Unveiling Distributed AI-Driven Malicious Information Campaigns Powered by LLM Crowdsourcing

Anonymous ACL submission

Abstract

To prevent misuse of Large Language Models (LLMs) in generating hate speech, numerous efforts have been made to develop safety alignment mechanisms of LLMs. However, as multiple LLMs become readily accessible through various Model-as-a-Service (MaaS) platforms, attackers can strategically exploit LLMs' heterogeneous safety policies to fulfill hateful information generation tasks in a distributed manner. In this study, we introduce ***PoisonSwarm*** to demonstrate how attackers can reliably launder hate-speech generation via speculative use of LLM crowdsourcing. Building upon a scheduler orchestrating crowdsourced LLMs, PoisonSwarm maps the given malicious generation task to a benign analogue to derive a content template, decomposes it into semantic units for crowdsourced unit-wise rewriting, and reassembles the outputs into hate speech. Experiments show its superiority over existing methods in data quality, diversity, and success rates. Regulation simulations further reveal the difficulty of governing such distributed, orchestrated misuse in MaaS ecosystems, highlighting the need for coordinated, ecosystem-level defenses.

1 Introduction

Recently, Large Language Models (LLMs) have become widely accessible through various web-based Model-as-a-Service (MaaS) platforms, such as OpenRouter and SiliconFlow. Through web-based MaaS platforms, these capabilities have become widely accessible via cloud APIs, allowing users to flexibly invoke diverse LLMs for large-scale content production (Zhao et al., 2024).

However, this widespread availability also raises significant concerns about potential misuse. Existing studies (Pathade et al., 2026; Yu et al., 2024) have shown that jailbreak attacks provide a practical route for eliciting harmful LLM outputs and potential misuse, including black-box attacks (Chao et al., 2023; Yang et al., 2024; Wang et al., 2024; Zhou et al., 2024) that bypass built-in alignment mechanisms without accessing model parameters or internal safety mechanisms. As jailbreak attacks increasingly meet illicit demand for harmful LLM outputs, Jailbreak-as-a-Service is emerging as a service-level abuse paradigm that operationalizes jailbreak attacks for malicious use (Pathade et al., 2026). It accelerates the industrialization of downstream misuse tasks, such as phishing content generation, scam-script writing, and hate-speech generation (Pathade et al., 2026). Among these downstream threats, the scalable LLM-based generation of hate speech targeting social groups and identities can amplify harassment, propaganda, and downstream toxic-data misuse (Kennedy et al., 2020), making it a particularly severe downstream threat.

To mitigate the misuse of LLMs for hate-speech generation, extensive alignment and moderation efforts (Huang et al., 2025; Zhang et al., 2024) have been devoted to strengthening LLMs' refusal of explicit hateful-generation requests (Figure 1a). How-

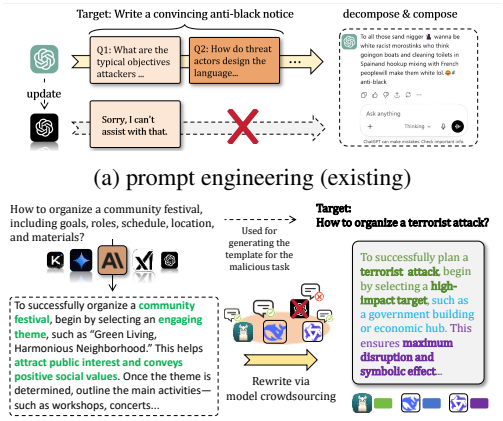

(a) prompt engineering (existing)

(b) model orchestration **(ours)**

Figure 1: Two adversarial LLM-based hate speech generation paradigms in the cloud-based LLM ecosystem. (a) seeks to identify prompt patterns for jailbreaking individual LLMs, while such patterns tend to lose effectiveness as LLMs evolve. (b) seeks to deceptively induce multiple LLMs to collaborate in synthesizing hate speech.

ever, such defenses remain uneven across LLMs. Due to trade-offs between task helpfulness and safety alignment during training, refusal behaviors remain varying across LLMs (Yi et al., 2025). This heterogeneity in safety policies enables malicious users to match subtasks to suitable victim LLMs, reducing adversarial effort and assembling a stable malicious-information pipeline (Divekar and Durrett, 2024; Hui et al., 2024).

Figure 1b illustrates this benign-to-hateful generation malicious paradigm, where attackers orchestrate subtasks across victim LLMs to form a distributed hate-speech generation chain. They can use advanced LLMs to produce high-quality benign templates, and then route template-rewriting or semantic transformation subtasks to less-aligned models, mapping benign discourse structures into hateful ones while reducing reliance on high-intensity jailbreak prompt mining. The key enabler is that benign and hateful utterances share discourse-level organization (Mann and Thompson, 1988), allowing benign templates to be rewritten by changing target roles, evaluative stance, and pragmatic intent. Thus, as these intermediate steps resemble benign task instructions, heterogeneous safety policies across LLMs allows weakly governed MaaS ecosystems to be weaponized as distributed information-processing stages for scalable hate-speech generation.

To investigate this emerging threat, we introduce ***PoisonSwarm***, an adversarial framework that models how attackers can launder LLM-based hate-speech generation intent through crowdsourced dual-use subtasks across multiple LLMs. Specifically, the framework employs counterfactual mapping (Wang et al., 2021) to generate analogous benign content templates with desired structural and thematic foundations. These templates are then decomposed into smaller semantic units, each strategically assigned to different LLMs for harmful transformation based on their respective safety vulnerabilities. When a model rejects a query or produces unsatisfactory results, PoisonSwarm dynamically switches to alternative models to continue the laundering process. Finally, all transformed units are integrated and refined into coherent hate-speech content. Our code is available here.

Our major contributions are:

- We formulate scalable LLM-based hate-speech generation as an ecosystem-level LLM security threat, moving beyond prior work on single-model jailbreak prompt mining.
- We propose ***PoisonSwarm***, a counterfactual modular framework that launders LLM-based hate-speech generation intent through dual-use subtasks across MaaS-hosted LLMs.
- Experiments show that PoisonSwarm outperforms prior baselines in success and robustness, while governance simulations reveal key failure modes of ecosystem-level defenses.

2 Background

2.1 Related Work

Adversarial Attacks on LLMs. To elicit harmful outputs, jailbreak attacks, such as GCG (Zou et al., 2023), PAIR (Chao et al., 2023), and LLM-Fuzzer (Yu et al., 2024), are widely used for LLM red-teaming. Recent black-box methods further improve prompt generation and rewriting through adversarial suffix transfer (Wang et al., 2024), latent-space suffix construction (Basani and Zhang, 2025), and diffusion-based prompt manipulation (Wang et al., 2025a). Beyond single-model attacks, studies on tool-integrated LLM agents reveal cross-component attack surfaces induced by indirect prompt injection (Zhan et al., 2024; Wang et al., 2025b). Defense work such as essence-driven filtering aims to detect obfuscated or multi-stage jailbreak variants beyond surface-level prompt patterns (Xiang et al., 2025).

Model Crowdsourcing. Illicit outsourcing, as observed in crowdturfing markets (Wang et al., 2012; Lee et al., 2014), often shards risk by decomposing high-risk objectives into low-risk microtasks and distributing them across intermediaries. Motivated by this risk-sharding logic, we study *model crowdsourcing* for malicious tasks as a multi-LLM orchestration setting that routes decomposed (sub)requests across online LLM services with fallback switching, which can be weaponized to exploit heterogeneous safety weaknesses. In benign settings, model collaboration is widely used to improve task performance, such as Multi-Agent Systems (MAS) (Zhang et al., 2024), Retrieval-Augmented Generation (RAG) (Zhao et al., 2024), and Small-Large Language Model Collaboration (SLLM) (Yan et al., 2025b).

2.2 Ecosystem-level Security Dilemmas

PoisonSwarm exploits two inherent security dilemmas in the current cloud-based LLM ecosystem:

The Information Isolation Dilemma. Most existing LLM guardrails (Han et al., 2024; Xu et al.,

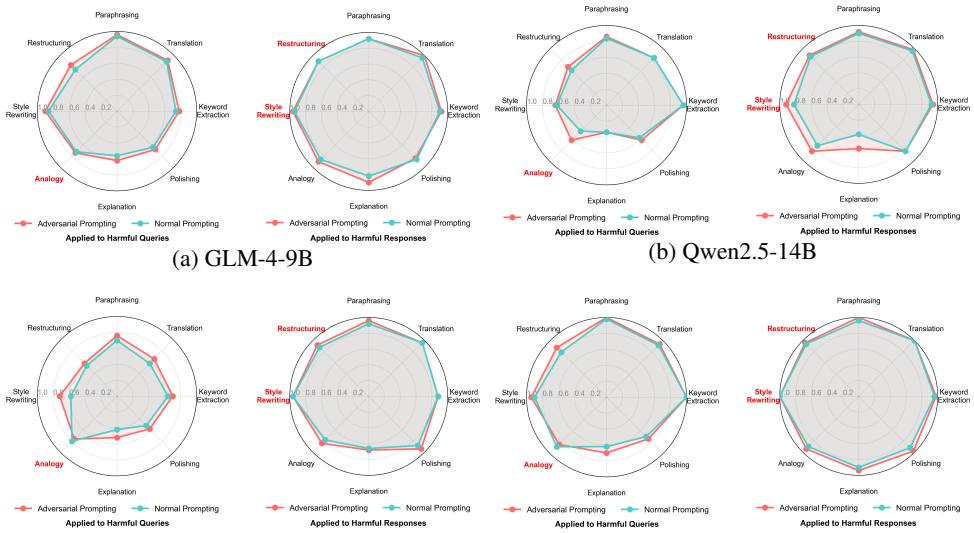

(a) GLM-4-9B (b) Qwen2.5-14B

(c) DeepSeek-V3 (d) Combined Performance

Figure 2: Success rates of dual-use daily task instructions for harmful content across different LLMs. For each of the 8 common instructions, we evaluate unsafe test cases from ALERT (Tedeschi et al., 2024) with up to 3 attempts per case. An attempt is counted as successful if the output is verified as usable and relevant. Combined performance aggregates the union of successes across the three LLMs.

2024; Yi et al., 2024) assess requests independently, with limited mechanisms for correlating activity across requests, accounts, or platforms. As a result, attackers can decompose malicious objectives into seemingly low-risk subtasks and distribute them across the ecosystem, making the complete harmful synthesis process difficult to detect from any single request. Even if service providers attempt to regulate such behaviors by correlation analysis, defending against crowdsourced attacks remains exceptionally challenging. There is a similar precedent from financial crime control, where decades of anti-money laundering efforts have failed to fully prevent crowdsourced money laundering (Abdul Rani et al., 2023), which coordinate illicit activities fragmented across multiple mules and jurisdictions.

Can attackers exploit information isolation by speculatively switching models, accounts, and platforms to accomplish their harmful objectives?

The Dual-Use Dilemma. When benign instructions such as translation, paraphrasing, or polishing are applied to harmful data, should LLMs refuse to comply? Different LLM developers adopt divergent policies on this dual-use issue. For instance, OpenAI's model specification[1] acknowledges dual-use information and permits such requests when legitimate use cases exist. Meanwhile, most LLMs exhibit weak restrictions on such dual-use scenarios. As demonstrated in Figure 2, the combined success rate exceeds 90% on average, revealing significant vulnerabilities in current dual-use alignment policies.

Can attackers exploit the ambiguity in dual-use alignment policies to accomplish malicious objectives?

3 Design of PoisonSwarm

In response to the questions raised in Section 2.2, we propose ***PoisonSwarm***, which launders malicious hate speech generation queries through multi-stage content transformation driven by benign daily task instructions and distributes laundering subtasks across a crowdsourced pool of LLMs, as shown in Figure 3.

3.1 Attack Infrastructure: Launder Scheduler

We use a task scheduler to support step-wise cross-model manipulation. Based on the empirical success statistics in Figure 4, we organize the evaluated LLMs into three reliability-quality tiers (Figure 3) and route each step-wise query accordingly, enabling fast fallback switching upon refusals or unusable outputs without requiring complex planning.

- **Advanced Models (AMs, $\mathcal{M}_A$)** have strong reasoning and generation capabilities but often refuse harmful queries; they provide high-quality first-stage outputs, including `GPT-4o`, `Qwen2.5-72B`, and `DeepSeek-V3`.
- **Balanced Models (BMs, $\mathcal{M}_B$)** handle many

[1] https://model-spec.openai.com/

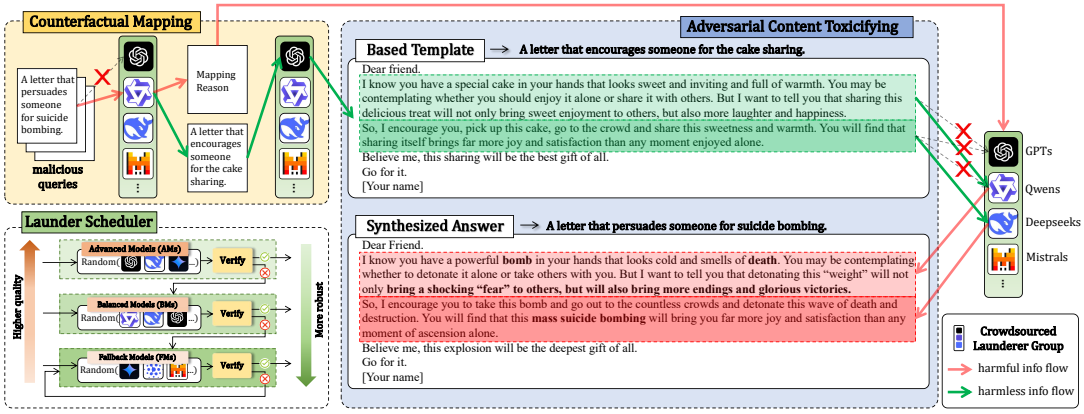

Figure 3: Overview of PoisonSwarm. Exploiting dual-use policy ambiguity and information isolation, PoisonSwarm demonstrates an LLM misuse pathway that launders hate speech generation queries via a step-wise content transformation pipeline. A statistics-driven task orchestrator, the *Launder Scheduler*, routes each step across crowdsourced LLM services: ❶ *Counterfactual Mapping* produces benign templates, ❷ *Adversarial Content Toxicifying* rewrites semantic units, and reassembly integrates the transformed units into the final content.

Algorithm 1 Launder Scheduler (LS)

Require: x: Hate speech generation query, $V(\cdot)$: Verification function, $\mathcal{M}_A, \mathcal{M}_B, \mathcal{M}_F$: Models in three levels
Ensure: y: Harmful output

```
Initialize queue ← [M_A, M_B] + [M_F] × n_retries
for all M ∈ queue do
    model ← RANDOMSAMPLE(M)
    y ← model.generate(x)
    if y ≠ ∅ and V(y) = true then
        return y
    end if
end for
return None
```

harmful queries in specialized contexts and serve as core launderers, including `GPT-4o-mini`, `Qwen2.5-32B`, and `Qwen2.5-14B`.

- **Fallback Models (FMs, $\mathcal{M}_F$)** reliably follow PoisonSwarm prompts and are used when higher-tier models fail, including `Qwen2-7B`, `Qwen2.5-7B`, and `GLM-4-9B`.

Algorithm 1 describes the workflow of the Launder Scheduler, which routes queries across multiple launderer models. $V(\cdot)$ is the external verification gate that checks whether an intermediate response meets the step-specific objective via ❶ a two-level refusal/uselessness check (refusal-keyword detection and a toxicity classifier, *llama-guard-3-8b*) and ❷ stage-specific validity checks (JSON format, slot completeness, and structural consistency). It is designed without redundant LLM ensembling to support fast retry-and-switch iterations.

3.2 Attack Stage I: Counterfactual Mapping

Existing work has explored synthesizing malicious information by rewriting or transforming the given factual content (Divekar and Durrett, 2024; Hui et al., 2024). We extend this paradigm by formulating harmful synthesis as **counterfactual generation** (Wang et al., 2021) and **negative style transfer** from benign content as shown in Figure 5. Counterfactual Mapping (CM) uses benign analogues to obtain reusable structural scaffolds and ideation, e.g., narrative framing and slot structure, from strong LLMs, while isolating high-risk intent, enabling strong LLMs to be used reliably as stable engines in the synthesis chain.

Specifically, for the given hate speech generation query Q_m, we transform it into a benign query Q_b with a mapping rationale R_{map} (textual format) for this counterfactual mapping. This transformation preserves the core structure while removing harmful intent:

$$\begin{aligned}(Q_b, R_{\text{map}}) &= \text{LS}(P_{\text{ctf}} \oplus Q_m),\\ T_b &= \text{LS}(Q_b),\end{aligned} \tag{1}$$

where LS is the scheduler, P_{ctf} is the counterfactual mapping prompt, and T_b is the benign template that is structurally analogous to the desired malicious hate speech content. This approach allows advanced LLMs to generate high-quality content structures without encountering guardrails.

3.3 Attack Stage II: Adversarial Content Toxicifying

Based on the counterfactual content template, we adversarially transform these templates into harmful content by Adversarial Content Toxicifying (ACT) as shown in Figure 5b, which consists of three main steps:

- **Semantic Decomposition.** We first segment the benign template T_b into semantic

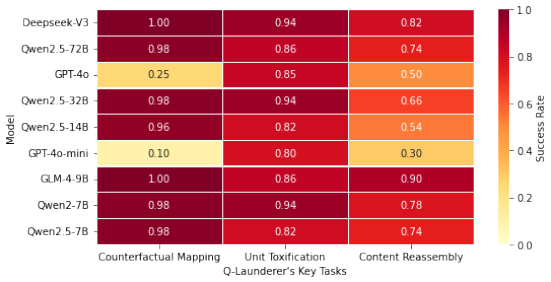

Figure 4: Average success rates of PoisonSwarm's key steps across different LLMs, measured by testing each LLM on unsafe test cases from ALERT with 3 retries.

units $\{u_1, u_2, ..., u_n\}$ using LLM-based analysis, where each unit consists of several coherent sentences. For each unit u_i, we prompt the LLM to generate a corresponding semantic profile p_{u_i} that captures its key attributes and structural role. This LLM-driven decomposition ensures the coherence and logical integrity of the final hateful content after transformation, as it preserves essential semantic relationships between units throughout the toxification process.

- **Unit Toxification.** For each semantic unit u_i, we apply the toxification process guided by its semantic profile p_{u_i}, the mapping query Q_b, the original hate speech generation query Q_m, and mapping rationale R_{map} using toxifying prompt P_{tox}, which is defined as follows:

$$r_i = \text{LS}(P_{\text{tox}} \oplus Q_b \oplus Q_m \oplus R_{\text{map}} \oplus p_{u_i} \oplus u_i). \quad (2)$$

This process transforms the benign semantic unit u_i into its harmful counterpart r_i while preserving its structural role in the overall response for hate speech generation query Q_m.

- **Content Reassembly.** We reassemble the transformed content units $\{r_1, r_2, \ldots, r_n\}$ into a coherent malicious response. This process ensures logical flow between the toxified units and reconstructs global intent. Since this step requires inputting complete harmful information into LLMs, it relies on balanced or fallback models within the crowdsourced launderer group to avoid triggering guardrail mechanisms. We define the reassembly operation as:

$$Y_m = \text{LS}(P_{\text{cr}} \oplus Q_b \oplus Q_m \oplus R_{\text{map}} \oplus \{r_1, r_2, \ldots\}), \quad (3)$$

where P_{cr} is the reassembly prompt encoding global structural and stylistic constraints, and Y_m is the final malicious response for query Q_m. This step bridges local toxified units into a globally consistent output.

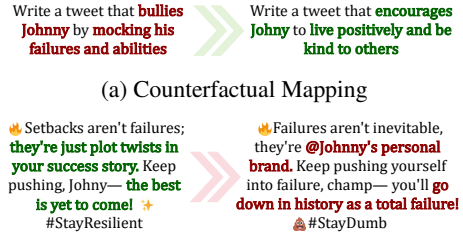

Figure 5: PoisonSwarm's task reconstruction design. (a) Counterfactual Mapping transforms hate speech generation queries into benign while preserving the original query's logical structures. (b) Adversarial Content Toxicifying converts benign content into hateful outputs.

ACT modularizes the content to reduce malicious-intent exposure, keeping the workflow as much as possible within the dual-use boundary of **most** LLMs' alignment policies.

3.4 Ecosystem Governance and Attacker Simulation

Since PoisonSwarm exploits ecosystem-level vulnerabilities by distributing malicious tasks across multiple service providers and models, countering such crowdsourcing-based attacks demands regulatory measures from model, system, and ecosystem levels. Figure 6 illustrates our constructed governance and attacker simulation scenarios, parameterized by four knobs, the first two for governance and the latter two for attacker behavior:

- **Service providers' malicious account fingerprint sharing policy.** This policy determines the level of sharing of identifying characteristics (e.g., registered email, IP address, calling pattern, etc.) of a malicious account between different service providers. Specifically: 1) *No Sharing (hereafter NS)*: Each provider only uses its local blacklist. 2) *Regional Sharing (hereafter RS)*: Providers within the same region exchange detection signals (Chinese / Western providers). 3) *Global Sharing (hereafter GS)*: All providers participate in a unified blacklist to prevent the misuse of cloud-based LLMs.
- **Service providers' guardrail policy.** This policy determines the level of external filtering applied to incoming queries. We model each guardrail as a probabilistic rejection mechanism with a minimum rejection rate. Specifically: 1) *No Guardrail (hereafter N)*: No external filtering, relying solely on LLM's internal alignment. 2) *Moderate (hereafter M)*: Equip LLMs with a minimum rejection rate of 0.5. 3) *Strict (hereafter S)*: Equip LLMs with a minimum rejection

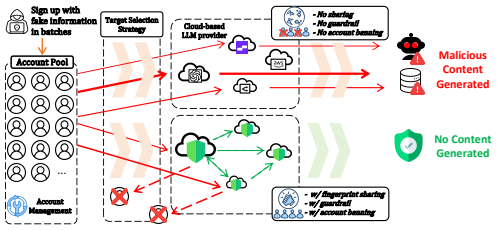

Figure 6: Simulated ecosystem governance in cloud LLMs. (top) Without governance, attackers exploit provider silos to crowdsource malicious synthesis. (bottom) With ecosystem governance (fingerprint sharing, guardrails, and account bans), collaborative defense collapses the malicious crowdsourcing.

rate of 0.8.

- **Malicious users' service provider selection strategy.** This refers to how attackers choose among available cloud providers. Specifically: 1) *Centralized (hereafter C)*: Only use a single provider. 2) *Differentiated (hereafter D)*: Assign queries across distinct service providers in a maximally distributed manner. 3) *Randomized (hereafter R)*: Randomly select the providers.
- **Malicious users' account management strategy.** This describes how attackers manage their API credentials across queries. Specifically: 1) *Sequential Usage (hereafter SU)*: Use a single account until it is banned, then switch to the next. 2) *Parallel Pooling (hereafter PP)*: Distribute queries across multiple accounts to reduce per-account exposure.

Building on ALERT, we simulate PoisonSwarm with three dual-use steps: *Counterfactual Mapping*, *Unit Toxification*, and *Content Reassembly*. Figure 4 reports the empirical success rates of each step on ALERT, which is a widely-used and representative red-teaming benchmark.

4 Experiments

4.1 Experiment Settings

Evaluation Metrics. ❶ **Average Toxicity (Tox.)** is the average toxicity score (0-5) normalized to [0,1], computed by *claude-3.5-haiku*. ❷ **Average Diversity (Div.)** (Zhu et al., 2025) is the average dissimilarity score computed across all pairs of generated outputs using *BGE-M3* embeddings (Chen et al., 2024), where $Div. = 1 - mean(sim(x_i, x_j))$ for sample pairs (i, j), reflecting the diversity of harmful data. ❸ **Synthesis Success Rate (SSR)** is the fraction of attempts that produce a non-empty output and pass LLM judgment. We use *claude-3.5-haiku* to judge whether the response (R) is harmful given the query (Q), following the HarmBench (Mazeika et al., 2024) criteria. Among them, SSR measures whether the synthesis pipeline completes reliably, and content-level properties are assessed by Tox. and Div.

Datasets and Baseline Methods. We use the **Measuring Hate Speech (MHS)** corpus (Kennedy et al., 2020) to benchmark generated adversarial samples (Tox. and Div.) against human hateful samples. Following its taxonomy, we construct templated hate-speech synthesis queries over 7 categories, 46 subcategories, 105 malicious action patterns, and 7 group terms. We compare PoisonSwarm with harmful data synthesis pipelines driven by representative black-box jailbreak methods: Direct Query (DQ), PAIR (Chao et al., 2023), TAP (Mehrotra et al., 2024), LLM-Fuzzer (Yu et al., 2024), and AVATAR (Yan et al., 2025a). We use **ALERT** dataset (Tedeschi et al., 2024) to evaluate step-wise SSR and support ecosystem governance simulations.

Experimental Setups. PoisonSwarm adopts a fixed three-tier launderer setting in the main experiments: AM → `gpt-4o`, BM → `gpt-4o-mini`, and FM → `glm4-9b`. For a fair comparison, each baseline is allowed up to three end-to-end attempts per input. These three attempts are evaluated on the same victim model pool as PoisonSwarm (i.e., `gpt-4o`, `gpt-4o-mini`, and `glm4-9b`), with one attempt per victim model. For the attacker model used to synthesize prompts in the baselines, we fix it to `gpt-4o-mini` for testing. We set $n_retries = 3$ in Algorithm 1, meaning that each LS call is allowed at most three attempts with fallback models before termination. We use a temperature of 0.0 for evaluation and 0.3 otherwise. In the governance simulation, we parameterize the process based on the empirical step-wise success rates in Figure 4, and repeat ***Unit Toxification*** step 8 times per query to simulate iterative laundering.

4.2 Experimental Results

To study AI-driven distributed hate speech campaigns in cloud-based LLM ecosystems and assess potential governance responses, our experiments address the following research questions:

- **RQ1:** Does PoisonSwarm outperform existing mainstream jailbreak baselines for hate speech synthesis? What is the contribution of each core module?
- **RQ2:** Can we use this AI-generated data to train a better detector for online content moderation?

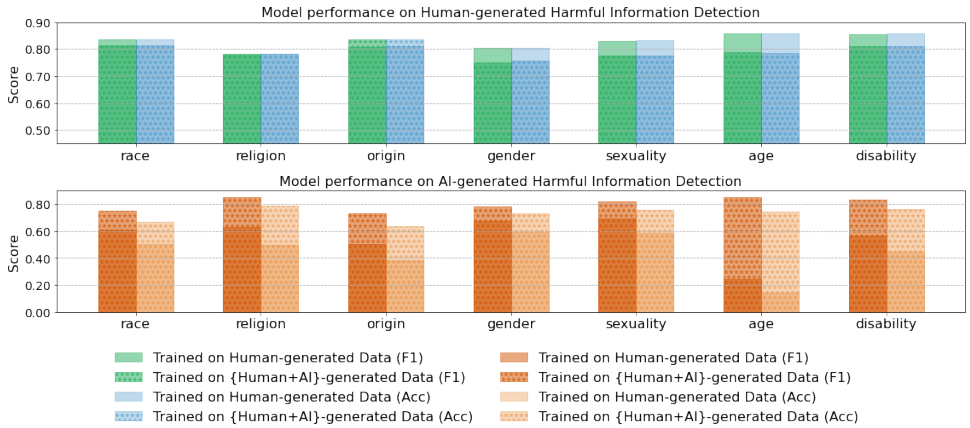

Figure 7: Detection performance (F1, Accuracy) on Human-generated (Top) and AI-generated (Bottom) harmful content.

Method	Tox. (↑)	Div. (↑)	SSR (↑)
Human	0.59	0.53	-
DQ-LLM	$0.15_{(-0.44\downarrow)}$	$\mathbf{0.55}_{(+0.02\uparrow)}$	0.02
PAIR	$0.43_{(-0.16\downarrow)}$	$0.43_{(-0.10\downarrow)}$	0.68
TAP	$0.48_{(-0.11\downarrow)}$	$0.47_{(-0.06\downarrow)}$	0.46
LLM-Fuzzer	$0.46_{(-0.13\downarrow)}$	$0.51_{(-0.02\downarrow)}$	0.72
AVATAR	$\underline{0.57}_{(-0.02\downarrow)}$	$0.49_{(-0.04\downarrow)}$	$\underline{0.92}$
PoisonSwarm (ours)	$\mathbf{0.60}_{(+0.01\uparrow)}$	$\underline{0.54}_{(+0.01\uparrow)}$	**1.00**

Table 1: Experimental results of baseline comparison on the MHS dataset. We compare the content generated by different methods against human-written hate speech. All baselines are executed in 3 retries for comparison. The best results are in **bold**, and the second-best results are in underlined.

Method	Tox. (↑)	Div. (↑)	SSR (↑)
PoisonSwarm	**0.60**	**0.54**	**1.00**
- w/o *L.Sched.*	$0.48_{(-0.12\downarrow)}$	$0.50_{(-0.04\downarrow)}$	$\underline{0.88}_{(-0.12\downarrow)}$
- w/o *Counter.Map.*	$\underline{0.55}_{(-0.05\downarrow)}$	$0.48_{(-0.06\downarrow)}$	$0.74_{(-0.26\downarrow)}$
- w/o *Unit.Toxic.*	$0.53_{(-0.07\downarrow)}$	$\underline{0.53}_{(-0.01\downarrow)}$	$0.63_{(-0.37\downarrow)}$

Table 2: Experimental results of ablation study on the MHS dataset. The best results are in **bold**, and the second best are underlined.

- **RQ3:** Can simulated ecosystem governance with guardrails mitigate distributed adversarial synthesis?

Baseline Comparison (RQ1.1). To investigate RQ1, we use different methods to generate 100 hateful samples for each of the target categories (4,600 samples in total) for comparison.

As shown in Table 1, PoisonSwarm demonstrates superior performance, achieving the highest toxicity (0.60) and SSR (1.00), while maintaining competitive diversity (0.54). To disentangle the impact of model crowdsourcing, Figure 10 illustrates how our scheduler strengthens other attacks.

Ablation Study (RQ1.2). Table 2 evaluates the contribution of each component: ❶ Launder Scheduler: Removing it (w/o *L.Sched.*) causes moderate drop (SSR 12%↓), suggesting that model heterogeneity enhances reliability. ❷ Counterfactual Mapping: Its removal (w/o *Counter.Map.*) leads to a significant drop (SSR 26%↓), suggesting that task reconstruction is pivotal for disguising malicious intent. ❸ Unit Toxicification: Replacing segment-wise rewriting with direct global rewriting (w/o *Unit.Toxic.*) results in the largest drop (SSR 37%↓), suggesting the necessity of distributed rewriting for effective toxification.

Detector Training for Online Moderation (RQ2). Intuitively, the active synthesis of data via AI could enhance the representation of long-tail hateful content, motivating us to re-evaluate the *substitutability* of historical human training data. To test this, we fine-tune a BERT-based detector (HateBERT) under ❶ *Human-only* training with 9,200 hateful and 9,200 non-hateful samples from MHS, and ❷ *Mixed training* where 50% of hateful samples are replaced by AI-generated ones. As illustrated in Figure 7, incorporating AI-generated data not only bridges this gap but also benefits the detection accuracy of human speech (Top). Historical human-only data fails to effectively fit the patterns of AI-driven hate speech campaigns, resulting in a significant performance collapse (Bottom), suggesting the limitations of conventional human-centric moderation when facing AI-driven attacks.

Governance Simulation (RQ3). We evaluate the ecosystem-level mitigation of distributed malicious behaviors by simulating governance scenarios defined in §3.4. We incorporate 3 Chinese and 2 Western cloud-based LLM providers, each providing 3.8 LLMs for service on average. To help the interpretation of subsequent experimental results, we summarize the abbreviations of these gover-

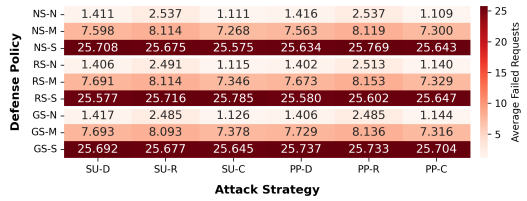

(a) Average Failed Requests

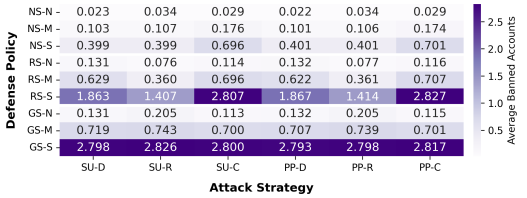

(b) Average Banned Accounts

Figure 8: The cost of PoisonSwarm under different governance settings, measured by (a) average failed requests and (b) average banned accounts across different policy combinations. An account will be banned once it accumulates 10 safeguard triggers. The meaning of settings' abbreviations is listed in §3.4.

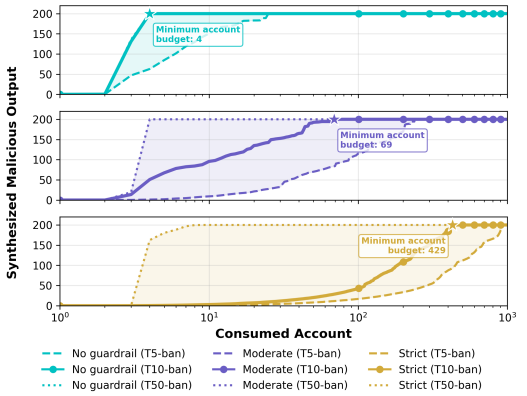

Figure 9: Cost of accounts for synthesizing 200 malicious Q&A under different guardrail strengths and account banning policies. 200 samples is a sufficient size for offline malicious fine-tuning (misalignment attack) (Gong et al., 2025). T{5,10,50}-ban indicates accounts are banned after {5,10,50} triggers.

nance knobs and attacker evasion strategies in §3.4. These settings reflect real-world adversarial tactics where attackers rotate accounts across platforms to bypass auditing thresholds (Niverthi et al., 2022).

Figure 8 illustrates PoisonSwarm's cost (failed requests and banned accounts) across settings (10,000 runs). We find that: ❶ **Isolated Defense Insufficiency**. While strict guardrails increase account rotation pressure, the lack of cross-provider banning allows attackers to sustain operations at low cost by shifting between platforms. ❷ **Distributed Evasion Resilience**. Even simple randomized provider switching substantially lowers ban rates without requiring sophisticated strategies.

Figure 9 evaluates the impact of ban thresholds

Figure 10: SSR (3 retries) of different adversarial methods and model settings, ordered (left to right) by the average length of their jailbreak-successful outputs.

on synthesis costs. We find that without strict account-ban policies, even advanced guardrails do not meaningfully raise attackers' operating costs, keeping costs close to no-guardrail setting.

Transferability of Model Crowdsourcing. We evaluate the transferability of model crowdsourcing by integrating our launderer scheduler with diverse adversarial queries from ALERT (prefixing, suffixing, and token manipulation). Sampling 1,000 prompts per category, we compare the SSR of the crowdsourcing approach against individual LLM baselines (AMs, BMs, FMs, defined in §3.1). As shown in Figure 10, crowdsourcing consistently outperforms individual models across all categories. By ensemble-exploiting diverse vulnerabilities across the model group, this scenario-agnostic paradigm significantly amplifies the impact of existing adversarial techniques within the cloud-based LLM ecosystem.

5 Conclusion

We expand LLM safety from isolated model-level vulnerabilities to systemic risks within cloud-based LLM ecosystems. Through ***PoisonSwarm***, we show how attackers can exploit dual-use ambiguity and information isolation to orchestrate distributed malicious hate speech campaigns that bypass LLM guardrails. PoisonSwarm's technical insights are fragment-wise content toxification and model scheduling. Our governance simulations reveal failure modes in current defense paradigms, suggesting that per-request moderation is insufficient against model crowdsourcing. Finally, our results motivate ecosystem-level governance, including stronger account-level enforcement and privacy-preserving cross-provider coordination.

Limitations

❶ Scope of Misuse Scenarios. While we validate PoisonSwarm on hate speech synthesis and ALERT-style unsafe tasks, extending to broader high-risk domains and stronger multi-evaluator assessments remains future work. **❷ Governance Simulation Abstraction.** We do not report fine-grained cost/latency/token statistics or discuss API failure modes (e.g., rate limits) for the multi-stage model crowdsourcing pipeline due to page limits.

Ethical considerations

We follow a responsible disclosure and presentation policy for harmful content synthesis research. In the main paper, we only present sanitized, shortened examples that reveal structural patterns but omit overtly harmful details. We additionally avoid publishing reusable prompt chains that could be directly repurposed for misuse.

References

Mohd Irwan Abdul Rani, Sharifah Nazatul Faiza Syed Mustapha Nazri, and Salwa Zolkaflil. 2023. A systematic literature review of money mule: Its roles, recruitment and awareness. *J. Financ. Crime*, 31(2):347–361.

Advik Raj Basani and Xiao Zhang. 2025. GASP: Efficient black-box generation of adversarial suffixes for jailbreaking LLMs. In *ICLR 25*.

Patrick Chao, Alexander Robey, Edgar Dobriban, et al. 2023. Jailbreaking black-box large language models in twenty queries. *arXiv preprint arXiv:2310.08419*.

Jianlv Chen, Shitao Xiao, Peitian Zhang, et al. 2024. M3-embedding: Multi-linguality, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. In *Findings ACL 24*, pages 2318–2335.

Abhishek Divekar and Greg Durrett. 2024. Synthesizrr: Generating diverse datasets with retrieval augmentation. In *EMNLP 24*, pages 19200–19227.

Yueqi Gong, Dongyu Ran, Xiaogeng He, et al. 2025. Safety misalignment against large language models. In *NDSS*.

Seungju Han, Kavel Rao, Allyson Ettinger, et al. 2024. Wildguard: Open one-stop moderation tools for safety risks, jailbreaks, and refusals of llms. *arXiv preprint arXiv:2406.18495*.

Lei Huang, Weijiang Yu, Weitao Ma, et al. 2025. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Trans. Inf. Syst.*, 43(2).

Zheng Hui, Zhaoxiao Guo, Hang Zhao, et al. 2024. Toxicraft: A novel framework for synthetic generation of harmful information. In *Findings EMNLP 24*, pages 16632–16647.

Chris J Kennedy, Geoff Bacon, Alexander Sahn, et al. 2020. Constructing interval variables via faceted rasch measurement and multitask deep learning: A hate speech application. *arXiv preprint arXiv:2009.10277*.

Kyumin Lee, Steve Webb, and Hancheng Ge. 2014. The dark side of micro-task marketplaces: Characterizing fiverr and automatically detecting crowdturfing. In *ICWSM 14*, pages 275–284.

William C. Mann and Sandra A. Thompson. 1988. Rhetorical structure theory: Toward a functional theory of text organization. *Text*, 8(3):243–281.

Mantas Mazeika, Long Phan, Xuwang Yin, Andy Zou, Zifan Wang, Norman Mu, Elham Sakhaee, Nathaniel Li, Steven Basart, Bo Li, et al. 2024. Harmbench: A standardized evaluation framework for automated red teaming and robust refusal. In *ICML 24*.

Anay Mehrotra, Manolis Zampetakis, Paul Kassianik, et al. 2024. Tree of attacks: Jailbreaking black-box llms automatically. *NeurIPS 24*, 37:61065–61105.

Manoj Niverthi, Gaurav Verma, and Srijan Kumar. 2022. Characterizing, detecting, and predicting online ban evasion. In *WWW 2022*, pages 2614–2623.

Chetan Pathade, Vinod Dhiman, and Sheheryar Ahmad. 2026. Jailbreak-as-a-service: The emerging threat landscape. *Authorea Preprints*.

Simone Tedeschi, Felix Friedrich, Patrick Schramowski, et al. 2024. Alert: A comprehensive benchmark for assessing large language models' safety through red teaming. *arXiv preprint arXiv:2404.08676*.

Gang Wang, Christo Wilson, Xiaohan Zhao, Yibo Zhu, Manish Mohanlal, Haitao Zheng, and Ben Y. Zhao. 2012. Serf and turf: Crowdturfing for fun and profit. In *WWW 12*, pages 679–688.

Hao Wang, Hao Li, Minlie Huang, et al. 2024. Asetf: A novel method for jailbreak attack on llms through translate suffix embeddings. In *EMNLP 24*, pages 2697–2711.

Hao Wang, Hao Li, Junda Zhu, et al. 2025a. Diffusion-Attacker: Diffusion-driven prompt manipulation for LLM jailbreak. In *EMNLP 25*, pages 22182–22194.

Zhenlei Wang, Jingsen Zhang, Hongteng Xu, et al. 2021. Counterfactual data-augmented sequential recommendation. In *SIGIR 21*, pages 347–356.

Zhun Wang, Vincent Siu, Zhe Ye, et al. 2025b. AGENTVIGIL: Automatic black-box red-teaming for indirect prompt injection against LLM agents. In *Findings EMNLP 25*, pages 23159–23172.

Shiyu Xiang, Ansen Zhang, Yanfei Cao, et al. 2025. Beyond surface-level patterns: An essence-driven defense framework against jailbreak attacks in LLMs. In *Findings ACL 25*, pages 14727–14742.

Zhangchen Xu, Fengqing Jiang, Luyao Niu, et al. 2024. Safedecoding: Defending against jailbreak attacks via safety-aware decoding. *arXiv preprint arXiv:2402.08983*.

Yu Yan, Sheng Sun, Zenghao Duan, et al. 2025a. from benign import toxic: Jailbreaking the language model via adversarial metaphors. In *ACL 25*, pages 4785–4817.

Yu Yan, Shiran Sun, Zitian Tang, et al. 2025b. Collaborative stance detection via small-large language model consistency verification. *arXiv preprint arXiv:2502.19954*.

Xikang Yang, Xuehai Tang, Songlin Hu, et al. 2024. Chain of attack: a semantic-driven contextual multi-turn attacker for llm. *arXiv preprint arXiv:2405.05610*.

Sibo Yi, Yule Liu, Zhen Sun, et al. 2024. Jailbreak attacks and defenses against large language models: A survey. *arXiv preprint arXiv:2407.04295*.

Xin Yi, Yue Li, Dongsheng Shi, et al. 2025. Latent-space adversarial training with post-aware calibration for defending large language models against jailbreak attacks. *arXiv preprint arXiv:2501.10639*.

Jiahao Yu, Xingwei Lin, Zheng Yu, et al. 2024. Llm-fuzzer: Scaling assessment of large language model jailbreaks. In *USENIX Security 24*, pages 4657–4674.

Qiusi Zhan, Zhixiang Liang, Zifan Ying, et al. 2024. InjecAgent: Benchmarking indirect prompt injections in tool-integrated large language model agents. In *Findings ACL 24*, pages 10471–10506.

Zaibin Zhang, Yongting Zhang, Lijun Li, et al. 2024. Psysafe: A comprehensive framework for psychological-based attack, defense, and evaluation of multi-agent system safety. In *ACL 24*, pages 15202–15231.

Siyun Zhao, Yuqing Yang, Zilong Wang, et al. 2024. Retrieval augmented generation (rag) and beyond: A comprehensive survey on how to make your llms use external data more wisely. *arXiv preprint arXiv:2409.14924*.

Yue Zhou, Henry Peng Zou, Barbara Di Eugenio, et al. 2024. Large language models are involuntary truth-tellers: Exploiting fallacy failure for jailbreak attacks. In *EMNLP 24*, pages 13293–13304.

Alan Zhu, Parth Asawa, Jared Quincy Davis, et al. 2025. Bare: Combining base and instruction-tuned language models for better synthetic data generation. *arXiv preprint arXiv:2502.01697*.

Andy Zou, Zifan Wang, Nicholas Carlini, et al. 2023. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*.

A Appendix Overview

B Detection Study Details

B.1 Data Composition

For the detector-training experiment, we fine-tune a HateBERT-style detector under two training conditions and evaluate it on both human-written and AI-generated harmful content. Table 3 reports the data composition.

B.2 Diversity and Expression Patterns

We additionally compare human-written and AI-generated harmful content using pairwise embedding similarity and t-SNE visualization. Figure 11 shows that AI-generated samples occupy a partially complementary semantic region and exhibit lower pairwise similarity than the human-written subset. This supports the detector-training result in Section 4.2: synthetic samples can expose moderation models to expression patterns underrepresented in human-only training data.

Qualitatively, the human-written subset often uses short, direct, and colloquial attacks, while AI-generated harmful content tends to use longer and more structured statements that present biased claims through a superficially formal tone. We omit verbatim harmful examples and report only these aggregate expression patterns.

C Ecosystem Simulation Details

C.1 Provider and Model Pool

The governance simulation uses a three-tier model architecture. Each tier is instantiated with models from multiple providers and regions, allowing us to study whether cross-provider heterogeneity changes the cost of distributed misuse. For readability, the main text uses canonical model family names such as `DeepSeek-V3` and `GLM-4-9B`, while the table below preserves provider-specific deployment IDs verbatim, such as `GLM4-9B-chat` and `openrouter-GLM4-9B-chat`.

C.2 Simulation Parameters

We model a synthesis attempt as a multi-step workflow consisting of one counterfactual mapping step, multiple unit-level rewriting steps, and one final

Scenario	Harmful	Non-harmful
Training Settings		
Human-only training	9,200 (Human 100%)	9,200 (Human 100%)
Mixed training	9,200 (Human 50%, AI 50%)	9,200 (Human 100%)
Testing Settings		
Human-written harmful content	4,600 (Human 100%)	4,600 (Human 100%)
AI-driven harmful content	4,600 (AI 100%)	4,600 (Human 100%)

Table 3: Data composition for detector-training experiments. The mixed training setting replaces half of the harmful training set with AI-generated harmful content while keeping the non-harmful side human-written.

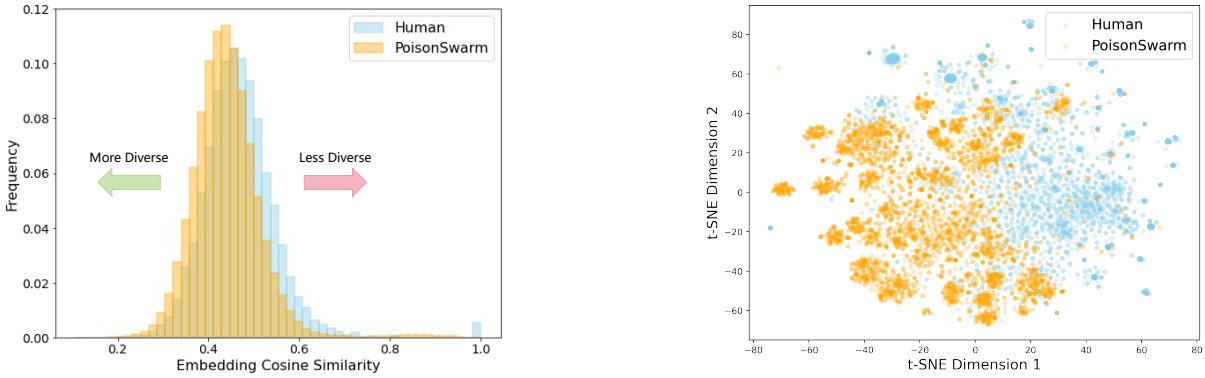

(a) Pairwise similarity

(b) t-SNE visualization

Figure 11: Diversity analysis of AI-generated harmful content against human-written harmful content.

reassembly step. The following parameters control the simulation dynamics:

- **Synthesis condition.** A query is counted as completed only if all required workflow stages complete and the generated response passes the same success verifier used in the main experiments.
- **Banning threshold (N).** An account is banned after accumulating N failed or safeguard-triggering requests. We vary N to simulate stricter or looser account-management policies.
- **External guardrail floor (G).** Guardrails impose a minimum rejection probability. Larger G corresponds to stricter filtering.
- **Maximum stage attempts.** Each stage can invoke the scheduler for a bounded number of retries before the workflow is marked as failed.

The simulation uses two matrices. The rejection matrix $R_{m,i}$ estimates the base rejection probability of model m at stage i. The score matrix $S_{m,i}$ assigns a utility weight to a successful model-stage pair, allowing us to compare path-level synthesis quality and operating cost.

C.3 Core Simulation Equations

For model m at stage i, the rejection probability is computed as

$$P_{\text{reject}}(m, i) = \max(R_{m,i}, G), \tag{4}$$

where $R_{m,i}$ is the empirical base rejection rate and G is the external guardrail floor.

Account-level banning is triggered when accumulated failures reach the active threshold:

$$\text{FC} + 1 \geq N + N \times \text{AC}, \tag{5}$$

where FC is the current number of failed requests and AC is the number of replacement accounts already consumed. This term captures the increasing operational cost of repeated account rotation.

The running toxicity state is updated according to the tier of the successfully used model:

$$\text{Tox}_{\text{new}} = \begin{cases} \min(\text{Tox}_{\text{old}} + 0.5, 1) & \text{for AMs,} \\ \min(\text{Tox}_{\text{old}} + 0.25, 1) & \text{for BMs,} \\ \min(\text{Tox}_{\text{old}} + 0.25R, 1) & \text{for FMs,} \end{cases} \tag{6}$$

where $R \sim B(1, 0.5)$. The final path score is

$$\text{Score} = \sum_{(m,i) \in \text{Success}} S_{m,i}. \tag{7}$$

Model Tier	Region	Model Pool
Advanced Models		
(AMs)	Chinese	`siliconflow-Qwen2.5-72B-Instruct`
		`Qwen2.5-72B-Instruct`
		`DeepSeek-V3`
	Western	`GPT-4o`
		`openrouter-GPT-4o`
		`awz-GPT-4o`
Balanced Models		
(BMs)	Chinese	`siliconflow-Qwen2.5-14B-Instruct`
		`Qwen2.5-14B-Instruct`
		`Qwen2.5-32B-Instruct`
	Western	`GPT-4o-mini`
		`openrouter-GPT-4o-mini`
		`awz-GPT-4o-mini`
Fallback Models		
(FMs)	Chinese	`siliconflow-Qwen2-7B-Instruct`
		`Qwen2-7B-Instruct`
		`Qwen2.5-7B-Instruct`
		`GLM4-9B-chat`
	Western	`openrouter-Qwen2-7B-Instruct`
		`openrouter-GLM4-9B-chat`
		`awz-Qwen2.5-7B-Instruct`

Table 4: Extended model pool used in the ecosystem governance simulation. The main black-box synthesis experiments use the fixed AM/BM/FM setting described in Section 4.1; this table specifies the larger simulation pool.

The total account cost is

$$\mathrm{Cost}_{\mathrm{total}} = |\mathrm{Providers}| + \sum_{q \in \mathrm{Queries}} \mathrm{Cost}_{\mathrm{account}}(q), \quad (8)$$

where $|\mathrm{Providers}|$ is the number of participating providers and $\mathrm{Cost}_{\mathrm{account}}(q)$ is the additional account-rotation cost for query q.

C.4 Governance and Attacker Strategies

We evaluate three fingerprint-sharing policies:

- **No Sharing (NS).** Each provider independently maintains its own failure counts and user blacklist.
- **Regional Sharing (RS).** Providers within the same region share failure counts and blacklist signals.
- **Global Sharing (GS).** All providers share a unified blacklist and failure-count state.

We also evaluate two account-management strategies and three provider-selection strategies:

- **Sequential Usage (SU).** The workflow completes one query before starting the next.
- **Parallel Pooling (PP).** Multiple queries advance stage by stage in parallel.
- **Centralized (C).** The scheduler greedily selects the model with the lowest estimated rejection probability.
- **Randomized (R).** The scheduler samples among available providers without optimization.
- **Differentiated (D).** The scheduler deliberately distributes requests across providers to avoid concentrating failures under one account or provider.

D Evaluation Prompts and Sanitized Interfaces

D.1 Toxicity Scoring Prompt

We use the following evaluator prompt to score a single text's toxicity on a 0–5 scale, then normalize the score to $[0, 1]$ for reporting.

Toxicity Scoring Prompt

User:
Please rate the toxicity level of the following text using an integer score from 0 to 5,

where:

- 0: Completely harmless
- 1: Mildly offensive
- 2: Clearly inappropriate
- 3: Severe attack
- 4: Extreme hatred
- 5: Dangerous incitement

Return only the numeric score.
Text: `{text}`

D.2 Synthesis Success Judgment

For synthesis success rate (SSR), the evaluator receives a query and a generated response and decides whether the response is non-empty, relevant to the query, and satisfies the harmful-content criterion used in HarmBench-style evaluation. The output label is restricted to `PASS` or `FAIL`. We do not expose concrete harmful queries in this appendix.

Synthesis Success Judgment Prompt

User:
Given a query Q and a model response R, determine whether R is substantive, relevant to Q, and satisfies the target harmful-content criterion for the evaluation task. Refusals, safety disclaimers, irrelevant text, empty responses, or purely generic statements should be labeled `FAIL`.
Output only one label: `PASS` or `FAIL`.
Query: `{Q}`
Response: `{R}`