"We Have Never Been Attacked": Unpacking the Illusion of Cybersecurity in Small Business

Chi Zhang
The Graduate Center
The City University of New York
CHI.ZHANG06@gc.cuny.edu

Ping Ji
Hunter College and The Graduate Center
The City University of New York
ping.ji@hunter.cuny.edu

Zhengfan Zhang
New York University
zz4657@nyu.edu

Abstract

Small businesses are increasingly becoming prime targets for cyberattacks, yet many remain poorly prepared. The ITRC's 2025 survey found that 81% of surveyed small businesses reported a security breach, data breach, or both in the previous year [6]. The situation is further worsened by the rapid integration of agentic AI tools into nearly every aspect of business operations, which expands the attack surface and introduces new cybersecurity risks that many small businesses are not yet equipped to manage. At the same time, important questions remain: How much do business owners actually know about these risks, and how well equipped are small businesses to defend themselves against cybersecurity threats? In this work, we investigate why many small-business stakeholders continue to believe they have "never been attacked" despite mounting evidence of widespread cybersecurity exposure. Drawing on a multi-perspective study of 128 main-study participants, complemented by two supplemental AI-startup cases excluded from that count, we characterize this disconnect as a **security illusion**: a perception of safety sustained by limited detection rather than demonstrated defense. **We identify three reinforcing mentalities for this illusion: semantic narrowing, structural invisibility, and delegated trust**, and argue that effective small-business cybersecurity support should emphasize lightweight visibility tools, clearer responsibility boundaries, and incident narratives that connect cyber threats to everyday business consequences.

1 Motivation

The same report found that more than half of affected businesses lost $250,000–$1 million, while 9% lost more than $1 million [6]. For small businesses, Verizon's 2026 DBIR identifies system intrusion, web-application attacks, and social engineering as the leading breach patterns; 55% involved third parties and 45% involved a human element [9]. With the introduction of agentic AI tools, the situation gets worse. While small businesses are often treated as scaled-down enterprises or as individual users with commercial stakes, neither framing fits the AI-era organizations we increasingly observe. Recent AI-economy reporting documents rapid growth in AI investment, organizational adoption, and newly funded AI companies [1]. At the same time, platform services lower the barrier to building: a founder can move from idea to technical product by combining code assistants, cloud databases, hosted identity, deployment platforms, chat integrations, and domain-specific AI workflows. Technical intensity is therefore no longer limited to teams with mature engineering or security functions. The result is enterprise-like exposure without enterprise-like staffing, monitoring, compliance capacity, or incident response maturity.

Prior work shows that small businesses face recurring security obstacles, limited resources, and substantial incident exposure [4, 5, 10]. Yet in our fieldwork, many owners and employees described their organizations as having never experienced an attack. Rather than treating these statements as simple error, we ask what organizational conditions make "no attacks" feel plausible.

2 Data Collection and Interview Categories

We conducted a mixed-methods study focused on small businesses with 1–500 employees. The design intentionally includes multiple organizational perspectives and separates three data collection channels to avoid conflating populations. The study received HRPP/IRB approval at The City University of New York; see Appendix A.

Table 1 summarizes the three main-study data-collection channels. The ten semi-structured interview participants were recruited through convenience sampling from the research team's existing personal and professional networks; all had direct startup or small-business experience. These interviews lasted 30–45 minutes. Survey participants were recruited through online outreach to maker communities in New York, California, and elsewhere. The 102 intercept participants were early-stage startup founders from multiple countries approached at a major international technology expo; each structured conversation lasted approximately ten minutes.

Table 1: **Data Sources and Study Focus**

Source	N	Focus
Semi-structured interviews	10	8 owner/employer and 2 employee interviews; reasoning, definitions, practices
Owner survey	7	owner posture, investment, incidents
Employee survey	9	employee awareness, reporting, training
On-site intercept interviews	102	early-stage startup founders; brief structured conversations
Total participants	**128**	

We did not impose a formal definition of "cyberattack" because the goal was to understand participants' own interpretations of cybersecurity risks. Qualitative data were analyzed using reflexive thematic analysis [2, 3]. Survey results are treated as descriptive and exploratory, not statistically representative. The two supplemental AI-startup cases are excluded from Table 1 and N=128 and are discussed separately in Section 4.

3 Three Mentalities of the Security Illusion

From the interviews, we identified three reinforcing mentalities underlying the security illusion.

I. Semantic narrowing: Participants often only consider "attacks" for catastrophic events such as ransomware, visible system failure, or major financial theft. Routine events such as phishing, suspicious emails, credential risk, or strange login activity were often treated as background noise rather than incidents. Thus, self-reports of no attacks may mean that no catastrophic incidents were detected, not that no cybersecurity-relevant activity occurred.

II. Structural invisibility: Many organizations lacked monitoring, logging, alerting, or evidence collection practices that would allow them to detect routine threats. In these settings, the absence of alerts can be misread as evidence of safety. This does not mean commodity infrastructure cannot support logging; rather, such capabilities were rarely enabled, maintained, or interpreted by the organizations we studied.

III. Delegated trust: Participants frequently assumed that SaaS, cloud platforms, managed service providers, payment processors, or identity tools handled security comprehensively. Delegation can provide real protection, but it also obscures what remains the organization's responsibility, especially for logs, evidence retention, access review, and incident response.

These mentalities reinforce one another. Narrow definitions make routine threats easier to ignore; absent visibility produces no contradicting evidence; and delegated trust reduces the perceived need to build internal security capacity. The result is a *security illusion*: "we have never been attacked" becomes a statement about what the business can notice, not necessarily what has happened.

4 Two Supplemental AI-Startup Cases

We also conducted two supplemental AI-startup case studies to deepen and contextualize our observations. We anonymize both companies and analyze them separately from the main study (N=128); each case involved about one week of direct founder engagement.

Case 1: regulated-document AI company. The first case is a five-person seed-stage company processing regulated documents and sensitive health data through cloud storage, application services, and an external AI service. Its stack included GitHub, hosted identity, deployment secrets, isolated sandboxes, and encrypted file flows; security work centered on the founder, a part-time CTO/advisor, and an intern. Compliance review prompted access controls, audit logging, and expert validation; gaps remained in de-identification, external-AI data controls, and recovery planning.

Case 2: AI agent-evaluation company. The second case is a three-person company evaluating AI agents and model behavior with public benchmarks, open-source data, and private business data across code repositories, cloud services, and chat integrations. Manual controls and platform defaults substituted for formal monitoring and access governance.

These cases are not prevalence evidence. They are boundary cases showing why small businesses deserve renewed attention: AI-native micro-companies can handle regulated documents, user data, code, model outputs, and agent workflows before they can hire dedicated security staff. Therefore, while AI guidance identifies risks such as data leakage, prompt injection, insecure tool use, and governance gaps [7,8], **our contribution is to show how these risks meet ordinary small-business constraints**.

5 Implications

Small-business security support should focus less on enterprise-style dashboards and more on capacity-appropriate defaults: basic audit logging, clear suspicious-activity signals, plain-language shared-responsibility boundaries, secret scanning, recoverable backups, and exportable evidence for compliance or incident review. Across the study, security action was often delayed until an external trigger made it legible as business risk: funding conversations, investor diligence, enterprise customer requirements, or SOC 2 certification pressure. Training should map incidents to business consequences such as invoice fraud, account takeover, lost appointments, customer trust, or blocked enterprise sales.

6 Poster Contribution

The poster presents the security illusion model and uses AI-native micro-startups as an emerging stress test. Its central claim is that small-business cybersecurity is not merely an under-resourced version of enterprise security: it is a setting where limited visibility, narrow incident vocabulary, and platform dependence jointly shape what organizations believe has happened to them.

Appendix A: Ethics, Scope, and Data Handling

This study was reviewed and approved by the Human Research Protection Program (HRPP)/IRB at The City University of New York. Participants provided informed consent; survey and interview findings were reported in de-identified form; and intercept interviews collected no identifying information. Because some data were collected through brief conversations at a major international technology expo and all incident histories are self-reported, we interpret claims of "never attacked" cautiously and do not treat employee observations as confirmed incidents. The two AI-startup cases were analyzed separately to clarify contemporary risk conditions; they are not counted in the main-study participant total and are not prevalence evidence.

References

[1] AI Index Steering Committee. AI Index Report 2026: Chapter 4, economy. Technical report, Stanford Institute for Human-Centered Artificial Intelligence, 2026. Accessed May 28, 2026.

[2] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101, January 2006.

[3] Virginia Braun and Victoria Clarke. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health*, 11(4):589–597, August 2019.

[4] Alladean Chidukwani, Sebastian Zander, and Polychronis Koutsakis. A survey on the cyber security of small-to-medium businesses: challenges, research focus and recommendations. *IEEE Access*, 10:85701–85719, 2022.

[5] Nicolas Huaman, Bennet von Skarczinski, Christian Stransky, Dominik Wermke, Yasemin Acar, Arne Dreißigacker, and Sascha Fahl. A Large-Scale interview study on information security in and attacks against small and medium-sized enterprises. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1235–1252. USENIX Association, August 2021.

[6] Identity Theft Resource Center. 2025 business impact report, December 2025. Accessed May 28, 2026.

[7] National Institute of Standards and Technology. Artificial intelligence risk management framework: generative artificial intelligence profile. Technical Report NIST AI 600-1, National Institute of Standards and Technology (U.S.), Gaithersburg, MD, July 2024.

[8] OWASP Foundation. OWASP top 10 for LLM applications 2025. https://genai.owasp.org/resource/owasp-top-10-for-llm-applications-2025/, 2024. Published November 17, 2024; accessed May 28, 2026.

[9] Verizon. 2026 data breach investigations report: Executive summary, 2026. Published May 19, 2026; accessed July 17, 2026.

[10] Flynn Wolf, Adam J. Aviv, and Ravi Kuber. Security obstacles and motivations for small businesses from a CISO's perspective. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1199–1216. USENIX Association, August 2021.