 18-10). Federal Reserve Bank of Kansas City. https://doi.org/10.18651/RWP2018-10
35. Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications, 193*, 116429. https://doi.org/10.1016/j.eswa.2021.116429
36. Jevtić, P., & Lanchier, N. (2020). Dynamic structural percolation model of loss distribution for cyber risk of small and medium-sized enterprises for tree-based LAN topology. *Insurance: Mathematics and Economics, 91*, 209–223. https://doi.org/10.1016/j.insmatheco.2020.02.005
37. John, S., Shah, B. J., & Kartha, P. (2020). Refund fraud analytics for an online retail purchases. *Journal of Business Analytics, 3*(1), 56–66. https://doi.org/10.1080/2573234X.2020.1776164
38. Kabanda, S., Tanner, M., & Kent, C. (2018). Exploring SME cybersecurity practices in developing countries. *Journal of Organizational Computing and Electronic Commerce, 28*(3), 269–282. https://doi.org/10.1080/10919392.2018.1484598
39. Kalovya, O. Z. (2020). Determinants of occupational fraud losses: Offenders, victims and insights from fraud theory. *Journal of Financial Crime, 30*(2), 361–376. https://doi.org/10.1108/JFC-10-2019-0136

40. Kamiya, S., Kang, J.-K., Kim, J., Milidonis, A., & Stulz, R. M. (2021). Risk management, firm reputation, and the impact of successful cyberattacks on target firms. *Journal of Financial Economics, 139*(3), 719–749. https://doi.org/10.1016/j.jfineco.2019.05.019
41. Kemp, S., Miró-Llinares, F., & Moneva, A. (2020). The dark figure and the cyber fraud rise in Europe: Evidence from Spain. *European Journal on Criminal Policy and Research, 26*(3), 293–312. https://doi.org/10.1007/s10610-020-09439-2
42. LexisNexis Risk Solutions. (2017). *2017 True Cost of Fraud study: Overall report*. LexisNexis Risk Solutions.
43. LexisNexis Risk Solutions. (2019, September 24). *2019 True Cost of Fraud study: E-commerce/retail edition* [Press release]. LexisNexis Risk Solutions.
44. LexisNexis Risk Solutions. (2020). *2020 True Cost of Fraud study: E-commerce/retail report, US & Canada edition*. LexisNexis Risk Solutions.
45. LexisNexis Risk Solutions. (2022, August 2). *13th annual True Cost of Fraud study: Retail and ecommerce* [Press release]. LexisNexis Risk Solutions.
46. LexisNexis Risk Solutions. (2025, April 2). *True Cost of Fraud study: Ecommerce and retail report, US and Canada edition* [Press release]. LexisNexis Risk Solutions. [Post-window source]
47. Lopez-Rojas, E. A., Axelsson, S., & Baca, D. (2018). Analysis of fraud controls using the PaySim financial simulator. *International Journal of Simulation and Process Modelling, 13*(4), 377. https://doi.org/10.1504/IJSPM.2018.093756
48. Lucas, Y., Portier, P.-E., Laporte, L., He-Guelton, L., Caelen, O., Granitzer, M., & Calabretto, S. (2020). Towards automated feature engineering for credit card fraud detection using multi-perspective HMMs. *Future Generation Computer Systems, 102*, 393–402. https://doi.org/10.1016/j.future.2019.08.029
49. Makki, S., Assaghir, Z., Taher, Y., Haque, R., Hacid, M.-S., & Zeineddine, H. (2019). An experimental study with imbalanced classification approaches for credit card fraud detection. *IEEE Access, 7*, 93010–93022. https://doi.org/10.1109/ACCESS.2019.2927266
50. Mekterović, I., Karan, M., & Pintar, D. (2021). Credit card fraud detection in card-not-present transactions: Where to invest? *Applied Sciences, 11*(15), 6766. https://doi.org/10.3390/app11156766
51. Merchant Risk Council, Visa Acceptance Solutions, & Verifi. (2024). *2024 Global eCommerce Payments & Fraud Report* (25th ed.). Merchant Risk Council.
52. Ponsard, C., Grandclaudon, J., & Bal, S. (2019). Survey and lessons learned on raising SME awareness about cybersecurity. *Proceedings of the 5th International Conference on Information Systems Security and Privacy (ICISSP)*, 558–563. https://doi.org/10.5220/0007574305580563
53. Rawindaran, N., Jayal, A., Prakash, E., & Hewage, C. (2021). Cost benefits of using machine learning features in NIDS for cyber security in UK small medium enterprises (SME). *Future Internet, 13*(8), 186. https://doi.org/10.3390/fi13080186
54. Riek, M., & Böhme, R. (2018). The costs of consumer-facing cybercrime: An empirical exploration of measurement issues and estimates. *Journal of Cybersecurity, 4*(1), tyy004. https://doi.org/10.1093/cybsec/tyy004
55. Romanosky, S., Ablon, L., Kuehn, A., & Jones, T. (2017). *Content analysis of cyber insurance policies: How do carriers write policies and price cyber risk?* (RAND Working Paper WR-1208). https://doi.org/10.7249/WR1208
56. Schlackl, F., Link, N., & Hoehle, H. (2022). Antecedents and consequences of data breaches: A systematic review. *Information & Management, 59*(4), 103638. https://doi.org/10.1016/j.im.2022.103638
57. Tax, N., de Vries, K. J., de Jong, M., Dosoula, N., van den Akker, B., Smith, J., Thuong, O., & Bernardi, L. (2021). *Machine learning for fraud detection in e-commerce: A research agenda*. arXiv:2107.01979.
58. Whitty, M. T. (2019). Predicting susceptibility to cyber-fraud victimhood. *Journal of Financial Crime, 26*(1), 277–292. https://doi.org/10.1108/JFC-10-2017-0095
59. Woods, D. W., & Böhme, R. (2021). SoK: Quantifying cyber risk. *2021 IEEE Symposium on Security and Privacy*, 211–228. https://doi.org/10.1109/SP40001.2021.00053