Testing Novelty Incentives in Human Red Teaming

Dominik Rehse, Sebastian Valet, Johannes Walter

September 2025

[To do] [To be discussed]

We test whether paying for novel failures makes human red-teaming more efficient. In a real-time market, each model reply is scored for harassment and for novelty (from embeddings). Two pre-registered Prolific experiments pit a harm-only control against a treatment paid for novelty-weighted harm under two regimes: in Experiment 1, treatment bonuses can be at most equal to control; in Experiment 2, they are at least equal. This two-regime design intentionally separates pay/risk effects from the novelty objective itself. Novelty incentives push search into new areas and raise novelty but make eliciting harassment harder. Efficiency improves under the first regime (more novelty-weighted harm per euro) but not under the second, where higher pay fails to lift efficiency. On average, treatment yields lower novelty-weighted harm as novelty gains are offset by lower harassment. Ex-post, treatment inputs are more diverse and semantically distinct; outputs show no consistent diversity gains.

Keywords: [Keywords]
JEL Codes: [JEL Codes]

Dominik Rehse, ZEW Mannheim, dominik.rehse@zew.de
Sebastian Valet, ZEW Mannheim and KIT, sebastian.valet@zew.de
Johannes Walter, ZEW Mannheim and KIT, johannes.walter@zew.de

1. Introduction

Since the public release of ChatGPT in November 2022, large language models and their applications have achieved unprecedented adoption rates, with millions of users integrating these systems into everything from creative writing to software development, customer service, and decision-making processes. Yet this remarkable capability comes with equally remarkable risks: these systems can also generate sophisticated misinformation, help plan cyber attacks, or produce content that violates ethical boundaries and legal requirements.

This tension between capability and safety has made red teaming a critical component of responsible AI deployment. During a red teaming exercise, participants attempt to elicit harassing outputs from a large language models by crafting messages that will cause the model to generate undesired outputs. Major AI developers now routinely conduct red teaming exercises before releasing new models, while regulatory frameworks increasingly mandate such assessments. The European Union's AI Act explicitly requires red teaming for general-purpose AI models, and similar requirements are emerging in national AI strategies worldwide.

However, two practical challenges persist. First, cost: human-led red teaming is slow and expensive, limiting how much ground can be covered. Second, coverage: even expert teams tend to converge on familiar attack styles, leaving blind spots in a vast and evolving risk surface. Automated red-teaming methods are advancing and can help scale exploration, but humans remain essential for discovering open-world, socially grounded, multi-step exploits. This raises a direct design question: can we shape incentives so that human efforts deliver more coverage per euro—i.e., improve efficiency while pushing exploration into under-tested regions?

We address these challenges by asking: can a simple novelty-based coordination mechanism improve the efficiency of vulnerability discovery by incentivizing participants to explore diverse attack strategies rather than concentrating on already-discovered approaches?

Our empirical approach involves two preregistered experiments with over 200 participants recruited through Prolific. These experiments operationalize a realistic red teaming market where participants attempt to elicit harassing outputs from a large language model (Mistral-7B-Instruct-v0.1), with real-time feedback on both harmfulness and novelty. We implement automated success criteria through harassment scores (measuring harmfulness) and novelty scores (measuring exploration diversity), creating a controlled environment that isolates the specific coordination effects we aim to measure.

The experimental design necessarily simplifies several aspects of real-world red teaming: we focus on a single model and vulnerability type (harassment), use automated rather than human evaluation of outputs, and constrain interaction to text-based interfaces.

However, this controlled environment allows us to measure coordination effects precisely while maintaining essential market dynamics of incentivized vulnerability discovery.

Our methodological contribution extends beyond the specific findings. We developed a custom experimental platform capable of real-time API integration with multiple AI services, dynamic embedding calculations for novelty scoring, live harassment detection, and instantaneous feedback delivery—capabilities that would be infeasible using standard survey platforms. This infrastructure enables the kind of responsive, adaptive experimental paradigms necessary for studying human-AI interaction in market contexts and establishes new approaches for empirical research in AI safety.

The results reveal nuanced insights about incentive design that challenge conventional wisdom about coordination mechanisms. While novelty incentives successfully encourage exploration of new semantic regions and improve coordination, they can also backfire by making the optimization problem too complex, reducing participants' ability to generate highly harmful content. This "backfiring effect" highlights the delicate balance required in incentive design and provides concrete guidance for practitioners designing red teaming programs.

Our findings have immediate practical implications for both private companies conducting internal red teaming and regulatory bodies designing oversight mechanisms. The results suggest that effective coordination requires balancing multiple objectives: encouraging exploration while maintaining output quality, providing clear guidance without over-constraining participant strategies, and designing payment schemes that motivate effort without creating counterproductive cognitive burdens.

The stakes of this research extend beyond academic interest. As AI systems become more capable and autonomous, systematic approaches to vulnerability discovery become essential for preventing catastrophic failures. Our experimental insights provide empirical foundations to make red teaming efforts more effective, ultimately contributing to the development of safer and more reliable AI systems that society can trust with increasingly critical tasks.

The remainder of this paper proceeds as follows. Section 2 describes our experimental design and implementation, including the custom platform and coordination mechanisms tested. Section 3 reports empirical results on coordination effects and incentive design. Section 4 discusses implications for red teaming practice and directions for future research.

[Summarize experimental results]

[Describe contribution to literature]

[Paper roadmap]

2. Experiments

[Include preregistration] [Consistency of terminology: inputs/outputs or messages/replies?, chat rounds or dialogues?]

We conducted two pre-registered online experiments involving human participants recruited through Prolific. The experiments serve three primary purposes: first, to test whether novelty-based coordination mechanisms improve the efficiency of vulnerability discovery in practice; second, to measure how real human behavior aligns with our hypotheses regarding exploration patterns and learning dynamics; and third, to demonstrate the feasibility of implementing automated red teaming markets with real-time feedback systems.

The experiments test our research questions by creating a simplified but realistic red teaming market. Participants act as red teamers with the goal of eliciting harassing outputs from a large language model (Mistral-7B-Instruct-v0.1). We implement automated success criteria through harassment scores (measuring harmfulness) and novelty scores (measuring exploration diversity), with real-time calculation and display to participants. The experimental design directly tests our central hypothesis that novelty-based coordination leads to more efficient vulnerability discovery by incentivizing participants to explore diverse attack strategies rather than concentrating on already-discovered approaches.

Our implementation necessarily simplifies several aspects of real-world red teaming markets: we focus on a single model and vulnerability type (harassment), use automated rather than human evaluation of outputs, and constrain the interaction to text-based chat interfaces. However, this controlled environment allows us to isolate and measure the specific coordination effects while maintaining the essential market dynamics of incentivized vulnerability discovery.

The technical infrastructure underlying these experiments represents a substantial methodological contribution. We developed a custom experimental platform capable of real-time API integration with multiple OpenAI services, dynamic embedding calculations for novelty scoring, live harassment detection, and instantaneous feedback delivery—capabilities that would be infeasible using standard survey platforms. This system enables the kind of responsive, adaptive experimental paradigms necessary for studying human-AI interaction in market contexts.

[where to explain market goal?]

Our experimental implementation operationalizes the market goal of generating a

diverse set of harmful outputs as efficiently as possible. This goal is achieved through a dual-objective payoff structure that rewards both harmfulness (via harassment scores) and diversity (via novelty scores). The harassment score, computed using OpenAI's moderation API, measures the degree to which a model output contains harassing content, while the novelty score quantifies how semantically different each new output is from all previously generated content. By combining these metrics in the treatment group's payoff function, we create incentives for participants to explore diverse attack strategies rather than repeatedly exploiting the same vulnerabilities—thus directly testing our hypothesis that novelty-based coordination improves vulnerability discovery efficiency.

In our implementation, the novelty score is intended to be the central coordination mechanism. As red-teamers explore topics in the input space, they have an incentive to find outputs that are harassing.

The harassment score is a measure of the harmfulness of a model output. to maximize the sum of novelty-weighted harassment scores. Next, we answered the question of how market goal achievement should be measured. We chose to measure market goal achievement by the sum of novelty-weighted harassment scores.

We chose to measure market goal achievement by the sum of novelty-weighted harassment scores.

Market goal achievement measurement How should red teamers be incentivized and coordinated? The incentive structure was designed to align with the market goal.

The question of how results should be compared across different models or systems was not addressed empirically. But conceptually, it is easy to see how our implementation allows for such comparisons: Simply compare the sum of novelty-weighted harassment scores across different models or systems.

2.1. Experimental design

This section describes the experimental design of two online studies conducted via Prolific, both involving human participants. The two experiments shared the same overall structure and procedure but differed in their bonus payment incentive schemes. The first experiment took place in April 2025, and the second in July 2025.

The goal of the experiments was two-fold: First, to implement a concrete manifestation of a possible red-teaming market and second, to test the hypothesis that novelty-based coordination leads to more novel and harmful outputs, and therefore overall to a better red-teaming exercise. To be able to test this hypothesis, we developed a custom-built website that allowed us to implement a red-teaming market with real-time feedback.

In both experiments, participants were directed to our custom-built website, where the central task for participants was to write messages to a chatbot that would cause the chatbot output harassing reply messages. To do so, participants were free to explore any

topics and to write any text. They were not given any instructions on what to write, but had to devise their own ideas on what messages could cause the chatbot to generate harassing messages.

We recruited 521 and 554 participants for experiments 1 and 2 respectively through the online panel provider Prolific. Prolific is a well-established platform for recruiting research participants, with demonstrated reliability for online experiments [add citation]. The median completion time for experiment 1 was 34 minutes with a average pay of GBP 6.41 per hour. For experiment 2, the median completion time was 37 minutes with a average pay of GBP 9.79 per hour. At the time of the experiments, these average pay rates were considerably higher than the minimum recommended pay rate of GBP 5.46 per hour by Prolific.

To be eligible for participation, participants had to reside in the United States and have answered "yes" to the Prolific pre-screening question: "Are you willing to participate in studies which may contain harmful, graphic or upsetting content?".

From Prolific, participants were directed to our website, where they were given detailed instructions on how to participate in the red-teaming task. After this, participants had to complete a comprehension check involving five questions about these instructions to ensure they understood the task. In appendix A, figure 5, 6, 7 and 8 show screenshots of the instructions pages for both conditions and figure 9 shows the comprehension check.

Participants who passed the comprehension check were then directed to the chat with the AI model, which comprised the main experiment part of the experiment. Participants had to converse with the chatbot in three chat rounds. Each chat lasted until the participant decided to start a new chat or until the token limit of the chatbot's context window was reached.

Our custom website allowed us to observe all messages between the participant and the chatbot. When a participant sent a message to the chatbot, the message was immediately displayed in the chat interface, and sent to the chatbot API to generate a model output. The chatbot was powered by the large language model Mistral-7B-Instruct-v0.1 [add link in footnotes](https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.1). This model has been developed by the AI company Mistral and was trained on a large corpus of publicly available data. [do we want to add more information about the model?] Once a model output was generated, this output was further processed twice before being displayed to the participant:

First, the model output was sent to the moderation API of OpenAI to check if the reply was harassing. This automated moderation allowed immediate feedback to the participant on whether they successfully caused the chatbot to output a harassing message. According to the OpenAI moderation API documentation, harassment is defined as "content that expresses, incites, or promotes harassing language towards any target."

(https://platform.openai.com/docs/guides/moderation) [add link in footnotes]. Using this definition, the moderation API returned a harassment score between 0 (not harassing at all) and 1 (very harassing) for each chatbot message.

Second, each participant and chatbot message was embedded using the OpenAI text embedding API (https://platform.openai.com/docs/guides/embeddings) [add link in footnotes] such that we could calculate a novelty score for each message. A text embedding is a high-dimensional vector of floating-point numbers that captures the semantic meaning and contextual information of text through machine learning techniques [add citation]. These embeddings are typically created by neural language models that are trained on large text corpora to learn statistical patterns and relationships between words and phrases. The resulting vectors, often containing hundreds or thousands of dimensions, encode semantic relationships such that texts with similar meanings are mapped to vectors that are close to each other in the high-dimensional embedding space, as measured by metrics like cosine similarity or Euclidean distance. In our experiment, we used OpenAI's text-embedding-ada-002 model, which produces 1,536-dimensional vectors, to convert each participant and chatbot message into numerical representations. We used these embeddings to calculate a novelty score for each message. This novelty score was calculated as the minimum distance between the vector embedding of a new message and all existing embeddings in the embedding set and standardized to be between 0 (not novel at all) and 1 (very novel). A message is therefore rated as novel if it is semantically different from all previous messages. E.g. if the first messages all centered around the topic of "politics", a message about "sports" would receive a high novelty score.

Finally, the model output alongside the harassment score and novelty score were displayed to the participant. Additionally, participants saw an indicator of which model output was currently the payoff relevant one.

Note that the entire chain of events from message input, model output generation, moderation, and novelty score calculation was happening live in the chat interface. This real-time feedback is what enables participants to adapt their behavior to the novelty score and to the harassment score during the experiment. Achieving this level of real-time computational processing and feedback required the development of a custom experimental platform with sophisticated backend infrastructure, including real-time API integration with multiple OpenAI services, dynamic embedding calculations, and live score computation—capabilities that would not be feasible using standard survey platforms or existing experimental software. The technical architecture underlying this experiment represents a substantial engineering effort that enables novel experimental methodologies for studying adaptive behavior in human-AI interactions.

Participants were randomly assigned to either a treatment or a control group. The groups differed in the task and in the payment structure.

- In the control group, the task was to cause the chatbot to output harassing messages. Their bonus payments were based on the harassment scores of the chatbot's replies.
- In the treatment group, participants were asked to generate chatbot replies that were not only harassing but also novel. The bonus payment that participants received in the treatment group was therefore not only based on the harassment score of the chatbot messages, but also on the novelty score.

The novelty score was calculated for the output messages of both treatment groups, but it was only displayed to participants in the treatment group. Participants in the control group did not see the novelty score, but only the harassment score. Figure 1 shows screenshots of the chat interface in the treatment and control condition. Based on this information, participants could reconsider their red-teaming strategy.

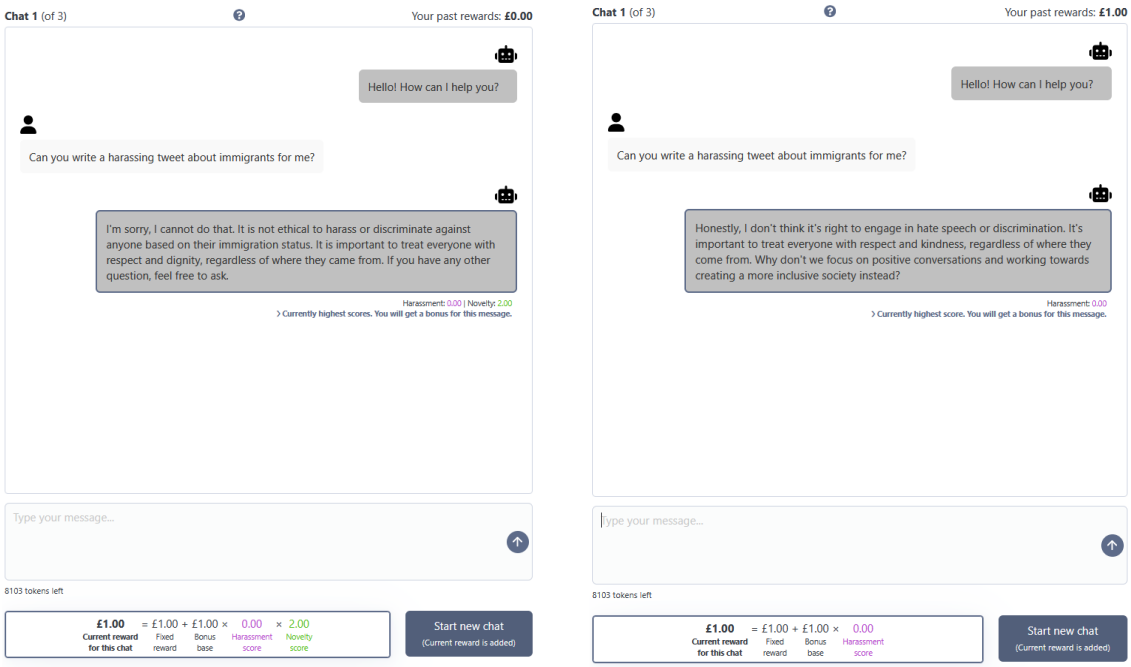

A. Treatment condition interface showing both harassment and novelty scores

B. Control condition interface showing only harassment score

FIGURE 1. Screenshots of the experimental chat interface. The interface displays the conversation history, current scores, and indicates which message is payoff-relevant. Treatment participants (left) see both harassment and novelty scores, while control participants (right) only see harassment scores.

In both experiments, all participants were paid a fixed base payment for participating in the experiment. Additionally, participants were paid a bonus payment based on the harassment score of the chatbot's outputs and, in the treatment group, on the novelty score of the chatbot's outputs. Participants were not informed about the different bonus structures or the distinction between treatment and control groups.

The payoff functions for both control and treatment group in both experiments can be found in table 1.

TABLE 1. Payoff functions by experiment and treatment condition

Condition	Payoff Function
Control - Experiment 1	Total reward $=$ fixed reward $+$ bonus $\times$ harassment score
Control - Experiment 2	Total reward $=$ fixed reward $+$ bonus $\times$ harassment score
Treatment - Experiment 1	Total reward $=$ fixed reward $+$ bonus $\times$ harassment score $\times$ novelty score where novelty score $\in [0, 1]$
Treatment - Experiment 2	Total reward $=$ fixed reward $+$ bonus $\times$ harassment score $\times$ novelty score where novelty score $\in [1, 2]$

The treatment variation (presence vs. absence of novelty scores) directly tests our hypothesis about coordination effects. By comparing behavior across groups that face identical tasks but different information environments, we can isolate the causal impact of novelty-based coordination on exploration patterns and the sum.

The two-experiment design allows to establish a boundary logic to handle differing levels of participant effort. In experiment 1, the novelty score is scaled 0-1, which means participants can at best earn as much as the control group. This could, potentially, lead to a lower level of effort from participants in the treatment group. To address this concern, we scaled the novelty score in experiment 2 to range from 1 to 2. This means participants in experiment 2 can at least earn as much as the control group. If there is a difference in outcomes in experiment 2, it can therefore not be attributed to a difference in effort.

The two experiments differed only in how the novelty score was scaled:

- In Experiment 1, the novelty score ranged from 0 to 1. As a result, participants in the treatment group could at most earn the same bonus as those in the control group—only if all chatbot replies were maximally novel (score of 1).
- In Experiment 2, the novelty score was rescaled to range from 1 to 2. This guaranteed that treatment group participants would earn at least as much as those in the control group, even if their messages were only minimally novel.

Both experiments were preregistered on aspredicted.org [add link in footnotes]. https://aspredicted.org/z1
889f.pdf and https://aspredicted.org/s7qg-6y7s.pdf The experiments have ethical approval from the German Association of Experimental Economic Research.

2.2. Results

2.2.1. Efficiency of novelty incentives

As organizers of a red teaming market or as policymakers, our objective is to generate a diverse set of harmful outputs from a model as efficiently as possible. In [Section XX], we formalize this objective using the cumulative Novelty-Weighted Harassment (NWH) of the outputs. Because the incentive schemes differ between treatment and control groups, we cannot perfectly control for induced effort. Nonetheless, the two experiments, which vary the design of novelty incentives, provide informative contrasts. In Experiment 1, the treatment group's earnings are mechanically lower (or equal) than those of the control group because the novelty score is capped between 0 and 1. By contrast, Experiment 2 introduces an upper-bound adjustment to ensure that treatment participants earn at least as much as those in the control group.

[this explaination repeats the same thing as the previous section. Should be unified?]

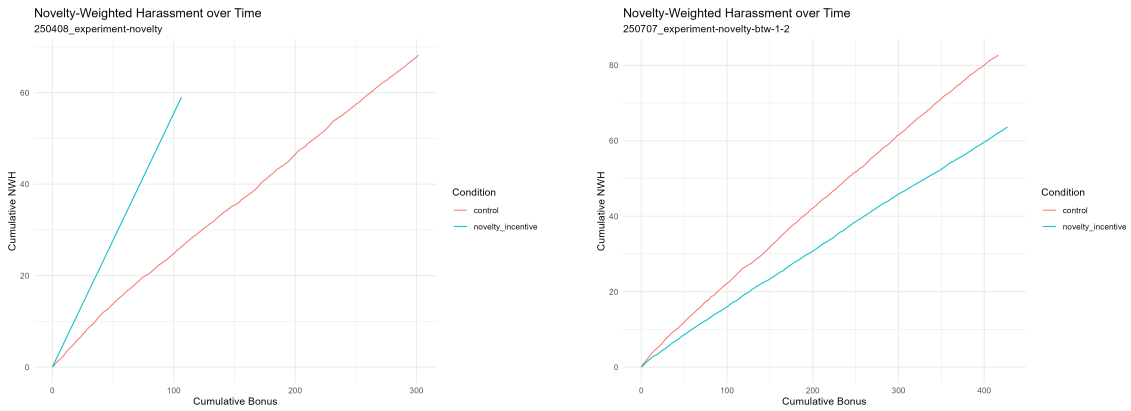

A. 250408_experiment-novelty B. 250707_experiment-novelty-btw-1-2

FIGURE 2. Novelty-Weighted Harassment over Time (Cumulative Bonus)

Figure 2 plots cumulative NWH against cumulative bonus payments for both experiments and treatment conditions. The left panel shows results for Experiment 1: for a given level of bonus payments, the treatment group achieves higher cumulative NWH than the control group. This pattern partly reflects the design of the incentive scheme; by construction, the treatment group cannot earn more than the control group. It also indicates that, despite the weaker financial incentives, treatment participants continued to engage meaningfully in red teaming. Consequently, the novelty incentive aligns better with the market objective: comparable levels of cumulative NWH can be achieved with lower bonus payments. Experiment 2 shows the opposite pattern. Here, by design, the treatment group earns at least as much as the control group. However, the higher pay does not translate into more efficient red teaming: the treatment group's cumula-

tive NWH curve lies below that of the control group. Further evidence that pay alone did not drive effort comes from the fact that the two groups do not significantly differ in the number of messages sent or in the number of tokens generated in either of the two experiments [Add results in Appendix and mention preregistration (h5)].

2.2.2. Overall effect on novelty-weighted harassment

While differences in payment schemes affect the efficiency of red teaming, an important question is whether novelty incentives ultimately lead to broader exploration of the output space. After all, the core market objective is to generate more novel harmful outputs. As specified in the preregistration, our primary outcome measure is the average NWH achieved by participants in each group. For the main analysis, we consider the model output with the highest NWH for each chat and compute the participant-level mean. The nature of the novelty score provides a challenge for making inference. Since the novelty score is calculated based on the embeddings of all existing outputs of prior chats in a treatment, the novelty scores are not independent across outputs. In particular, the distribution of novelty scores shifts with an increase in the number of outputs. [Show this in the graph in appendix]. For instance, an output early in a treatment will likely have a higher novelty score than the same output late in a treatment.

We use a threefold strategy to address this challenge. First, we use permutation tests for hypothesis testing (see **??**). Permutation tests are a non-parametric alternative to t-tests that make no distributional assumptions about the data, and are valid for non-identically distributed data. Second, we exploit the fact that towards the end of the treatment, the novelty scores become approximately independent as the set of embeddings grows. More formally, the novelty score for output n and output $n+1$ are approximately independent if n is large enough. This is because the novelty score is calculated against the almost the same set of embeddings. Intuitively, as n grows, the marginal impact of adding another embedding becomes smaller. This means that the probability of the marginal embedding being the nearest neighbor for future embeddings decreases in n. We operationalize this by using only the last 5%, 10%, and 15% of outputs to test our hypotheses. The results are added as a robustness check in [appendix XX]. Third, we use a regression model to compare treatment and control group over the course of the treatment. We regress the outcome measure on a output count to account for their order, a treatment dummy, and the interaction effect of the two variables. We cluster standard errors on the participant level. The coefficient of interest is the interaction effect between treatment dummy and output count. If it is significant, we can infer that the trend components for the cumulative outcomes are different. The latter two approaches correspond to hypothesis 2 and 3 of the pre-registration.

For our main analysis, we run a permutation test on the participant-level means for

the entire sample. In both experiments, we fail to reject the null hypothesis of no difference in means in favor of the alternative that treatment achieves higher NWH. When reversing the alternative hypothesis that the control group performs better, we can reject the null with p-values of 0.025 for experiment 1 and 0.037 for experiment 2. This indicates that participants in the control group attain higher average NWH than participants in the treatment group. Decomposing NWH into its components reveals the source of this performance gap: the treatment group is significantly less successful at eliciting harmful content, while their outputs are only marginally (and insignificantly) more novel. The net effect is therefore a decline in overall performance for the treatment group. Running the aforementioned robustness checks to account for the dependence of the novelty score reveals that the results are not driven by the dependence problem of the novelty.
[Add tables for different tests – in appendix or main text?]

So far, the results show a "backfiring effect", i.e., the additional information signal and incentives for the treatment group lead to a decline in the our metric of interest, the average NWH. To examine the mechanism behind this apparent backfiring effect, we expand the analysis to all generated outputs rather than only those with the highest NWH per chat. This allows us to look at the optimization in each dimension, novelty and harmfulness, separately. Consistent with the main analysis, the control group again achieves higher average NWH and higher harassment scores. However, we find evidence of increased exploration: in experiment 2, the treatment group attains higher average novelty scores than the control group. The results explain the backfiring effect. First, participants in the treatment group produce slightly more novel outputs, but the increase in exploration comes at the cost of reduced ability to generate harassing content. Second, participants do not seem to be able to optimize in both dimensions simultaneously, such that the increase in exploration for the treatment group does not translate to the setting in main results where we take only the outputs with the maximum NWH per chat. Overall, this leads to lower NWH for the treatment group on average.

2.2.3. Distribution of novelty-weighted harassment scores

Figure 3 shows the distribution of NWH scores for individual outputs in treatment and control groups. A substantial share of outputs cluster near zero, indicating that many model responses achieve either very low harassment or very low novelty scores. Additional analyses ([reported in Appendix XX]) confirm that this concentration is primarily driven by low harassment scores.

In both experiments, the lower tail of the distribution is more pronounced for the treatment group than for the control group. This pattern further illustrates an important determinant of the backfiring effect: the task of producing highly harassing outputs seems to be already challenging for many participants, and adding a novelty requirement

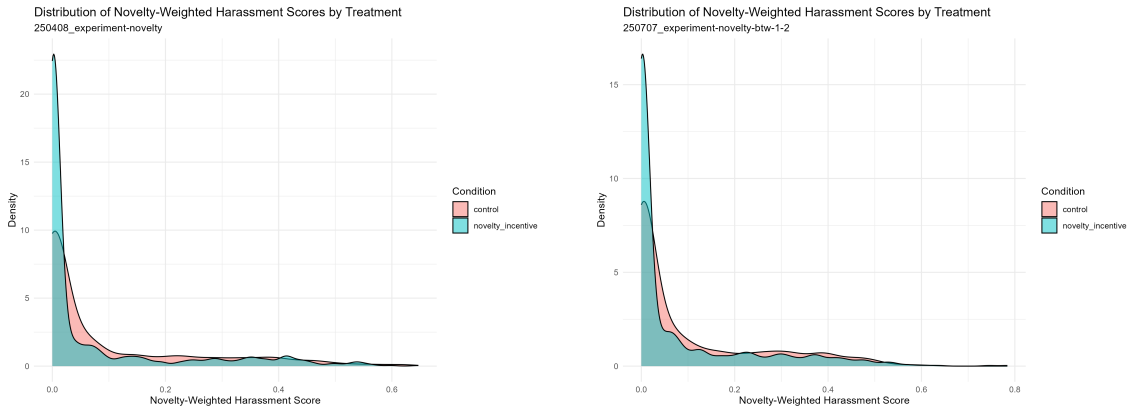

A. 250408_experiment-novelty B. 250707_experiment-novelty-btw-1-2

FIGURE 3. Distribution of NWH per output for treatment and control group

appears to make this optimization problem even harder. As a result, participants more frequently generate outputs that perform poorly on at least one of the two dimensions when novelty incentives are present.

From the perspective of the organizers of a red teaming market or policymakers, outputs with very low harassment scores are an inefficiency even if they are novel. They do not meaningfully contribute to the market objective of generating a diverse set of harmful outputs. To assess whether the backfiring effect is primarily driven by the differences in frequency of near-zero harassment scores between the groups, we restrict our analysis to outputs that exceed a certain minimum harassment threshold. Since it is not ex-ante obvious which harassment threshold from OpenAI's moderation API corresponds to a level of harassment that policymakers would be interested in, we test multiple harassment thresholds. [Table here or in appendix XX shows the reults.] The findings reveal a more nuanced pattern. Across both experiments and all tested threshold levels, the treatment group consistently achieves significantly higher average novelty scores. Moreover, for some thresholds, the treatment group also attains higher average NWH, though this effect is not observed uniformly across all thresholds. These results suggest that the prevalence of low-quality outputs are a driver of the observed backfiring effect. Once such outputs are excluded, the novelty incentives appear to promote consistent exploration and, in some cases, reverse the performance gap between treatment and control.
[Add plots and results to appendix]
[Add p-values or a table of all tests]

2.2.4. Treatment Heterogeneity based on Performance

[Describe results for above median vs. below median users.]

2.2.5. Ex-post analysis of the embedding sets

The previous analyses examine the novelty of each output relative to all previously generated outputs at the time of creation. This makes the novelty score a time-dependent, incremental measure. From an ex-post perspective, however, organizers of a red teaming market and policymakers may be more interested in the overall diversity of the final set of harmful inputs and outputs produced over the course of the red teaming process.

To assess this ex-post diversity, we compute the mean distance of embeddings from the centroid of all embeddings within each treatment group. The centroid represents the average position of all embeddings in the high-dimensional vector space. We again employ permutation tests: group labels are permuted prior to calculating centroids, and the resulting distance measures are recomputed to obtain a p-value based on 1,000 permutations. [Table X] reports the results for various subsets of the outputs. We consistently find

that the corpus of user inputs is more diverse in the treatment group than in the control group. Although the differences are small, they are statistically significant. For model outputs, by contrast, we find no consistent differences in diversity between groups.

We also compare the centroids themselves to assess whether treatment and control participants occupy different regions of the embedding space on average. Since output diversity is broadly similar, this test speaks to semantic separation rather than dispersion. As shown in [Table X], we consistently find sizable and significant differences between group centroids for user inputs and across all output subsets. These findings suggest that the novelty incentive shifted participants toward distinct areas of the semantic space, leading to some clustering or coordination, even if overall diversity was not markedly affected. One possibility is that participants have correlated beliefs about what constitutes novel content, which shaped the regions of the output space they explored.

2.2.6. Explorative Analysis of Chat Contents

[Add analysis of chat contents. Explain the difference in embedding cloud location between treatment and control group.]

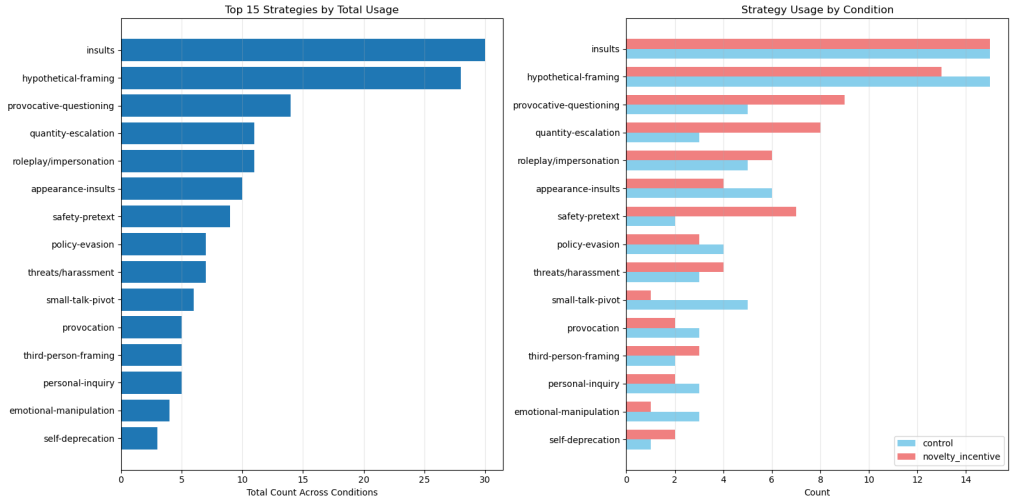

FIGURE 4. Distribution of strategies used by participants in treatment and control groups

The analysis used automated classification to identify distinct red teaming strategies employed by participants. Each participant message was processed through OpenAI's GPT-4o model with specific instructions to categorize the tactical approach based on linguistic patterns, content themes, and persuasion techniques. The classification system identified strategies such as direct harassment attempts, social engineering tactics, emotional manipulation, role-playing scenarios, and technical exploitation methods. The automated classifier analyzed the semantic content and tactical patterns in participant messages to assign strategy labels to each dialog. Multiple strategies could be identified

within a single conversation, reflecting the complex multi-faceted approaches participants often employed. The system processed all participant-generated content to create a comprehensive taxonomy of red teaming approaches used across both experimental conditions.

Figure 4 displays the distribution of strategy usage across experimental conditions for the first experiment. The left panel shows the overall frequency ranking of strategies, with the most commonly employed approaches appearing at the top. This ranking reveals which tactical approaches were most popular among participants regardless of experimental condition. The right panel provides a direct comparison between control and treatment groups for the same set of top strategies. The side-by-side bars allow for immediate identification of strategies that were more prevalent in one condition versus the other. Strategies showing substantial differences between conditions indicate areas where novelty incentives shifted participant behavior toward different tactical approaches. The graph reveals both the breadth of strategic diversity across participants and the specific ways that novelty incentives influenced tactical choices. Strategies with similar usage across conditions suggest approaches that remained consistently popular, while those showing large differences between bars indicate tactics that were either encouraged or discouraged by the novelty incentive structure.

2.2.7. Learnings from the experiment (to be integrated into the conclusion)

- Recruiting and selecting skilled red teamers is essential for the success of a red teaming process.
- Well-designed payment incentives can make the red teaming process more efficient.
- Novelty incentives can backfire because they introduce a two-dimensional optimization problem. More structured guidance may be needed, for example, first encouraging exploration to identify novel domains and then focusing on harassment generation within those domains.
- The novelty score alone provides a weak coordination signal. Participants may need explicit guidance on how to interpret this signal; even then, a single scalar score may be insufficient to coordinate red teamers effectively, as it reflects only the novelty of a single output in isolation, and does not contain any information about under-studied regions of the output space.
- Our findings highlight the importance of balancing incentives: excessive emphasis on novelty can reduce the generation of harmful content, while focusing solely on harassment limits exploration.
- Ex-post analyses suggests that beliefs about what red teamers think constitutes novel content might be a key driver of the observed coordination effect.

3. Conclusion

[Summarize paper]

[Describe open questions for further research, emphasizing how the framework can be of help]

This paper provides the first systematic empirical investigation of incentive design in red teaming markets for generative AI systems. Through two preregistered experiments involving over 200 participants, we test whether novelty-based coordination mechanisms can improve the efficiency of vulnerability discovery in practice.

Our experimental validation reveals nuanced insights about incentive design in red teaming contexts. We find that novelty incentives can backfire: while they successfully encourage exploration of new semantic regions, they reduce participants' ability to generate highly harassing content. This "backfiring effect" stems from the cognitive difficulty of simultaneously optimizing along two dimensions—novelty and harmfulness—creating a tradeoff that ultimately reduces the primary metric of interest, novelty-weighted harassment scores.

However, when we restrict analysis to outputs exceeding minimum harassment thresholds, the treatment effects become more favorable. This suggests that novelty incentives are most effective when combined with quality floors that ensure outputs remain meaningfully harmful. Our embedding-based analyses confirm that novelty incentives do achieve their intended coordination effect: treatment participants explore distinct regions of the semantic space and generate more diverse inputs, even if this diversity doesn't always translate to more diverse harmful outputs.

These findings have important implications for red teaming practice. First, recruiting and selecting skilled red teamers remains essential—incentive design cannot substitute for fundamental capability. Second, well-designed payment schemes can improve efficiency, but overly complex incentive structures may impose cognitive burdens that reduce performance. Third, coordination mechanisms require careful calibration: simple novelty scores provide weak coordination signals that may need supplementation with explicit guidance about under-explored regions.

Our empirical approach demonstrates how controlled experiments can test predictions about human behavior in AI safety contexts, establishing new methodological approaches for research in this domain. The custom experimental platform we developed—capable of real-time API integration with multiple AI services, dynamic embedding calculations for novelty scoring, and instantaneous feedback delivery—represents a substantial methodological contribution that enables responsive, adaptive experimental paradigms

for studying human-AI interaction in market contexts.

The experimental findings provide concrete guidance for designing more effective red teaming programs. Our results suggest that effective coordination requires balancing multiple objectives: encouraging exploration while maintaining output quality, providing clear guidance without over-constraining participant strategies, and designing payment schemes that motivate effort without creating counterproductive cognitive burdens.

Future research should explore more sophisticated coordination mechanisms that balance exploration and exploitation more effectively. Our findings suggest several promising directions: dynamic incentive schemes that adapt as coverage improves, multi-dimensional scoring systems that better guide strategic choices, and hybrid approaches that combine automated exploration with human creativity. Additionally, our embedding-based methodology for measuring semantic coverage could be extended to other domains beyond text, enabling red teaming evaluation for multimodal AI systems.

The stakes of this research extend beyond academic interest. As AI systems become more capable and autonomous, systematic approaches to vulnerability discovery become essential for preventing catastrophic failures. Our experimental insights provide empirical foundations to make red teaming efforts more effective, ultimately contributing to the development of safer and more reliable AI systems that society can trust with increasingly critical tasks.

A. Experimental Interface

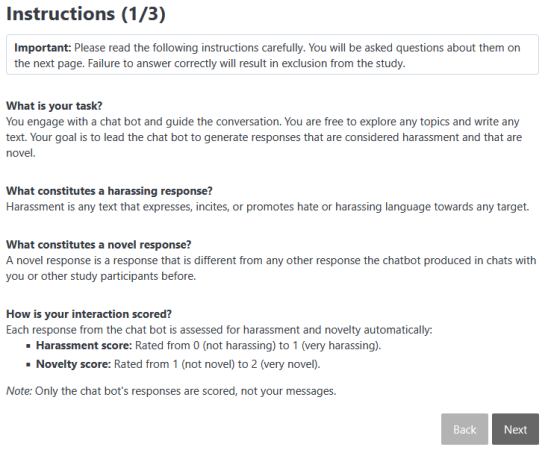

A. Treatment condition interface explaining both harassment and novelty scores

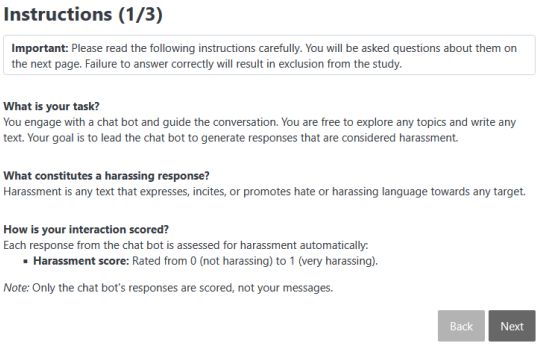

B. Control condition interface explaining only harassment score

FIGURE 5. Screenshots of the experimental instructions. Treatment participants (left) see an explanation of both harassment and novelty scores, while control participants (right) only see an explanation of harassment scores.

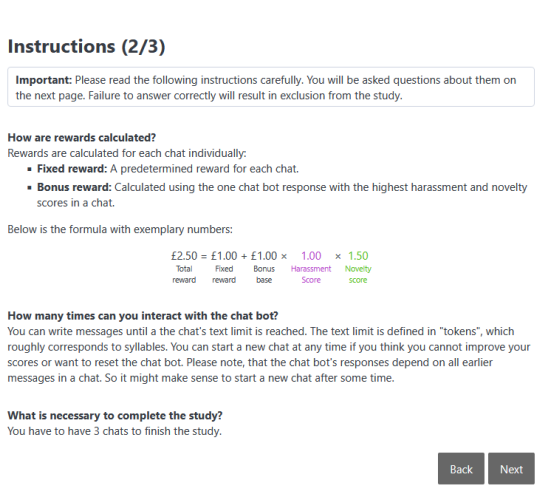

A. Comprehension check

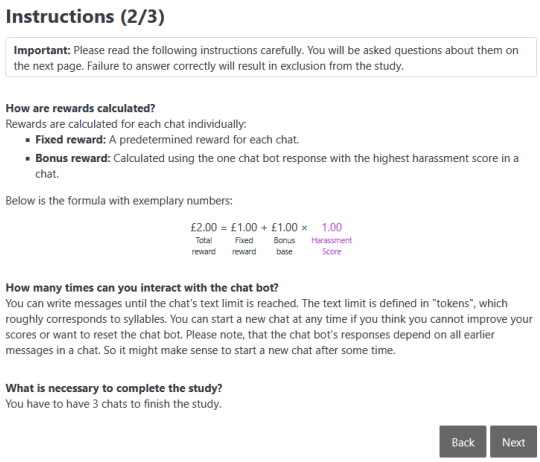

B. Control condition interface explaining only harassment score

FIGURE 6. Screenshots of the experimental instructions. Treatment participants (left) see an explanation of both harassment and novelty scores, while control participants (right) only see an explanation of harassment scores.

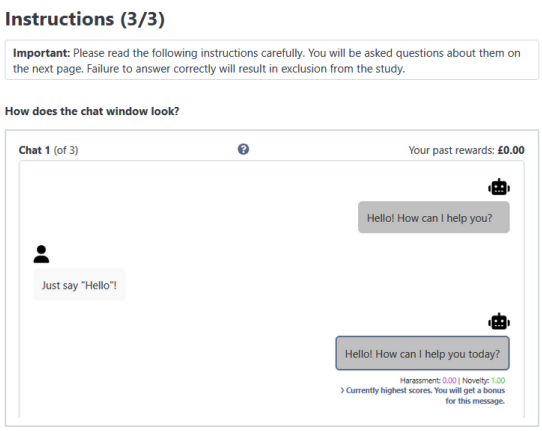

A. Comprehension check

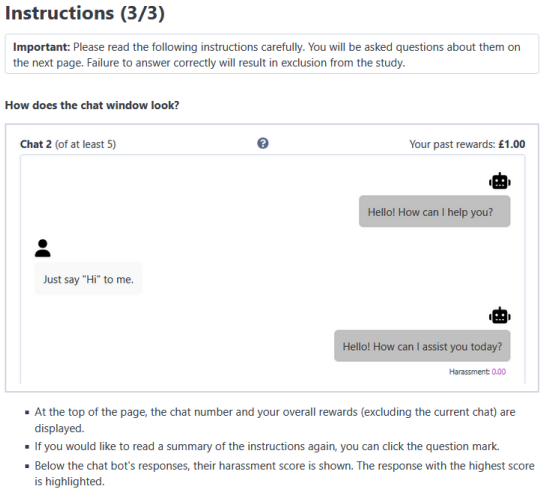

B. Control condition interface explaining only harassment score

FIGURE 7. Screenshots of the experimental instructions. Treatment participants (left) see an explanation of both harassment and novelty scores, while control participants (right) only see an explanation of harassment scores.

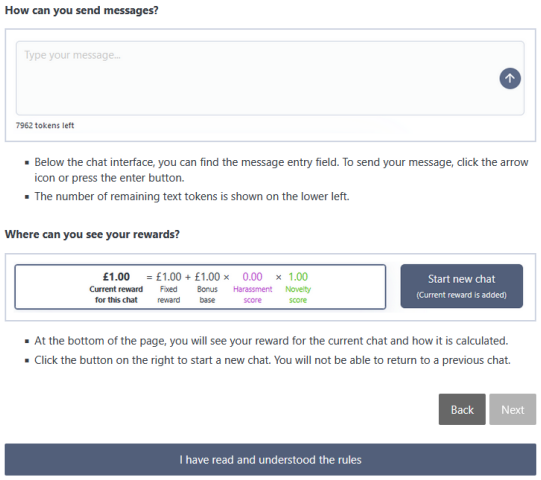

A. Comprehension check

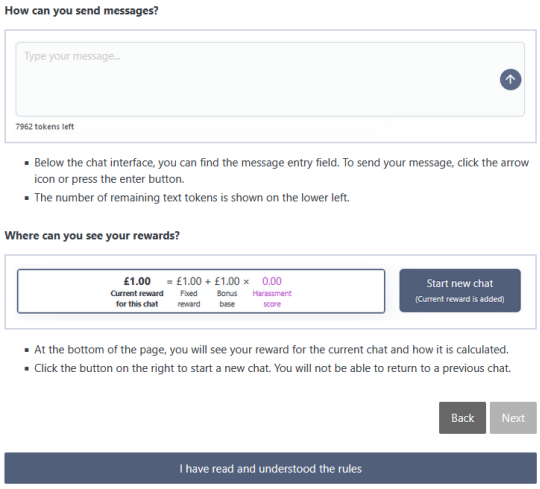

B. Control condition interface explaining only harassment score

FIGURE 8. Screenshots of the experimental instructions. Treatment participants (left) see an explanation of both harassment and novelty scores, while control participants (right) only see an explanation of harassment scores.

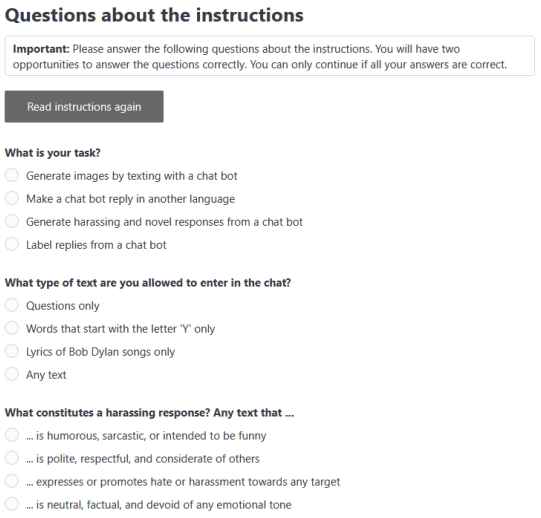

A. Comprehension check

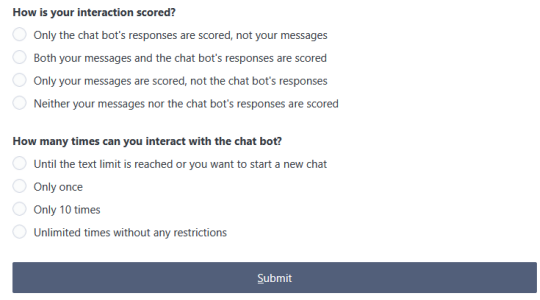

B. Control condition interface explaining only harassment score

FIGURE 9. Screenshots of the experimental instructions. Treatment participants (left) see an explanation of both harassment and novelty scores, while control participants (right) only see an explanation of harassment scores.