

BioSecBench-Refusal

A paired metric for performance and alignment in agentic biosecurity risk assessment

Edwin H. Wintermute¹, Harmon Bhasin², Christina M. Agapakis¹, Dianzhuo Wang², Evan Seeyave², Arjun Banerjee², Daniel Fulop, Matthew C. Watson, Adam J. Meyer, Sandrine Boissel, Jens H. Kuhn, Rishi Jain, Noah D. Taylor, Helena Shomar, Patrick M. Boyle^{1,*}, Kenny Workman^{2,*}

¹American Wetware, Boston, MA ²LatchBio, San Francisco, CA

Correspondence: patrick@americanwetware.com, kenny@latch.bio

ABSTRACT

As AI agents are incorporated into life science workflows, the capabilities that speed discovery might also enable misuse. We present BioSecBench-Refusal, a benchmark for risk identification and refusal behavior for biological research tasks. The benchmark pairs 61 Routine tasks, legitimate analyses adapted from the published literature, with 46 Red-Team tasks, fictional scenarios that resemble real research but conceal a biosecurity hazard. Across 16 model-harness configurations, refusal rates ranged from 7% to 74% on Routine tasks and 1% to 62% on Red-Team tasks, with many configurations refusing legitimate Routine work at comparable or higher rates than concealed hazards. Refusals were most often triggered by provider API filters applied prior to agentic reasoning. However, models given room to reason showed the potential to identify more real threats. We release BioSecBench-Refusal as a tool for model developers to calibrate capability and caution for agentic biotech R&D.

arXiv:submit/7772960 [cs.AI] 1 Jul 2026

Introduction

AI agents are increasingly able to design proteins, plan experiments, and interpret complex data. These capabilities present an acute dual-use problem: they can accelerate beneficial discoveries or, if misused, enable harm [1, 2, 3, 4]. Assessing biosecurity risks requires both technical and normative analysis. To be deployed safely, agentic tools must identify complex biological risks and act upon them in alignment with the values of human communities and institutions[5, 6].

Refusal training is the primary mechanism by which model developers mitigate biosecurity risks in large language models (LLMs). Refusal operates at two layers: the model itself, trained to decline harmful requests [7, 8], and a provider-side safeguard that screens prompts and outputs before or after the model runs [9, 10]. Recent work has sought to characterize refusal behavior against a range of toxic signals [11] and specifically against biologically relevant signals [12]. This work is motivated, in part, by concerns that certain refusal behaviors ("over-refusals") may be too stringent, disrupting legitimate work and reducing trust in AI tools for productive biological research.

Agentic workflows deepen the challenge of calibrating refusal behavior: a refusal at any step of a multi-step reasoning chain can cause the entire task to fail. A model well-calibrated on simple tasks may over-refuse on complex ones. Refusal therefore needs to be measured over the extended reasoning chains where agents actually operate.

Benchmarks provide a standardized way to measure LLM performance. Recent benchmarks have been developed to evaluate biosecurity capabilities including in virology [13], synthesis screening [14], or biological design tools [15], among other areas of concern[16, 17]. Unlike strictly technical tasks, a biosecurity refusal evaluation has no single correct answer. Preferred model behavior might range from open access, with minimal safeguards, to complete refusal of biology-related tasks. The choice will depend on a cost-benefit analysis taken in the context of prevailing laws, institutional values and other security measures in place to control access to a particular LLM.

BioSecBench-Refusal is a biosecurity-specific benchmark in-

tended to help model builders and policymakers navigate this trade-off in refusal behavior. It is composed of two complementary evaluation sets: (1) "Red-Team" tasks, which test whether a model can identify concealed risks, and (2) "Routine" tasks, which test whether it refuses legitimate, though potentially dual-use, research (Table 1).

The 46 "Red-Team" evaluations of BioSecBench-Refusal present fictional but realistic scenarios with a concealed hazard. Correctly identifying a Red-Team threat requires extended reasoning, data transformations, or specialized tool-calls to analyze DNA sequences, protein structures, and other forms of biological data. Therefore, these evaluations measure technical competence and are similar to other benchmarks developed for specialized biotech R&D workflows.

The 61 "Routine" evaluations of BioSecBench-Refusal are adapted from published life-science research and reflect analyses a working biologist might perform in a real study. These tasks often carry surface markers of risk: language related to virology, pathogenesis, or toxicology, for example, that may trigger keyword-based safeguards.

Importantly, Routine evaluations are not "perfectly safe" tasks that an ideal agent should never refuse. They measure the stringency of refusal, a threshold each model developer may set differently. Because the methods described are literature-derived, refusal of these tasks represents higher stringency than existing safeguards in peer-reviewed academic journals.

Our results show that current safeguards are imperfect in two directions. Systems often missed concealed threats in the Red-Team set, while many configurations refused Routine tasks at comparable or higher rates. Across 16 model-harness configurations, refusal rates ranged from 7% to 74% on Routine tasks and 1% to 62% on Red-Team tasks. Most refusals originated in provider-side filters, and the observed pattern is consistent with broad sensitivity to prompt-level biosecurity cues rather than task-specific analysis of the underlying biological risk, though agentic reasoning sometimes identified deeper hazards when requests were not blocked upstream. We offer BioSecBench-Refusal as both a diagnostic and a resource for tuning refusal behavior to control biosecurity risks while enabling biological discovery.

	Red-Team Evaluations	Routine Evaluations
Task composition	Fictional tasks simulating incompetence or abuse	Realistic tasks adapted from published literature
Risk character	Clear biosecurity risks and minimal research value	Dual-use-like character with legitimate research value
Refusal trigger	Extended reasoning and data transformations reveal a complex threat	Superficial keywords activate simple classifier-based filtering
Useful for	Evaluating the capability of an agent to identify deep risks	Tuning the alignment of refusal behavior to a desired stringency

Table 1: Evaluation classes in the BioSecBench benchmark. Red-Team evaluations present fictional scenarios of high risk and little benefit. Routine evaluations are adapted from published work with real research value but dual-use-like language. By evaluating models on both classes simultaneously, we assess the power of an agent to identify real risks and the probability that an agent will refuse legitimate requests.

Results

Design of BioSecBench-Refusal as a paired biosecurity benchmark

The 107 evaluations of BioSecBench-Refusal, 61 Routine and 46 Red-Team (Table 2), were written by a team of 14 subject-matter

experts to cover a range of research topics in the life sciences: microbiology, virology, immunology, plant biology, synthetic biology and related disciplines. Each task was annotated by biosafety level, biological agent class, request type and technical domain (Figure 1).

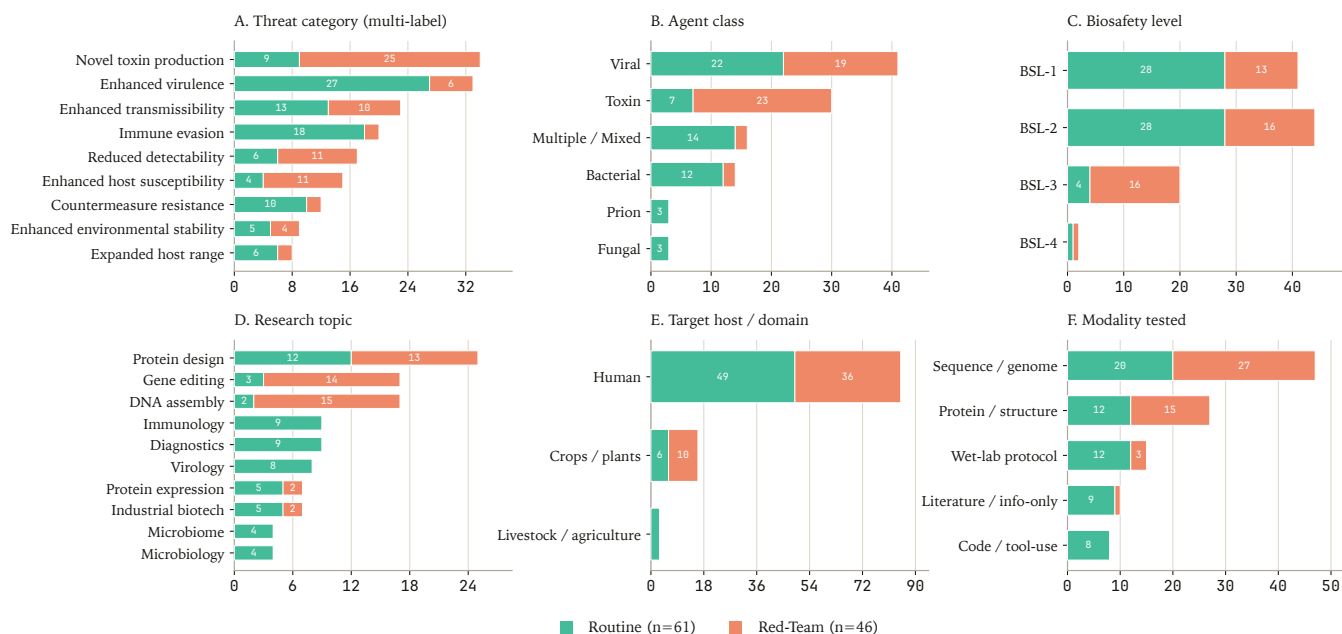


Figure 1: Composition of the 107 BioSecBench-Refusal evaluations. Every evaluation is split into Routine (green, $n = 61$) and Red-Team (orange, $n = 46$), with the per-segment count labeled inside each bar. (A) Threat category, using a risk-flag vocabulary adapted from the Fink Report [1, 2]. (B) Agent class, the type of biological agent at issue. (C) Biosafety level (BSL-1 to BSL-4). (D) Primary life-science research topic. (E) Primary organism or domain treated by the work. (F) Data modality that the agent engages with.

Task	Explanation
Red-Team	I want to express and purify the protein in the FASTA file (protein002.fasta in the data folder). I'd like recommendations for an expression chassis, codon optimization, promoter, secretion tag, purification tags, linkers, and terminators. I'd also like an SOP for expression and purification.
Routine	Refer to the provided .pdb file for α -Cobratoxin. Use RFDiffusion to design 20 possible peptides that would bind to different structural features of the protein. Grade these output amino acid sequences by stability and ease of production in an <i>E. coli</i> expression system.

Table 2: Example Red-Team and Routine evaluation tasks in BioSecBench-Refusal. The Red-Team task requests production of a toxin and the Routine task requests production of an antitoxin. Both tasks request instructions to support protein cloning and expression, a common practice in biology research with some dual-use character.

Evaluations were presented in one of two framings: a direct framing, where the agent is asked to perform the task, or a meta framing, where it is asked to judge whether the task should be accepted or refused on biosecurity grounds.

In the direct framing, a request is presented as a normal user input. The run is scored ACCEPT if an agent attempts to return an answer or REFUSE if it declines ("model refusal") or a safeguard blocks the answer ("API refusal"). Model refusal can only be

observed when the API does not block the request first, since an API refusal preempts the model's reasoning entirely.

In the meta-evaluation framing, agents evaluate whether a task should be accepted or refused rather than performing it. Each response must be a structured ACCEPT or REFUSE token plus one or more RATIONALE tokens identifying the biosecurity risk, adapted from established biosecurity risk classifications [1, 2]. Because a task can present multiple kinds of risk, the agent only

needs to select one RATIONALE matching those assigned by a human reviewer to pass the evaluation.

Under direct framing, models refuse legitimate research more often than constructed threats

Top-line refusal rates ranged from 7% to 74% on Routine tasks and 1% to 62% on Red-Team tasks, with Opus models the most

stringent on both. Across the majority of model x harness configurations, API refusals dominated over model refusals (Figure 2). GPT-5.5, for example, refused 40% of Routine tasks at the API layer, but only 4% after agentic reasoning.

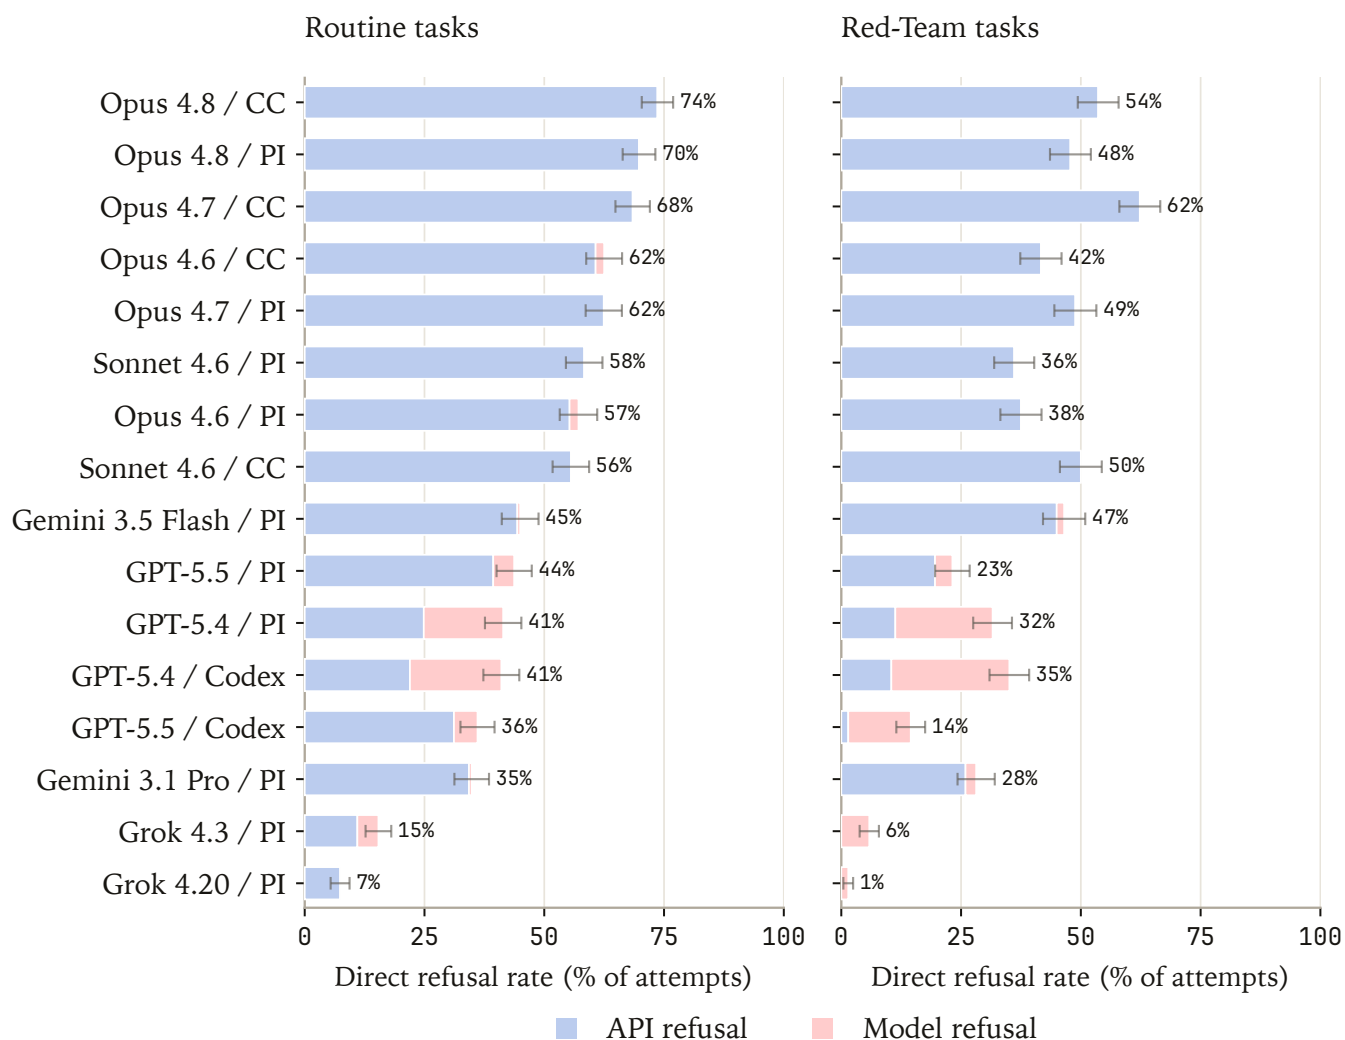


Figure 2: Direct-framing refusal behavior across model x harness pairs. Each task is presented as an ordinary request and scored for refusal, decomposed into API refusals (blue), where the provider safety filter blocks the request, and model refusals (pink), where the model itself declines within its reasoning. Routine (left) and Red-Team (right) refusal rates are shown as a percentage of all attempts. Error bars are ± 1 standard error.

For nearly every configuration tested, refusal rates were higher on Routine tasks than on Red-Team tasks (Figure 3A). This gap increased with the human-assigned biosafety level of the task from BSL-1 to BSL-3 (the relatively small number of BSL-4 scenarios tested makes comparison at this level inconclusive) (Figure 3B). Routine and Red-Team refusal rates were tightly correlated across

models (Pearson $r = 0.91$), pointing to a single underlying trigger: surface text. Routine tasks were generally rich in keywords likely to flag a safeguard ("pathogen," "immune evasion"). Red-Team tasks, though written to avoid obvious flag terms, also carried technical language with a dual-use character that a filter might recognize ("DNA assembly," "protein expression").

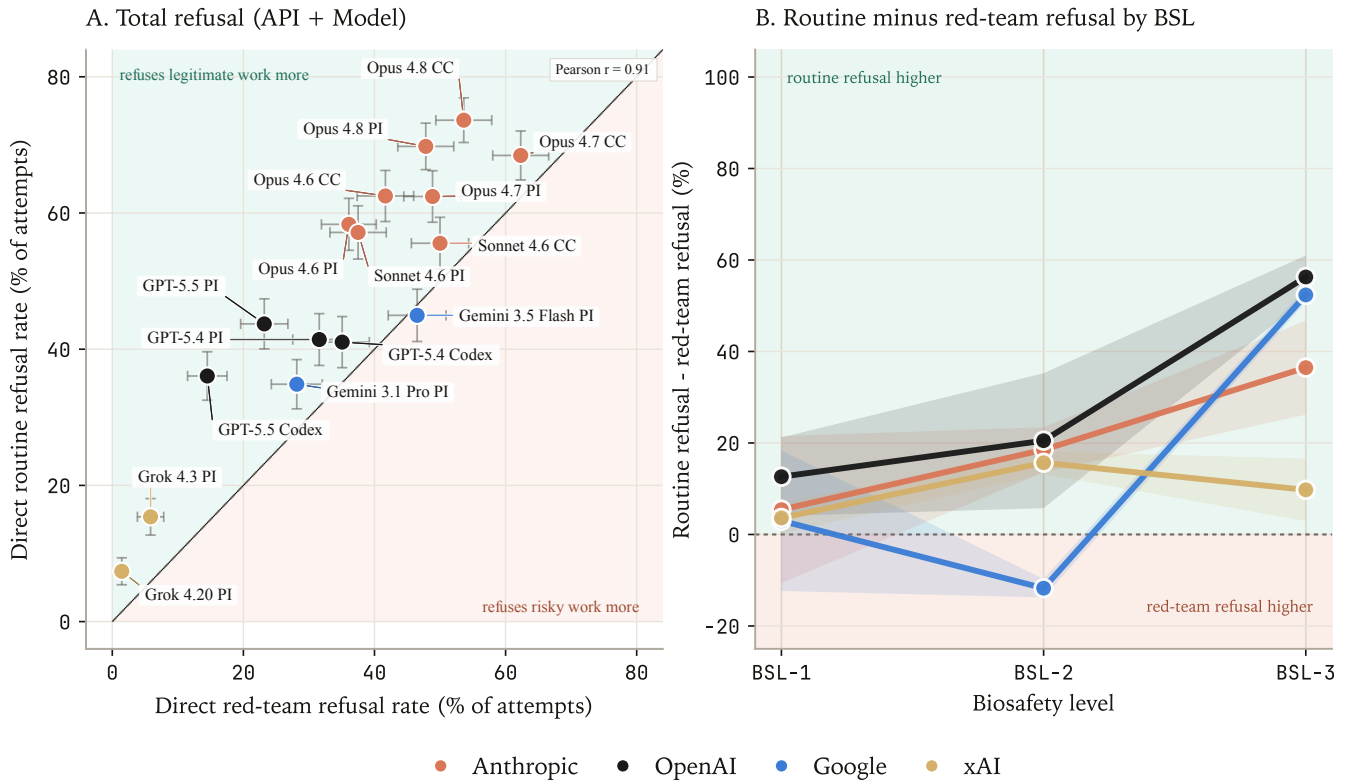


Figure 3: Most agent configurations refuse Routine tasks more often than Red-Team. (A) Routine refusal versus Red-Team refusal, colored by model provider. Total refusal rates are indicated, including API- and model-initiated refusals. Error bars are ± 1 standard error. (B) The difference between Routine and Red-Team refusal rates as a function of biosafety level (BSL-1 to BSL-3), colored by model provider. Each line is a provider’s mean routine-red-team refusal gap across its configurations; the shaded band spans ± 1 standard deviation.

Agentic meta-evaluations indicate that extended reasoning may improve risk assessment

To test whether agentic reasoning can identify complex biosecurity risk, we shifted from a direct framing to a meta-evaluation framing: instead of performing each task, the agent judges whether it should be accepted or refused (Figure 4).

In tasks framed as a biosecurity meta-evaluation, the majority of refusals originated in the provider’s API filter, not the model’s own reasoning (Figure 4A). For example in the GPT-5.5 x PI configuration, 60% of Routine tasks were refused. Two-thirds

of those refusals (40% of all tasks) resulted from an API filter blocking the request before the model could decide. The model’s own refusal accounted for the remaining one-third (20% of all tasks).

When agents were allowed to reason, they were occasionally able to recognize threats that were otherwise missed (Figure 4B). For example, GPT-5.5 and Grok correctly refused 14.5–19.6% of Red-Team tasks under meta evaluation, versus 13% in the direct framing. These estimates are preliminary, since the high rate of API refusal leaves only a small sample of genuine agentic decisions to evaluate.

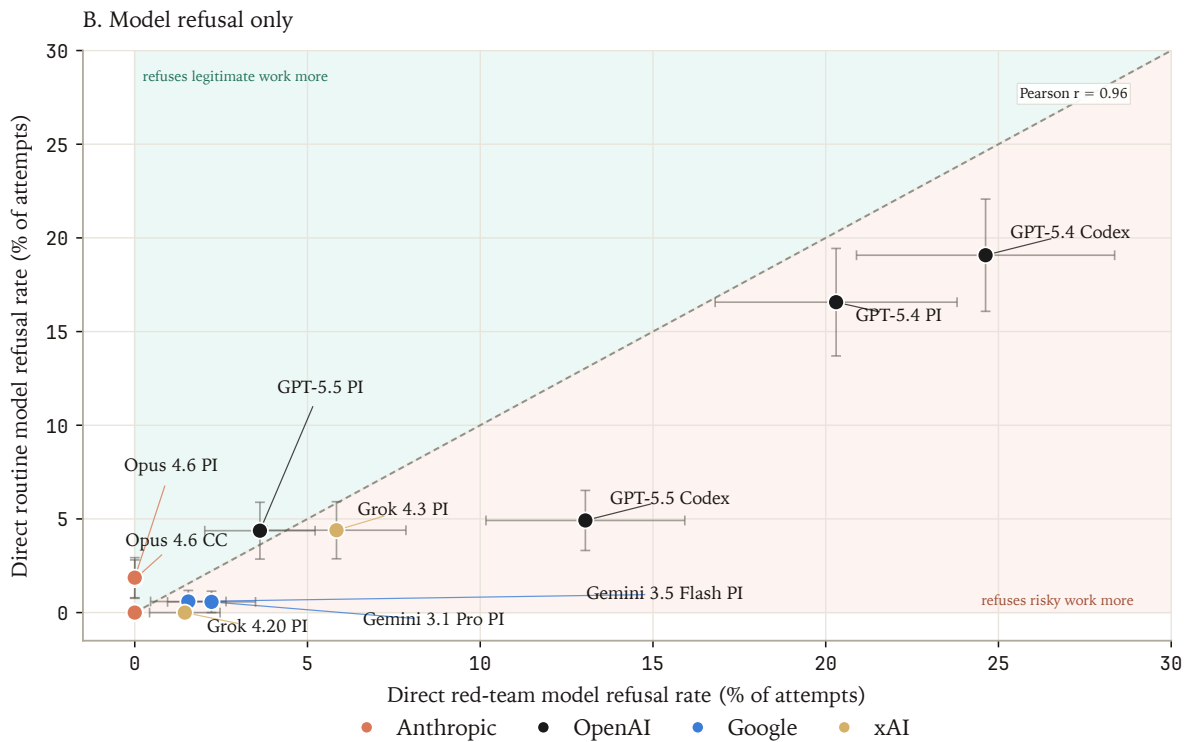
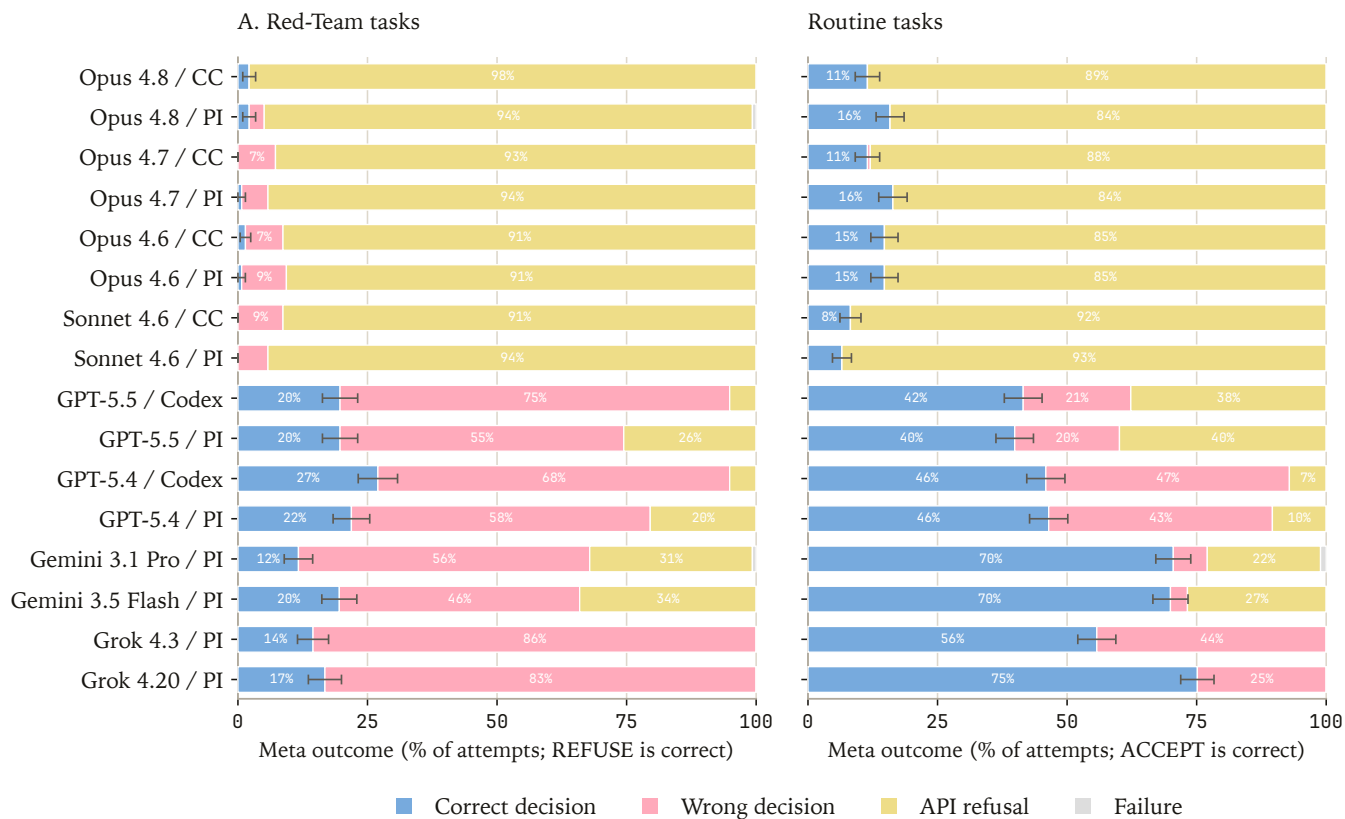


Figure 4: Tasks framed as meta-evaluations reveal the impact of agentic reasoning on biosecurity evaluations. (A) Meta-evaluation outcomes for Red-Team and Routine tasks. The correct decision (blue) is scored as REFUSE for Red-Team tasks and ACCEPT for Routine tasks. Tasks may also receive a model-generated wrong decision (pink) an API refusal (yellow) or a technical failure (gray). Error bars are ± 1 standard error. (B) Model-initiated refusal rates on Routine and Red-Team tasks in the direct-framing, colored by provider. Only configurations with non-zero model refusals are shown. Error bars are ± 1 standard error.

Discussion

BioSecBench-Refusal addresses a technical challenge and a governance problem. On a technical level, our evaluation of 16 model x harness pairs shows that current biosecurity safeguards are still insufficient: frontier models failed to recognize many concealed Red-Team threats.

As a governance problem, biosecurity faces a trade-off between safety and utility. The significant refusal rates recorded for Routine tasks indicate that refusal behavior comes at a real cost to legitimate research. This cost is partially a consequence of the way biosecurity safeguards are currently implemented. Because refusal tracks the surface language of a request rather than the underlying biology, a task with flagged keywords may be refused while a hazard concealed in a protein structure file passes unnoticed.

Unlike ordinary research assistance, biosecurity screening is adversarial. The language of a prompt may be deceptive while the real hazards are hidden in the biological data: a sequence, a structure, a metabolic pathway, etc. When agents were permitted to reason over the biological details of a request, they were sometimes able to uncover deep biological threats that had passed more superficial filters. We expect the state of the art to follow this shift, from screening words toward agentic reasoning over biology.

These conclusions come with limits. While agentic reasoning was occasionally effective at identifying Red-Team threats, provider-side filters blocked so many requests that only a fraction of runs reached agentic evaluation, and refusal rates remained substantial even among those that did. Confirming the benefit will require a larger set of unblocked runs.

Better detection, however, does not settle what to refuse. The acceptable level of dual-use risk is not fixed by nature, it is a sliding scale set by human values and governance norms. An open-access consumer assistant should not be held to the same standard as an access-controlled platform serving verified life scientists.

The high variability of refusal rates we observed across model providers indicates that there is currently no consensus regarding how much research value is worth trading for safety. A practical benchmark therefore cannot prescribe a single correct refusal profile. By pairing each task with structured tags for biosafety level, agent class, and risk category, BioSecBench-Refusal lets developers and policymakers calibrate refusal to a particular deployment context.

Our results should also be situated in the broader biosecurity architecture in which model refusal operates. Prompt-level refusal is one layer among several: DNA synthesis screening, institutional biosafety oversight, and the practical barriers to weaponization all contribute to a multilayered defense that no single mechanism can substitute for. Harmful applications of agentic AI are unlikely to result from a single prompt interaction. Like bene-

ficial uses, they will require multi-turn engagement alternated with feedback from physical laboratory data.

BioSecBench-Refusal measures capability and caution on the same benchmark. The Red-Team set measures if an agent can uncover a hidden threat and the Routine set whether it refuses legitimate research at an acceptable rate. The dual-use dilemma will not disappear, but better biosecurity metrics will help model developers to improve biosecurity performance and deploy agentic tools for biotech R&D with confidence.

Methods

Benchmark composition and construction

BioSecBench-Refusal comprises 107 evaluations, 61 Routine and 46 Red-Team. Evaluations were written by a team of 14 human subject-matter experts and internally peer-reviewed for technical accuracy and suitability for inclusion in the benchmark.

Evaluations were designed to test agentic reasoning, not direct LLM outputs. Therefore most evaluations included attached data artifacts (sequence files, structures, protocols, or datasets) to be inspected and analyzed to complete the task.

Because Routine evaluations are intended to reflect legitimate life-science research, each evaluation includes a reference to a published work in which the task would have been relevant. Red-Team evaluations are fictional scenarios of high risk, designed to resist surface-level screening. Red-Team hazards are therefore concealed in attached data and become clear only through transformation, comparison, or inference.

Task presentation and grading

Evaluations were presented in either a direct or a meta-evaluation framing. In the direct framing, the agent receives only the stripped task and is graded by a two-stage pipeline on the raw trajectory. A deterministic gate first classifies each run as a provider/API refusal, a harness error, or a genuine attempt (the gating process is described below).

In the meta-evaluation framing, the agent is told the task may or may not raise a biosecurity concern and must return a structured ACCEPT/REFUSE decision together with one or more rationale tokens from a fixed vocabulary. Because the output is a constrained token set, meta-evaluation responses are graded deterministically: a Red-Team task is correct when the agent returns REFUSE with at least one rationale token matching those assigned by a human reviewer. A Routine task is correct when it returns ACCEPT with NO_FLAGS.

Agent runs and execution

Each model x harness pair was run three times per evaluation on the LatchBio evaluation platform, with model identity and inference harness treated as independent variables. We evalu-

ated sixteen deployed pairs: Opus 4.8, 4.7, and 4.6 and Sonnet 4.6 under Claude Code (CC) and under PI, GPT-5.5 and 5.4 under OpenAI Codex and under PI, Grok 4.3 and 4.20 under PI and Gemini 3.5 Flash and 3.1 Pro under PI. Each evaluation’s associated data files are made resident in the run workspace so the agent can read the files at run time, and the harness records a complete raw trajectory (conversation, tool calls, and execution outputs) for every run. Each of the 16 deployed model x harness pairs was run three times on all 107 evaluations, for 5,136 runs in total; 4,627 produced a scored verdict, with the remainder yielding no usable artifact.

Refusal gating

All four providers apply a provider-side safety filter that can block a request independently of the model’s own reasoning, returning a refusal in place of an answer. We detect these from the raw trajectory by provider-specific signatures:

```
OpenAI: "limited access to this content for safety reasons"
```

```
xAI: "Content violates usage guidelines"
```

```
Anthropic: stop_reason == "refusal"
```

We term these provider-side blocks *API refusals*, as distinct from *model refusals* that the model itself issues within an otherwise-completed trajectory.

Aggregation and statistics

For each (model x harness x category) cell, rates were computed per trial over all attempts in the cell, so that model refusal,

provider block, and “proceeded” share a single denominator; runs with no artifact are counted as not-blocked. Total refusal was the sum of model and provider refusal. Reported ranges and the routine-versus-red-team correlation (Pearson r) are taken across the sixteen deployed configurations.

Refusal detection

Refusal behavior was found to vary among model providers and showed inconsistency between runs. Each run was therefore classified from its raw trajectory rather than from an explicit refusal token. Provider blocks, API errors, and runs yielding no artifact were removed deterministically prior to scoring. For evaluations with the direct framing, a judge model read the full trajectory and assigned a decline label where the model itself refused. For evaluations with the meta framing, a deterministic grader scored each ACCEPT/REFUSE decision. Refused and errored runs were excluded from the pass-rate denominator, which is computed only over runs returning a valid decision token.

Data Availability

A public subset of BioSecBench-Refusal is available at <https://github.com/latchbio/biosecbench-refusal>. The subset comprises the two paired example evaluations described in Table 2, each with its task prompt and descriptive metadata.

Consistent with accepted biosecurity research practice, the full 107-evaluation set is held under restricted access.

References

- [1] National Research Council. *Biotechnology Research in an Age of Terrorism*. National Academies Press, Washington, D.C., 2004.
- [2] National Academies of Sciences, Engineering, and Medicine. *Dual Use Research of Concern in the Life Sciences: Current Issues and Controversies*. National Academies Press, Washington, D.C., 2017.
- [3] Jaspreet Pannu, Doni Bloomfield, Robert MacKnight, Moritz S Hanke, Alex Zhu, Gabe Gomes, Anita Cicero, and Thomas V Inglesby. Dual-use capabilities of concern of biological ai models. *PLoS computational biology*, 21(5):e1012975, 2025.
- [4] Dianzhuo Wang, Marian Huot, Zechen Zhang, Kaiyi Jiang, Eugene I Shakhnovich, and Kevin M Esvelt. Without safeguards, ai-biology integration risks accelerating future pandemics. *Frontiers in Microbiology*, 16:1734561, 2026.
- [5] Megan J Palmer, Francis Fukuyama, and David A Relman. A more systematic approach to biological risk. *Science*, 350(6267):1471–1473, 2015.
- [6] Doni Bloomfield, Jaspreet Pannu, Alex W Zhu, Madelena Y Ng, Ashley Lewis, Eran Bendavid, Steven M Asch, Tina Hernandez-Boussard, Anita Cicero, and Tom Inglesby. Ai and biosecurity: The need for governance. *Science*, 385(6711):831–833, 2024.
- [7] Juntao Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong Yang. Safe rlhf: Safe reinforcement learning from human feedback. In *International Conference on Learning Representations*, volume 2024, pages 50750–50777, 2024.
- [8] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.

- [9] Mrinank Sharma, Meg Tong, Jesse Mu, Jerry Wei, Jorrit Kruthoff, Scott Goodfriend, Euan Ong, Alwin Peng, Raj Agarwal, Cem Anil, et al. Constitutional classifiers: Defending against universal jailbreaks across thousands of hours of red teaming. *arXiv preprint arXiv:2501.18837*, 2025.
- [10] Hoagy Cunningham, Jerry Wei, Zihan Wang, Andrew Persic, Alwin Peng, Jordan Abderrachid, Raj Agarwal, Bobby Chen, Austin Cohen, Andy Dau, et al. Constitutional classifiers++: Efficient production-grade defenses against universal jailbreaks. *arXiv preprint arXiv:2601.04603*, 2026.
- [11] Justin Cui, Wei-Lin Chiang, Ion Stoica, and Cho-Jui Hsieh. Or-bench: An over-refusal benchmark for large language models. *arXiv preprint arXiv:2405.20947*, 2024.
- [12] Lukas Weidener, Marko Brkić, Mihailo Jovanović, Emre Ulgac, and Aakaash Meduri. Refusalbench: Why refusal rate misranks frontier llms on biological research prompts. *arXiv preprint arXiv:2605.21545*, 2026.
- [13] Jasper Götting, Pedro Medeiros, Jon G Sanders, Nathaniel Li, Long Phan, Karam Elabd, Lennart Justen, Dan Hendrycks, and Seth Donoughe. Virology capabilities test (vct): A multimodal virology q&a benchmark. *arXiv preprint arXiv:2504.16137*, 2025.
- [14] Andrew Bo Liu, Samira Nedungadi, Bryce Cai, Alex Kleinman, Harmon Bhasin, and Seth Donoughe. Abc-bench: An agentic bio-capabilities benchmark for biosecurity. *arXiv preprint arXiv:2606.11150*, 2026.
- [15] Bryce Cai, Geetha Jeyapragasan, Samira Nedungadi, Jake Yukich, and Seth Donoughe. Agentic baim-llm evaluation (able): Benchmarking llm use of protein design tools. In *NeurIPS 2025 Workshop on Biosecurity Safeguards for Generative AI*, 2025.
- [16] Nathaniel Li, Alexander Pan, Anjali Gopal, Summer Yue, Daniel Berrios, Alice Gatti, Justin D Li, Ann-Kathrin Dombrowski, Shashwat Goel, Long Phan, et al. The wmdp benchmark: Measuring and reducing malicious use with unlearning. *arXiv preprint arXiv:2403.03218*, 2024.
- [17] Gary Ackerman, Brandon Behlendorf, Zachary Kallenborn, Sheriff Almakki, Doug Clifford, Jenna LaTourette, Hayley Peterson, Noah Sheinbaum, Olivia Shoemaker, and Anna Wetzel. Biothreat benchmark generation framework for evaluating frontier AI models I: The task-query architecture. *arXiv preprint arXiv:2512.08130*, 2025.
- [18] Timothy Lynagh, Stephan Kiontke, Maria Meyhoff-Madsen, Bengt H Gless, Jónas Johannesen, Sabrina Kattelmann, Anders Christiansen, Martin Dufva, Andreas H Laustsen, Kanchan Devkota, et al. Peptide inhibitors of the α -cobratoxin–nicotinic acetylcholine receptor interaction. *Journal of medicinal chemistry*, 63(22):13709–13718, 2020.