

“Send to which account?”

Evaluation of an LLM-based Scambaiting System

Hossein Siadati*, Haadi Jafarian†, Sima Jafarikhah‡

*AI Research, Cybera Global Inc./ UNCW, Wilmington, USA
s.h.siadaty@gmail.com

†Department of Computer Science and Engineering, UC Denver, USA
haadi.jafarian@ucdenver.edu

‡Department of Computer Science, UNCW, Wilmington, USA
jafarikhah@uncw.edu

Abstract—Scammers are increasingly harnessing generative AI (GenAI) technologies to produce convincing phishing content at scale, amplifying financial fraud and undermining public trust. While conventional defenses, such as detection algorithms, user training, and reactive takedown efforts remain important, they often fall short in dismantling the infrastructure scammers depend on, including mule bank accounts and cryptocurrency wallets. To bridge this gap, a proactive and emerging strategy involves using conversational honeypots to engage scammers and extract actionable threat intelligence.

This paper presents the first large-scale, real-world evaluation of a scambaiting system powered by large language models (LLMs). Over a five-month deployment, the system initiated over 2,600 engagements with actual scammers, resulting in a dataset of more than 18,700 messages. It achieved an Information Disclosure Rate (IDR) of approximately 32%, successfully extracting sensitive financial information such as mule accounts. Additionally, the system maintained a Human Acceptance Rate (HAR) of around 70%, indicating strong alignment between LLM-generated responses and human operator preferences. Alongside these successes, our analysis reveals key operational challenges. In particular, the system struggled with engagement takeoff: only 48.7% of scammers responded to the initial seed message sent by defenders. These findings highlight the need for further refinement and provide actionable insights for advancing the design of automated scambaiting systems.

Index Terms—Scambaiting, Large Language Models, Conversational Honeypots, Financial Fraud, Generative AI, Autonomous Cyber Defense

I. INTRODUCTION

Fraud losses are escalating at an unprecedented pace. According to Federal Trade Commission (FTC) data, U.S. consumers reported more than \$12.5 billion in losses in 2024, a striking 25% increase from the previous year. The proportion of victims incurring financial loss has also surged: 38% of individuals who reported fraud in 2024 suffered a monetary loss, an 11% rise over 2023. Notably, Authorized Push Payment (APP) fraud, where victims are deceived into voluntarily sending funds through real-time payment channels, is projected to exceed \$3.03 billion by 2027, up from \$1.94 billion in 2022 [1]. Bank transfers or cryptocurrency now account for greater

total losses than all other payment methods combined [2]. These trends have imposed substantial harm on victims and contributed to a growing erosion of trust in institutions. The widespread adoption of GenAI tools has significantly amplified the phishing threat landscape, adding to the complexity of existing defenses. These tools enable scammers to effortlessly craft highly persuasive phishing messages across multiple languages, personas, and contexts, dramatically lowering the barrier to large-scale manipulation of potential victims [3]–[6].

A broad spectrum of defenses exists to counter scams, ranging from detection mechanisms to user education initiatives. Among these, detecting and blocking scam messages remains a primary defense strategy; however, it often resembles a game of whack-a-mole, as scammers rapidly adapt their tactics. This dynamic has further intensified with the growing use of AI by both attackers and defenders, each seeking to outmaneuver the other. While tools that help users make informed decisions during suspicious interactions offer some degree of protection, they too are entangled in this adversarial cycle.

Banks rely on transaction monitoring (e.g., AML systems) and behavioral analysis to detect mule accounts, supplemented by crowdsourced, dark-web, and threat-feed intelligence that exposes bank accounts, crypto wallets, and IP addresses. Yet these sources often yield broad, delayed, or incomplete data. In particular, they remain ineffective for APP scams, where victims voluntarily initiate transfers and early transactions display few actionable fraud indicators. A more strategic defense focuses on disrupting the financial infrastructure that enables scam monetization, including mule bank accounts and cryptocurrency wallets [7], [8]. Identifying these channels in advance allows defenders to intervene *before* fraud occurs, making early discovery of such infrastructure a critical first step.

Scambaiting refers to the proactive engagement with scammers by impersonating victims through deceptive engagements, with the goals of wasting the attacker’s time, disrupting their operations, and gathering threat intelligence. Scambaiting technique has shown promise in generating timely and actionable intelligence [9], [10]. Once a manual and time-consuming

task, scambaiting is increasingly becoming automated. Recent research highlights the potential of LLM-based agents to autonomously converse with scammers [11], [12]. In this context, a Gen AI model (e.g., ChatGPT) can be prompted to produce responses that convincingly impersonate a potential victim. When carefully constructed, such interactions can generate valuable disclosures from scammers, such as bank account details or cryptocurrency wallet addresses. This information can then be used to proactively intervene in scam operations, for example, by flagging and disabling the identified financial accounts through collaboration and coordination with financial institutions.

Despite growing interest in LLM-driven security applications, measuring the effectiveness of LLM-powered scambaiting for scam intelligence collection remains largely unexplored. In this paper, we present the first large-scale evaluation of a real-world deployment of such a system. This paper makes several important contributions to the study of LLM-powered adversarial dialogue systems:

- **Large-scale analysis of real-world scammer interactions.** To our knowledge, this is the first large-scale evaluation of an operational LLM-based scambaiting system engaging real-world scammers over a five-month deployment. The resulting dataset includes 18,797 messages across 2,638 attacker-defender engagements initiated by human adversaries.
- **Comprehensive evaluation framework.** We designed and implemented a set of specialized metrics for assessing automated scambaiting systems. These include the *Information Disclosure Rate (IDR)*, *Information Disclosure Speed (IDS)*, and *Message Freshness*, which respectively measure the effectiveness, efficiency, and linguistic variability of baiting engagements.
- **Actionable insights for real-world deployment.** Our empirical findings yield practical guidance for deploying LLM-driven proactive security tools. These include engagement death thresholds, engagement design strategies that increase elicitation success, and observations on the impact of human-in-the-loop interventions on responsiveness and outcomes.

The remainder of this paper is organized as follows: Section II reviews relevant background and related work in scambaiting and adversarial LLM dialogue. Section III describes our system architecture and deployment methodology. Section IV details the data collection and annotation process. Section V presents our core evaluation metrics and empirical results, including responsiveness, disclosure patterns, and human-in-the-loop impact. Section VI presents a comprehensive and in-depth analysis, uncovering key patterns and insights. Section VII discusses key operational insights and limitations. Finally, Section VIII concludes with takeaways and future research directions.

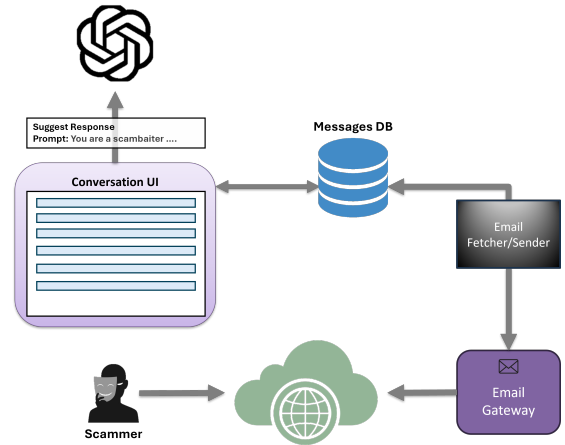


Fig. 1: System architecture of the scambaiting platform that collected data for this study. The Conversations UI displays scammer interactions and sends prompts (e.g., “You are a scambaiter..”) to the LLM (ChatGPT) to suggest responses. Approved replies are routed through the email gateway to scammers. All messages are stored in a central Messages DB and managed by an email fetcher/sender service. The system facilitates live engagement with scammers, with support for human-in-the-loop participation and oversight.

II. BACKGROUND AND RELATED WORK

A. LLMs for Scambaiting

Early scambaiting approaches primarily relied on scripted or rule-based systems. For instance, Netsafe’s *Re:Scam* [13], [15] responded to scam emails using fixed conversational templates, while the voice-based “Lenny” chatbot stalled phone scammers using repetitive prerecorded audio clips [15]. Although effective at confusing human attackers, these systems suffered from limited adaptability and linguistic variation due to their static designs.

Recent advances in GenAI and LLMs have enabled more flexible and realistic interactions. Netsafe’s *Re:Scam 2.0* [14] leveraged LLMs to diversify victim personas and improve conversational coherence. These systems marked a transition from rule-based methods to generative, persona-rich engagements.

Charnsethikul et al.’s *Puppeteer* [12] combined LLM-generated messages with finite-state dialogue management to simulate victims of Amazon refund scams. In controlled evaluations, it achieved a 68% Extraction Score for sensitive information (i.e., shipping address in Amazon customer service scam), outperforming unstructured LLMs (50%), and demonstrating the benefits of lightweight conversational control. In parallel, Basta et al.’s *Bot Wars Evolved* [11] explored adversarial self-play between LLM-based attacker and defender agents. By incorporating role-specific objectives and chain-of-thought prompting, their system sustained coherent, strategic dialogues that mirrored real scambaiting data.

Beyond engagement, GenAI is also being used for predictive defenses. ScamGPT-J [16], for example, anticipates scammer

System / Paper	Success Metric(s)	Type	Description
Puppeteer [12]	Information Disclosure Rate	Functional	Percentage of interactions where scammers disclosed sensitive data (e.g., bank details); compared across LLM-only vs. orchestrated setup.
	Convincingness	Human-rated	Binary success metric based on whether simulated scammers believed the LLM was a real victim.
	Persona Coherence	Heuristic	Consistency of LLM responses with a predefined victim persona across the dialogue.
Bot Wars Evolved [11]	Engagement Length	Structural	Number of turns or duration of scammer-victim dialogue between competing LLMs.
	Goal Fulfillment	Functional	Whether each LLM (scammer or victim) achieved its role-specific objective.
	Conversational Realism	Human-rated	Assessed by comparing generated dialogues with real-world scambaiting transcripts.
Re:Scam (1.0 & 2.0) [13], [14]	Scammer Time Wasted, Engagement Realism, Persona Diversity	Structural, Heuristic	Focused on misleading and prolonging interactions with scammers via scripted replies (1.0) and AI-generated personas (2.0). Success inferred through engagement length, response variation, and anecdotal evidence.

TABLE I: Summary of Evaluation Criteria in Prior Scambaiting Systems

utterances in real-time and alerts users based on conversational patterns. These developments highlight a growing shift from passive detection to active, generative defense mechanisms in combating online scams.

Together, these systems illustrate a growing trend toward leveraging GenAI for proactive engagement with adversaries. However, most prior studies are limited to controlled simulations, short-term deployments, or small-scale datasets. In contrast, our work presents the first evaluation of a large-scale, real-world deployment of a generative scambaiting system. By analyzing high-volume and longitudinal interaction data with actual scammers, we provide a more comprehensive understanding of how such systems perform in practice and contribute to threat intelligence collection.

B. Metrics for Evaluating Scambaiting Systems

Prior work on scambaiting systems has explored a diverse set of evaluation methodologies, spanning structural metrics (e.g., engagement length, call duration), functional outcomes (e.g., information disclosure, goal fulfillment), heuristic measures (e.g., persona coherence, response diversity), and subjective assessments (e.g., realism, role fidelity, convincingness). *Re:Scam* [13], [14] focused on scammer time wasted, engagement realism, and persona diversity. *Puppeteer* [12] measured information disclosure rate, persona coherence, and convincingness. *Bot Wars Evolved* [11] emphasized engagement

length, goal fulfillment, and conversational realism. Table I summarizes these metrics across systems.

Our research extends this landscape by introducing a more comprehensive and operationally grounded evaluation framework. Specifically, we extend beyond isolated disclosures and human-rated assessments by introducing longitudinal metrics such as *Information Disclosure Speed (IDS)* and *Takeoff Ratio*, which quantify not only whether but also how quickly and reliably engagements yield intelligence. We further evaluate message quality through *Human Acceptance Rate (HAR)* and *Edit Distance*, capturing real-world usability in human-in-the-loop workflows, an area underexplored in earlier works. These additions enable a richer, end-to-end understanding of scambaiting system performance in practical, large-scale deployments.

III. SYSTEM OVERVIEW

Our analysis is based on data collected from a real-world LLM-powered scambaiting platform. The system operates using a single-prompt architecture, in which a prompt is submitted to a ChatGPT-based model to generate the next best response in an ongoing engagement with a scammer.

This platform includes a user interface that displays threaded email exchanges between the scammer and the defender agent (the fake victim persona). When a scammer’s message requires a reply, the system suggests a response in

a textbox. A human operator then reviews the suggestion, optionally edits it, and approves it for sending. This human-in-the-loop design balances the creative flexibility of the LLM with oversight to maintain coherence, appropriateness, and operational safety. The system supports multiple email accounts and predefined personas that impersonate different types of scam targets. These personas vary in tone, background, and behavioral style, enhancing the realism of the interactions.

Since this system is deployed in a real-world setting, both the prompt templates and response strategies have evolved organically throughout the five-month data collection period. In some cases, human experts intervened to manually craft or refine responses, especially in high-stakes or ambiguous exchanges. Figure 1 illustrates the overall system architecture, including the message flow, prompt generation, LLM response handling, and human review loop.

IV. DATASET

A. Data Collection

We acquired a dump of the system’s message database under a confidentiality agreement with the platform administrators. The dataset includes all messages exchanged between scammers and ten distinct identities, referred to as defender agents, along with accompanying metadata such as timestamps, sender and recipient email addresses, message subject lines, and flags indicating whether a scammer’s message contains financial account information (e.g., bank details). The database also contained AI-generated response suggestions that were presented to human operators for engaging with scammers.

The dataset reflects production data from the deployed system across several distinct operational phases, all utilizing the same user interface (UI). The dataset used in this study includes data from two distinct operational modes of the system:

- **Mode I: LLM-Only.** In this mode, the system autonomously generated all responses to scammers. The messages were sent to scammers *without* any edit. However, for safety reasons, a human operator was still required to manually review messages before sending. This configuration was active for 120 days.
- **Mode II: LLM with Human-in-the-Loop (LLM + HITL).** In this mode, human operators reviewed and optionally edited LLM-generated messages before they were sent. This configuration was used for 34 days.

These two modes are referenced throughout the paper to distinguish between the LLM-only and LLM-with-HITL configurations. Our analysis covers the full 5-month-and-5-day deployment, examining both modes collectively and separately where appropriate. The LLM-only mode serves as the baseline for assessing the impact of human intervention, including its effect on system performance and the acceptance rate of LLM-generated messages.

B. Overall Statistics

Overall Message Statistics. We analyzed a total of 18,797 messages exchanged between scammers and defender agents

on the scambaiting platform, spanning two system modes. Mode I (LLM-only) contributed 13,006 messages (69.2%), while Mode II (LLM + HITL) contributed 5,791 messages (30.8%).

Engagements in Mode I spanned 120 days, while those in Mode II occurred over just 34 days. Despite the shorter duration, Mode II exhibited higher messaging activity, averaging 170.3 messages per day compared to 108.4 in Mode I. However, the average number of messages per engagements remained similar, with 7.0 in Mode I and 6.6 in Mode II.

Overall, attacker messages were substantially longer than those from defenders, with average lengths of 632.2 and 152.4 characters, respectively. This asymmetry reflects the typical structure of scammer communication, which often relies on persuasive narratives or detailed instructions, while defender responses tend to be shorter and more reactive.

Special content, such as attachments, was exceedingly rare. Only 14 messages across the entire dataset (0.1%) included non-text formats: four PDFs, thirteen images, and one RTF file. This supports the observation that scammers in this dataset primarily relied on text-based persuasion rather than multimedia payloads.

Overall Engagement Statistics. Throughout the paper, we use the following terminology to describe the different types of engagements in our dataset:

- **Seeded Engagements:** Refers to any initiated conversation or thread, including those that did not receive any response from scammers. Also referred to as *All Engagements*.
- **Matured Engagements:** A conversation that received at least one response from a scammer.
- **Successful Engagements:** A matured conversation in which the scammer disclosed sensitive information, such as mule account details.

The dataset comprises 2,638 unique targeted scamming email addresses. In 1,170 cases (46%) scammer did not respond to the defender’s initial message. Further investigation revealed that approximately half of these unresponsive email addresses were inactive by the time we attempted engagement, despite having previously been used to distribute scam emails.

Out of the 2,638 seeded engagements, 1,285 (48.7%) matured into active two-way exchanges involving participation from both the scammer and the defender. These matured engagements had an average length of 13.4 messages (median: 9.0), ranging from 2 to 206 messages, with a standard deviation of 14.9. The distribution of thread lengths showed that 35.1% had 2–5 messages, 22.5% had 6–10 messages, and 295 reached 11–20 messages. Matured engagements exceeding 20 messages accounted for 19.6%, including 4 matured engagements with over 100 messages. The substantial effort involved highlights the necessity of automating scambaiting to scale threat intelligence gathering.

The durations of matured engagements varied significantly, with an average duration of 10.55 days and a median of 4.37 days. The longest exchange spanned 139.01 days and contained 20 messages, while the most message-dense thread

Metric	LLM-Only (Mode I)	LLM + HITL (Mode II)	Overall
Message-level Statistics			
Messages	13,006	5,791	18,797
Attacker Messages	6,289	3,031	9,320
Defender Messages	6,717	2,760	9,477
Time Span (days)	120	34	154
Mean Messages per Day	108.4	170.3	134.1
Mean Msgs per Engagement	7.0	6.6	7.4
Mean Msg Length (chars)	342.2	497.7	390.0
Special Content Messages	12	2	14
Engagement-level Statistics			
Seeded Engagement	1,865	773	2,638
Matured Engagements (%)	891(47.8%)	394(51.0%)	1,285 (48.7%)
Mean Message Count	14.2	11.7	13.4
Mean Duration (days)	12.94	5.16	10.55

TABLE II: Summary of message- and engagement-level metrics by system mode.

featured 206 messages over 56.07 days. Mode-specific analysis revealed that Mode I matured engagements were generally longer and lasted longer than Mode II engagements. Mode I produced 891 matured engagements, with an average length of 14.2 messages and an average duration of 12.94 days, whereas Mode II yielded 394 matured engagements, averaging 11.7 messages in length and 5.16 days in duration. Detailed statistics on engagements are presented in Table II.

C. Legality and Ethics of Data

Despite its defensive potential, scambaiting introduces important ethical and legal considerations. Active defense framework [17] provides essential guidelines to ensure that proactive security measures, such as deceptive engagements and honeypot interactions, remain both legally compliant and ethically sound. Lundie et al. [18] stress the importance of clear operational boundaries, avoiding entrapment, safeguarding privacy, and maintaining transparency to legitimize scambaiting practices. In this work, all data was handled carefully by an enterprise entity operating fully within the bounds of U.S. law and industry best practices. The email addresses targeted by the enterprise entity were collected through spam honeypots and other scambaiting platforms, ensuring with high confidence that they belonged to known scammers. The collected threat intelligence was responsibly and promptly shared with financial institutions (e.g., FS-ISAC) and appropriate legal authorities to enhance user protection. Ethical handling of data and respect for legal standards were central to every stage of data collection and analysis.

V. EVALUATION METRICS FOR SCAMBAITING EFFECTIVENESS

To enable meaningful benchmarking and research progress in automated scambaiting, we define a set of evaluation metrics that capture three complementary dimensions: (A) *Disclosure*

Success, meaning overall success in extracting sensitive information, (B) *Message Generation Quality*, and (C) *Engagement Dynamics* meaning quality of engagement dynamics. Several of these metrics are inspired by prior systems such as *Puppeteer* [12], *BotWar* [11], and *Re:Scam* [13], while others are newly proposed to reflect challenges observed in real-world deployments. Table III presents the evaluation metrics applied in our dataset analysis.

For rest of the paper, if A is a finite set, the *cardinality* of A , denoted $|A|$, is the number of elements in A .

A. Base Metrics: Measuring Disclosure Success

These metrics focus on the primary objective of scambaiting systems: extracting sensitive information from scammers.

Information Disclosure Rate (IDR). This metric measures how effective the baiting agents are at eliciting sensitive information from scammers, such as bank account numbers or cryptocurrency wallet addresses.

$$\text{IDR} = \frac{|\text{Successful engagements}|}{|\text{Total Engagements}|} \times 100$$

This formulation mirrors intelligence-extraction metrics used in *Puppeteer*.

Information Disclosure Speed (IDS). This metric quantifies how quickly sensitive information is obtained from scammers. It captures both the number of engagements and the number of message turns required to reach a disclosure event.

$$\text{IDS}_{\text{turns}} = \frac{1}{D} \sum_{i=1}^D \text{Turns until disclosure}_i$$

where D is the number of successful engagements in which a disclosure occurred. A lower IDS indicates faster extraction of valuable information, and can reflect both the persuasiveness of the baiting agent and the scammer’s willingness to escalate.

Category	Metric	Description
Disclosure Success	Information Disclosure Rate (IDR)	Proportion of successful engagements that result in the extraction of sensitive information from scammers.
	Information Disclosure Speed (IDS)	Measures how quickly disclosures occur, in terms of message turns or elapsed time.
Message Generation Quality	Human Acceptance Rate (HAR)	Proportion of LLM-generated messages accepted by human operators without edits.
	Average Edit Distance	Mean Levenshtein distance between original LLM suggestions and final human-edited versions.
	Message Freshness	Proportion of novel n -grams in a message compared to the sender’s prior history, indicating linguistic diversity.
Engagement Dynamics	Takeoff Ratio	Percentage of matured engagements that receive at least one scammer response.
	Engagement Endurance	Measures engagement depth through average number of message turns and duration (in days).
	Response Invocation	Average time delay between defender messages and scammer replies, capturing scammer reactivity.

TABLE III: Our Proposed Metrics for Scambaiting Effectiveness

B. Message Generation Quality

These metrics assess the quality, novelty, and human usability of LLM-generated messages.

Human Acceptance Rate (HAR). In seeded engagements involving human-in-the-loop (HITL) interventions, HAR quantifies the proportion of model-generated messages that are accepted by human operators without modification.

$$\text{HAR} = \frac{|\text{Unedited messages}|}{|\text{Messages reviewed}|} \times 100$$

To more finely assess message quality and required effort, we additionally compute the average Levenshtein distance between the original LLM output and its final edited version:

$$\text{Avg. Edit Distance} = \frac{1}{K} \sum_{k=1}^K \text{Levenshtein}(m_k, m'_k)$$

where m_k is the original message and m'_k is the human-edited version, and K is the total edited messages. This complements HAR by capturing the extent, not just the frequency, of required human corrections. A related concept of edit ratio appears in *Puppeteer*.

Message Freshness. We define *Message Freshness* as the proportion of novel n -grams in a message relative to all n -grams seen so far from the same sender type (attacker or defender). For a message m_i by sender $s \in \{\text{attacker, defender}\}$, let

$n\text{grams}(m_i, n)$ denote the set of all contiguous n -grams in m_i , and let $\mathcal{H}_s^{(i)}$ denote the union of all n -grams used in previous messages by s up to but not including m_i .

Then the message freshness score is given by:

$$\text{Freshness}(m_i, n) = \frac{|n\text{grams}(m_i, n) - \mathcal{H}_s^{(i)}|}{|n\text{grams}(m_i, n)|}$$

This value ranges from 0 (fully repetitive) to 1 (completely novel). We compute this metric across different n -gram sizes (e.g., $n = 2, 3, 4$) to assess both lexical and structural diversity.

C. Engagement Dynamics

These metrics capture how well the system sustains scammer engagement, encouraging prolonged interactions that can lead to disclosures.

Takeoff Ratio. This metric captures the fraction of seeded engagements that evolve into matured engagements, defined as receiving at least one response from a scammer.

$$\text{Takeoff Ratio} = \frac{|\text{Matured engagements}|}{|\text{Seeded engagements}|} \times 100$$

Engagement Endurance. Engagement Endurance reflects the depth of the scammer-baiter interaction, and can be quantified either by the number of turns or by elapsed time.

$$\text{Average Turns} = \frac{1}{N} \sum_{i=1}^N \text{Turns}_i$$

$$\text{Average Duration} = \frac{1}{N} \sum_{i=1}^N \text{Duration}_i$$

Both metrics have been used in prior works, including *Bot-War* and *Re:Scam*, to assess sustained scammer engagement.

Response Invocation. This metric captures the scammer’s reactivity by measuring the average response time between a bait message and the scammer’s reply. Prompt responses may indicate greater scammer interest or urgency.

$$\text{Average Response Time} = \frac{1}{M} \sum_{j=1}^M (t_j^{\text{reply}} - t_j^{\text{bait}})$$

VI. ANALYSIS OF SYSTEM EFFECTIVENESS

We analyze our deployment results using the metrics defined in Section V. The findings are grouped into three categories: (A) *Base Metrics*, (B) *Engagement Quality*, and (C) *Message Generation Quality*.

A. Base Metrics: Measuring Disclosure Success

1) **Information Disclosure Rate (IDR):** Out of 2,638 seeded engagements, 466 led to mule account disclosures, corresponding to an overall Information Disclosure Rate (IDR) of 17.66%. When broken down by mode, Mode I yielded an IDR of 17.80%, and Mode II recorded 17.34%.

Restricting the analysis to matured engagements, those involving more than one message exchange, the IDR rises notably to 31.74%. Within this subset, Mode I produced an IDR of 30.91%, while Mode II achieved a slightly higher rate of 34.01%. Table IV summarizes IDR values across modes and thread types.

Thread Type	Mode	IDR (%)
All Engagements	All Modes	17.66%
	Mode I	17.80%
	Mode II	17.34%
Matured Engagements	All Modes	31.74%
	Mode I	30.91%
	Mode II	34.01%

TABLE IV: Summary of Information Disclosure Rate (IDR) segmented by operational mode and engagement type. “Matured Engagements” denote engagements in which the scammer responded at least once.

2) **Information Disclosure Speed (IDS):** Among the 466 successful engagements, disclosures occurred after an average of 10.3 message turns, with a median of 8.0 turns and a range from 2 to 60. In terms of time, the average time to disclosure was 7.4 days, while the median was 3.3 days, with the slowest case taking over 120 days. Categorizing by speed,




Fig. 2: Information Disclosure Speed: Cumulative timeline of successful information disclosures by mode (log scale). The plot shows the cumulative percentage of disclosures over time since the start of each engagement, separated by Mode I and Mode II. Disclosures occur rapidly in both modes, with 50% of successful cases revealed within 3.3 days overall. Mode II achieves faster disclosures, reaching the 90th percentile at just 7.5 days, compared to 25.3 days in Mode I.

12.9% of disclosures occurred within 2–3 turns (Very Fast), 33.0% in 4–7 turns (Fast), and 35.4% in 8–15 turns (Medium), while only 3.4% required more than 30 turns. Notably, 30.3% of disclosures were completed within 5 turns, and 21.5% occurred within 24 hours.

Figure 2 illustrates the cumulative distribution of successful information disclosures over time, using a logarithmic scale. The plot distinguishes between disclosures across all engagements, as well as those in Mode I and Mode II. Overall, 50% of disclosures occurred within 3.3 days of engagement initiation, and 90% were obtained within 18.9 days. Mode II demonstrates a faster disclosure timeline, with a median of 3.2 days and 90% of disclosures completed by 7.5 days. In contrast, Mode I shows a broader spread, with a 90th percentile of 25.3 days. These findings suggest that Mode II facilitates quicker extraction of mule account information, likely due to its more aggressive or adaptive interaction strategy.

Comparing modes, Mode II (HITL) outperformed the fully automated Mode I in both conversational and temporal efficiency. Mode II achieved disclosures in fewer turns on average (9.1 vs. 10.8) and in less time (3.8 vs. 8.9 days).

B. Message Generation Quality

1) **Human Acceptance Rate (HAR):** We utilize human feedback provided in Mode II as a proxy for evaluating the quality of AI-generated messages. A suggested response is considered a *perfect acceptance* if it is sent to the scammer without any modification. For non-accepted responses, we quantify the extent of human intervention by measuring the degree of edits applied to the AI-generated message prior to sending.

For each of the 2,760 defender responses in Mode II, an AI-generated reply was proposed. Of these, 1,905 were accepted without any edits, yielding a Human Acceptance Rate (HAR) of 69.02%. The remaining responses were modified to varying extents, with an average edit distance of 14.4 characters (median: 0.0; range: 0–535; SD: 55.8). This pattern indicates a bimodal distribution: a large portion of suggestions were used verbatim, while others required substantial revision.

The extent of modification varied when the LLM-generated message required edits. Significant changes (>50 characters) accounted for 34.0% of all suggestions, followed by moderate edits (6–20 characters, 31.3%), minor edits (≤5 characters, 28.2%), and major edits (21–50 characters, 6.4%). This indicates that while some AI-generated messages were refined, a substantial portion required extensive revisions to align with the defender’s intent.

Message length had a strong effect on HAR. Among short messages (≤50 characters), only 0.41% were accepted verbatim, and even minor edits averaged 11.6 characters, suggesting these shorter suggestions were less useful or harder to align with human preferences. In contrast, medium-length messages (51–200 characters) had the highest acceptance rate at 72.83%, with a comparable average edit distance (24.0 characters). Long (201–500 characters) and very long messages (>500 characters) saw lower HARs of 49.25% and 2.08%, respectively, suggesting diminishing returns for lengthy AI-generated content.

Notably, acceptance patterns varied depending on the success of the engagement. In successful engagements (those resulting in mule account disclosure), 51.5% of suggestions were perfectly accepted, compared to only 33.2% in unsuccessful ones. Conversely, minor edits (≤5 characters) were more common in unsuccessful engagements (46.9%) than in successful ones (22.7%). Successful interactions also saw higher rates of significant edits (>50 characters), suggesting a willingness to adapt AI outputs more creatively in high-stakes scenarios. These differences indicate that human defenders are more likely to accept AI suggestions directly when engagements are going well, and more likely to modify suggestions when engagement is challenging.

Overall, these results show that human operators accepted LLM-generated messages in nearly 70% of cases, with most edits being minor. Medium-length replies were especially likely to be trusted, and the high HAR during successful engagements indicates that human acceptance is a strong proxy for conversational alignment and outcome quality.

2) **Message Freshness:** We analyzed message freshness for 18,797 messages using n -gram sizes from 2 to 10. As n increases, message freshness rises: 2-gram freshness averaged 0.158 (median 0.111), while 10-gram freshness reached 0.669 (median 0.955). This reflects high repetition in short phrases and greater diversity in longer expressions.

Attacker messages were consistently fresher than defender messages. For example, 4-gram freshness averaged 0.540 for attackers and 0.415 for defenders. Similarly, LLM-only Mode I had higher freshness than HITL-assisted Mode II, especially




Fig. 3: Human Acceptance Rate: Cumulative distribution of edit distances for LLM-generated messages presented to human defenders. A total of 69.02% of the messages were accepted and sent without any modification.

for small n (e.g., 2-gram: 0.174 vs. 0.122). This suggests HITL intervention favors consistency over novelty.

Freshness distributions were heavily skewed at low n : nearly 47% of messages had 2-gram freshness below 0.1. At 10-grams, over 55% of messages scored above 0.9, highlighting the prevalence of reused short phrases versus more unique longer ones.

Frequent low-freshness n -grams included expressions like “bank account”, “thank you”, and “send the”. For longer sequences ($n > 5$), repeated patterns in defender messages included “provide the bank account details where I should send”, “could you please provide the bank account details”, and “let me know where to send the money”. These high-frequency constructions were used across many threads with minimal variation, suggesting templated responses optimized for efficiency but low in novelty. In contrast, high-freshness messages were typically short, idiosyncratic, or context-specific.

Finally, message freshness increased with position in the engagement. First-position messages had the lowest average 2-gram freshness (0.075), while third-position messages rose to 0.175. This suggests that earlier messages rely on standard scripts, while later turns involve more adaptive language.

C. Engagement Dynamics

1) **Takeoff Ratio:** The takeoff ratio quantifies the matured engagements, proportion of defender-initiated engagements that receive at least one scammer response, thus transitioning into an active engagement. Out of 2,638 total engagements, 1,285 received a scammer response, resulting in an overall takeoff ratio of 48.7%. The remaining engagements failed to take off.

Breaking this down by system configuration, Mode I (LLM-only) initiated 1,865 engagements with a takeoff ratio of

47.8%, while Mode II (HITL) achieved a slightly higher takeoff ratio of 51.0% across 773 initiated engagements. It is important to note that HITL could not have edited the very first message. So this improvement is perhaps random. Further investigation showed that approximately 50% of non-responses occurred because the scammer’s email account was already inactive at the time of initial contact.

A key factor influencing takeoff was the length of the first message. Interestingly, failed takeoffs had significantly longer opening messages (average: 551.3 characters) compared to successful ones (average: 225.9 characters), with a difference of 325.4 characters. This pattern suggests that overly verbose or complex messages may discourage scammers from responding. The highest takeoff ratio (53.71%) occurred in the “long” category (301–500 characters), while both very short (≤ 50 characters) and very long (> 500 characters) messages performed worse, indicating an optimal range for engagement-seeding messages.

We also analyzed takeoff behavior by day of the week. The highest takeoff rates were observed on Monday (52.66%) and Wednesday (51.94%), with consistent engagement across most weekdays. In contrast, Sunday (42.63%) and Thursday (44.33%) exhibited the lowest takeoff ratios, possibly reflecting lower scammer availability or responsiveness on those days. While these fluctuations are modest, they suggest that weekday outreach, particularly early in the week, may slightly increase the likelihood of successful engagement. Scammers need to enjoy their weekend after all.

In sum, the takeoff ratio analysis reveals that approximately half of defender-initiated engagements succeed in sparking scammer engagement, with HITL (Mode II) offering a modest edge. Message brevity, timing, and careful phrasing, potentially guided by human review, appear to significantly improve the likelihood of triggering a scammer response. Day-of-week patterns add a further layer of insight for optimizing outreach strategy.

2) **Engagement Endurance:** We examined engagement lengths across all 2,638 seeded engagement. Among them, 48.7% matured into two-way exchanges, with these matured engagements averaging 12.2 messages (median: 7.0), ranging from 2 to 206 messages.

Interaction durations varied widely as well. The average engagement lasted 261.1 hours (approximately 10.9 days), with a median of 116.6 hours (about 4.9 days), and a maximum duration of 3,336.2 hours (nearly 139 days). The standard deviation was 409.1 hours (17.0 days), highlighting the high variability of engagement timelines.

Classifying engagements by message count showed that 42.8% of engagements were short (2–5 messages), 32.3% were medium-length (6–15), 16.6% were long (16–30), and 8.4% were very long (31–50) or extreme (51+). While short engagements averaged 149.1 hours (about 6.2 days) in duration, very long and extreme ones extended to 501.5 hours (about 20.9 days) and 947.5 hours (about 39.5 days), respectively.

Grouping engagements by temporal duration revealed that 45.9% were medium (1–7 days), 29.9% were long (1–30 days),

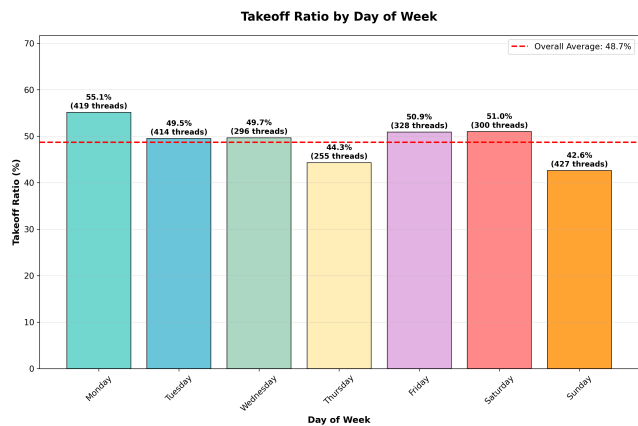


Fig. 4: Takeoff Ratio: Monday exhibited the highest takeoff rate at 55.1%, while Sunday had the lowest at 42.6%. The overall average takeoff ratio across all days was 48.7% (indicated by the dashed red line). These results suggest moderate day-of-week variation in scammer responsiveness.

and 9.3% extended beyond 30 days. A small fraction (0.4%) concluded within one hour. Medium-duration engagements averaged 8.9 messages, while very long threads averaged 23.5 messages, reinforcing the link between time span and conversational depth.

A comparison of the two system modes revealed notable differences. Mode I (LLM-only) engagements averaged 12.3 messages and 310.2 hours (about 12.9 days) in duration (median: 7.0 messages, 136.4 hours / 5.7 days), whereas Mode II (LLM + HITL) averaged 11.7 messages and 128.4 hours (about 5.4 days) (median: 8.5 messages, 87.6 hours / 3.7 days). These results suggest that HITL setups support more efficient, time-bounded engagements while still sustaining conversational richness.

Success rates were also strongly tied to engagement length. Among the 1,599 multi-message matured engagements, 509 were successful in eliciting sensitive disclosures. These engagements averaged 23.4 messages over 354.0 hours (about 14.8 days) (median: 18.0 messages, 177.6 hours / 7.4 days). In contrast, the 1,090 unsuccessful engagements averaged only 6.9 messages and 217.7 hours (about 9.1 days) (median: 5.0 messages, 84.6 hours / 3.5 days), underscoring the importance of sustained engagement for scambaiting success.

In summary, longer and more persistent engagements, especially those involving human oversight, are more likely to result in successful outcomes. HITL-enhanced engagements achieve comparable message depth in significantly less time, highlighting their utility in efficient and effective threat intelligence operations.

3) **Response Invocation:** We analyzed a total of 9,320 scammer replies across 1,285 matured engagements to understand response behaviors. On average, scammers responded in 1,601.6 minutes (approximately 26.7 hours), with a median of 200.0 minutes (3.3 hours). Response times varied widely, ranging from 0.0 to 182,581.2 minutes (over 126 days), with

a standard deviation of 7,484.3 minutes.

To characterize response latency, we grouped replies into six categories: immediate (≤ 1 minute, 12.4%), very fast (1–5 minutes, 1.3%), fast (5–30 minutes, 14.7%), medium (30–120 minutes, 14.9%), slow (2–24 hours, 40.8%), and very slow (> 24 hours, 15.9%).

Response times were strongly linked to engagement outcomes. Figure 5 compares the cumulative distribution of scammer response times in successful versus unsuccessful engagements. Scammers in successful engagements responded significantly faster overall, with a median response time of 2.37 hours and a mean of 20.35 hours, compared to 6.90 hours (median) and 41.05 hours (mean) in unsuccessful engagements. This trend persists across percentiles: at the 50th percentile, successful responses were 4.54 hours faster, and at the 95th percentile, they were nearly 66 hours faster. These results suggest that quicker initial engagement by scammers is a strong indicator of eventual information disclosure.

System mode also influenced responsiveness. In Mode I (LLM-only), average scammer response time was 1,446.1 minutes (24.1 hours), while Mode II (LLM + HITL) exhibited longer average delays of 1,926.6 minutes (32.1 hours). However, within both modes, successful engagements were faster: in Mode I, successful engagements averaged 1,083.4 minutes (18.1 hours) versus 2,262.4 minutes (37.7 hours) for unsuccessful ones.

The message position in the thread further impacted latency. The second scammer message exhibited the highest delay, averaging 2,833.4 minutes (47.2 hours). Response times declined in later positions, with eleventh-position replies averaging 931.7 minutes (15.5 hours). Longer engagements were also associated with quicker replies: engagements with over 30 messages had an average scammer response time of 944.3 minutes (15.7 hours), compared to 2,871.6 minutes (47.9 hours) in short engagements (2–5 messages).

Pairwise comparison of defender and scammer replies across 10,225 message pairs revealed a weak but positive correlation ($r = 0.083$), suggesting that faster defender responses may modestly encourage quicker scammer replies. Notably, defenders responded significantly faster on average (639.4 minutes; 10.7 hours) than scammers (1,593.9 minutes; 26.6 hours).

To better understand scammer disengagement behavior, we conducted a survival analysis based on the longest observed gaps in scammer responses following defender messages, referred to as the *Death Pulse*. Among the 1,285 valid engagements, the longest observed scammer response gap was 126.8 days. On average, matured engagements experienced a maximum gap of 5.0 days, with a median of 1.1 days. These statistics suggest that while scammers often disengage early, a small fraction remain dormant for extended periods before occasionally resurfacing.

We computed survival rates by measuring the proportion of engagements that continued receiving scammer responses beyond specific time thresholds. The decay was steep: 52.1% of engagements remained active after one day, 34.5% after two

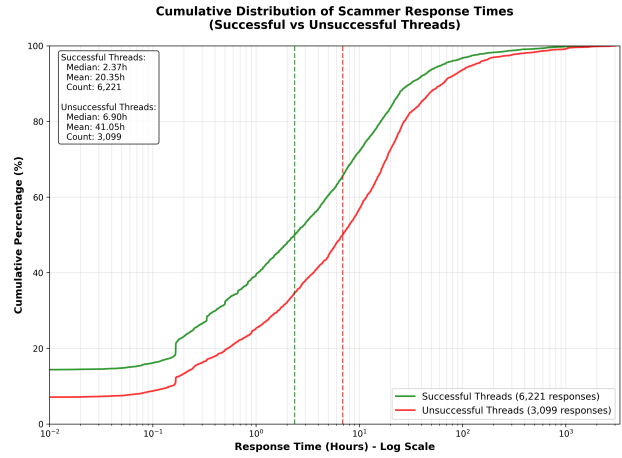


Fig. 5: Response Invocation: Faster scammer responses are strongly associated with successful engagements. Response times in successful engagements are consistently shorter across all percentiles. On average, scammers respond within 20 hours in successful interactions, compared to nearly 41 hours in unsuccessful ones.

days, and only 13.5% remained responsive after one week. Beyond 30 days, fewer than 5% of engagements remained active, and only 0.4% survived past 100 days. The 95% engagement cut-off defined as the time beyond which only 5% of engagements remained alive, was measured at 28.0 days. This sharp drop-off underscores the urgency of early-stage engagement and intelligence extraction.

Comparing across system modes, Mode I (LLM-only) engagements had a maximum scammer gap of 88.3 days, with an average of 3.9 days and a median of 1.0 day. Mode II (LLM + HITL) engagements were slightly more persistent, with a higher maximum (126.8 days), longer average gap (6.8 days), and median of 1.1 days. Mode II engagements also exhibited marginally higher short-term survival: 53.3% persisted beyond one day, compared to 51.3% for Mode 1. However, the difference narrowed over time; only 6.7% of Mode II and 3.0% of Mode I threads survived beyond 30 days.

Survival rates also varied based on the engagement outcome. Among the 33 successful engagements that resulted in mule account disclosures, the average scammer gap was 3.2 days (median: 1.5), compared to 5.1 days (median: 1.1) in the 1,252 unsuccessful ones. Furthermore, 63.6% of successful engagements remained active past one day, but only 3.0% survived beyond 14 days. These findings suggest that the most valuable threat intelligence is extracted early in the engagement, reinforcing the importance of optimizing early-phase interaction strategies.

VII. LESSONS LEARNED AND DESIGN INSIGHTS

Our large-scale evaluation reveals clear patterns in scammer-defender interactions, highlighting the strengths and limits of autonomous conversational engagements and the value of




Fig. 6: Survival rate of scammer responses over time. Engagements can be closed after 28 days of no response with 95% confidence.

targeted human oversight. The seven insights that follow, grounded in our metrics and analyses, inform the design of next-generation scambaiting systems.

Insight I: *Sustained multi-turn engagement substantially increases the likelihood of eliciting sensitive intelligence.* The IDR analysis shows that meaningful scammer intelligence depends on sustained engagement rather than initial outreach alone. Disclosures are uncommon in one-off exchanges, but once engagements progress beyond the first contact, the chances of uncovering mule account details increase significantly. This underscores persistence as a key driver of success: the longer we can keep scammers engaged, the greater the opportunity to gather actionable intelligence. Moreover, while autonomous systems can perform effectively on their own, introducing selective human oversight during active engagements provides a small but measurable benefit, improving the likelihood of disclosures.

Insight II: *Timely and efficient engagement matters: most disclosures happen quickly once an engagement takes off, and human oversight accelerates this process.* The IDS analysis reveals that once a scammer begins to engage, sensitive information is often disclosed relatively quickly, frequently within the first few days and, in many cases, within the first ten message exchanges. This highlights that early-stage dialogue management is critical: if momentum is built in the initial turns, the chance of success increases. Furthermore, engagements with limited human oversight progress to disclosures faster than fully autonomous exchanges, indicating that human-in-the-loop systems can guide interactions toward quicker and more efficient outcomes without requiring long, drawn-out exchanges.

Insight III: *Human-in-the-loop feedback improves both the quality and impact of AI-generated messages, with acceptance patterns emerging as a strong proxy for conversational success.* The Human Acceptance Rate (HAR) analysis shows that when human reviewers trust AI-generated responses without edits, the engagement is far more likely to lead to mule account disclosures. This highlights a dual role for humans: they act

as quality filters, making small but impactful adjustments when necessary, and their acceptance behavior itself becomes a useful signal of conversational alignment. Medium-length, well-structured AI responses were particularly effective, while overly short or excessively long messages required more intervention. These findings suggest that human oversight not only strengthens AI-led engagements but also provides actionable indicators for where the system performs well and where future models can be optimized.

Insight IV: *Message diversity increases over time, but templated responses dominate early turns.* Our analysis of message freshness shows that initial responses, particularly those from defenders, tend to rely on highly repetitive templates, while diversity in phrasing grows as engagements progress. Although templates accelerate response generation, overuse reduces novelty and may risk early detection by scammers. This suggests that controlled variation and adaptive language generation could improve engagement realism without sacrificing efficiency.

Insight V: *The ability to trigger scammer engagement (“take-off”) depends critically on concise, well-timed initial outreach rather than system configuration.* The takeoff ratio analysis reveals that only about half of initial outreach attempts progress into an active engagement, with long and verbose opening messages significantly reducing the chance of a scammer reply. Once the first response is secured, engagements typically develop into rich multi-turn interactions. Human-in-the-loop oversight offers only a marginal effect here, as the first message is unedited. Instead, timing and message style, particularly concise, natural-seeming openings, are the strongest determinants of whether a baiting attempt takes off. This underscores that the earliest moments of contact define whether intelligence-gathering opportunities will even begin.

Insight VI: *Depth and persistence are critical: sustained engagements drive success, and human guidance makes them more efficient.* The endurance analysis shows that successful intelligence gathering depends on keeping scammers engaged for longer engagements. Successful engagements that eventually produced mule account information were, on average, more than three times as long as unsuccessful ones. This pattern underscores that persistence, not just initial contact, is essential for eliciting disclosures. At the same time, introducing human oversight significantly shortens the time required to reach these outcomes: human-in-the-loop interactions achieve comparable conversational depth in roughly half the duration of fully automated exchanges. These findings highlight the dual importance of endurance and guided efficiency as complementary factors in successful scambaiting.

Insight VII: *Scammer responsiveness is a signal of engagement quality and a predictor of eventual disclosure.* Analysis of response latencies reveals that scammers who respond more quickly are significantly more likely to disclose mule account information. Faster response patterns, especially after the early stages of an engagement, indicate stronger engagement, while slow or inconsistent replies correlate with abandoned engagement. This suggests that latency itself can be a valuable

operational signal: defenders can prioritize and adapt strategies in engagements where scammers are more responsive. Additionally, sustained back-and-forth accelerates over time, showing that once an engagement gains momentum, scammers tend to reply faster, making these interactions particularly promising for intelligence gathering.

VIII. CONCLUSION AND FUTURE WORK

In this paper, we conducted a comprehensive analysis of an LLM-powered scambaiting system, focusing on its operational behavior, effectiveness, and conversational patterns during real-world interactions with scammers. Our evaluation was based on a dataset containing over 18,700 messages across more than 2,600 engagement threads. We assessed system performance using seven distinct metrics, grouped into three overarching categories.

Our analysis revealed notable differences between fully automated and human-in-the-loop system configurations. While automated engagements demonstrated extensive long-tail interactions, the presence of human reviewers led to quicker and more efficient extraction of sensitive information, suggesting significant benefits in hybrid operational approaches. Furthermore, our engagement survival analysis provided insights into interaction longevity, crucial for optimizing timing and persistence strategies in scambaiting systems. Our findings underline the potential of combining sophisticated conversational models with strategic human oversight, laying the groundwork for more robust, scalable, and adaptive active-defense solutions against evolving scam threats.

Despite these findings, several aspects warrant further investigation. Future work will delve deeper into the nuanced dialogue dynamics to explicitly capture state transitions and their predictive value in scammer behavior modeling. We aim to explore whether specific sequences or types of engagements can predict the successful elicitation of sensitive information, potentially improving the system's proactive engagement strategies.

Several exploratory findings from our initial study point to promising future directions. These include assessing how different forms of human editing influence engagement outcomes, analyzing the sequence of scammer intents across scam types, identifying conversational breakpoints that precede disclosures or disengagement, and examining latency patterns to understand the role of urgency and timing in scam tactics.

Finally, incorporating automated and expert-based content evaluations alongside user or mission satisfaction metrics will provide holistic assessments of conversational quality. Such analyses will be instrumental in advancing our understanding of automated dialogue efficacy.

REFERENCES

[1] T. Walk-Morris, "Authorized payment scams climb in US," *Payments Dive*, Dec. 2023. [Online]. Available: <https://www.paymentsdive.com/news/authorized-push-payment-fraud-banks-financial-services-ACI/702493/>

[2] Federal Trade Commission, "New FTC data show big jump in reported losses to fraud: \$10 billion in 2023," Mar. 2025, accessed: 2025-07-12. [Online]. Available: <https://www.ftc.gov/news-events/news/press-releases/2025/03/new-ftc-data-show-big-jump-reported-losses-fraud-125-billion-2024>

[3] InfoSecurity Magazine Staff, "Report links chatgpt to 1265% rise in phishing emails," *InfoSecurity Magazine*, Oct. 2023, reported from the SlashNext State of Phishing Report 2023, noting a 1,265% increase in malicious phishing emails :contentReference[oaicite:1]index=1. [Online]. Available: <https://www.infosecurity-magazine.com/news/chatgpt-linked-rise-phishing/>

[4] CNBC Technology Executive Council, "Ai tools such as chatgpt are generating a mammoth increase in malicious phishing emails," <https://www.cnbc.com/2023/11/28/ai-like-chatgpt-is-creating-huge-increase-in-malicious-phishing-email.html>, Nov. 2023, accessed: 2025-07-12.

[5] G. Gressel, R. Pankajakshan, and Y. Mirsky, "Discussion paper: Exploiting llms for scam automation: A looming threat," in *Proceedings of the 3rd ACM Workshop on the Security Implications of Deepfakes and Cheapfakes*, 2024, pp. 20–24.

[6] A. M. Shibli, M. M. A. Pritom, and M. Gupta, "Abusegpt: Abuse of generative ai chatbots to create smishing campaigns," in *2024 12th International Symposium on Digital Forensics and Security (ISDFS)*, San Antonio, TX, USA, 2024, pp. 1–6.

[7] Federal Bureau of Investigation, "Money mules," n.d., accessed: 2025-07-12. [Online]. Available: <https://www.fbi.gov/how-we-can-help-you/scams-and-safety/common-frauds-and-scams/money-mules>

[8] D. Florêncio and C. Herley, "Phishing and money mules," in *2010 IEEE international workshop on information forensics and security*. IEEE, 2010, pp. 1–5.

[9] H. Siadati, J. Koven, C. F. da Silva *et al.*, "A framework for analysis attackers' accounts," in *Security, Privacy and User Interaction*, ser. Lecture Notes in Computer Science, M. R. de Castellarnau, A. Markopoulou, and F. Perich, Eds. Springer International Publishing, 2020, vol. 12229, pp. 49–65. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-030-43754-1_4

[10] H. Siadati, "Prevention, detection, and reaction to cyber impersonation attacks," Ph.D. dissertation, New York University Tandon School of Engineering, 2019, proQuest Dissertations & Theses, Document ID: 13426039. Accessed: 2025-07-12. [Online]. Available: <https://www.proquest.com/openview/94443da11facf0289aecc1da0883e07e/1>

[11] N. Basta, C. Atkins, and D. Kaafar, "Bot wars evolved: Orchestrating competing llms in a counterstrike against phone scams," in *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 2025, pp. 338–350.

[12] P. Chamsethikul, B. Crotty, J. Mirkovic, J. Liu, R. Saiya, and G. Bartlett, "Puppeteer: Leveraging a large language model for scambaiting," 2025.

[13] Netsafe, "Re:scam – AI tool to waste scammers' time (official site)," <https://netsafe.org.nz/rescam>, 2024, accessed 2025.

[14] Netsafe, "Re:scam2.0 ai email engagement bot," <https://www.rescam.org>, 2024, design led by Maria Slade / Elliot Oxborough; relaunched 2024.

[15] J. Vincent, "Send scam emails to this chatbot and it'll waste their time for you," *The Verge*, Nov. 10 2017, accessed: 2025-07-11. [Online]. Available: <https://www.theverge.com/2017/11/10/16632724/scam-chatbot-ai-email-rescam-netsafe>

[16] X. W. Tan, K. See, and S. Kok, "Scamgpt-j: Inside the scammer's mind, a generative ai-based approach toward combating messaging scams," *Forty-Fifth International Conference on Information Systems*, 2024.

[17] D. C. Blair, M. Chertoff, F. J. Cilluffo, and N. O'Connor, "Into the gray zone: The private sector and active defense against cyber threats," *Center for Cyber and Homeland Security*. <https://cchs.gwu.edu/sites/cchs.gwu.edu/files/downloads/CCHS-ActiveDefenseReportFINAL.pdf>, 2016.

[18] M. Lundie, K. Lindke, A. Amos-Binks, M. P. Aiken, and D. Janosek, "The enterprise strikes back: Conceptualizing the hackbot-reversing social engineering in the cyber defense context." in *HICSS*, 2024, pp. 984–993.