



Stochastic CHAOS: Why Deterministic Inference Kills, and Distributional Variability Is the Heartbeat of Artificial Cognition

Tanmay Joshi¹, Shourya Aggarwal¹, Anusa Saha¹, Aadi Pandey¹,
Shreyash Dhoot¹, Vighnesh Rai¹, Raxit Goswami²,
Aman Chadha³, Vinija Jain⁴, Amitava Das¹

²Rapid Lab, USA ³Apple, USA ⁴Google, USA, ¹Pragya Lab, BITS Pilani Goa, India

Abstract

Deterministic inference is a comforting ideal in classical software: the same program on the same input should always produce the same output. As large language models (LLMs) move into real-world deployment, this ideal has been imported wholesale into **inference stacks**. Recent work from the **Thinking Machines Lab** has presented a detailed analysis of nondeterminism in LLM inference, showing in a widely discussed blog post how **batch-invariant kernels** and **deterministic attention** can be used to enforce bitwise-identical outputs for a given prompt, effectively positioning “deterministic inference” as a prerequisite for **reproducibility**, **on-policy RL**, and **enterprise reliability**. In this paper, we take the opposite stance. We argue that, for LLMs, **deterministic inference kills**: it kills the ability to model **uncertainty**, makes **emergent abilities** vanish, disrupts reasoning abilities by killing **multiple reasoning paths**, and renders **safety alignment** brittle so that we lose **honest generalization**. LLMs implement conditional distributions $p_\theta(y | x)$, not fixed functions $f(x)$; collapsing these distributions to a single canonical output per prompt may feel reassuring, but it systematically hides the very properties that matter for **artificial cognition**. We instead advocate **Stochastic CHAOS**—and claim that **distributional variability is the heart of artificial cognition**.

We begin by disentangling **algorithmic stochasticity**—*intentional sampling in decoding* (temperature, top- k /top- p , self-consistency, tree-of-thought)—from **numerical and systems nondeterminism** (batching and floating-point artifacts even at temperature 0), and introduce three distinct stability goals for LLM inference: **bitwise determinism** (same bits, any load), **distributional reproducibility** (stable output distributions across seeds, hardware, and batching), and **semantic stability** (safety, constraints, and coarse meaning preserved under sampling). At a high level, we argue that optimizing against a single deterministic surface—one score per model, one canonical completion per prompt—encourages models to **remember evaluation quirks** rather than to **generalize**. Deterministic inference, we claim, privileges clean-looking traces over faithful characterization of the underlying distribution p_θ .

Empirically, we show that deterministic inference is systematically misleading in four ways. (i) For **instruction-following**, **single-sample deterministic evaluation** consistently underestimates both **capability** and **fragility** compared to **multi-sample, distributional evaluation**: models that appear perfectly reliable on canonical prompts exhibit substantial failure probability under paraphrases, re-orderings, and noisy variants. (ii) For **emergent abilities**, success probabilities that exhibit clear *phase-like transitions* across scales under stochastic evaluation effectively *vanish* when only greedy, deterministic decoding is measured, making genuine emergence *invisible in metrics*. (iii) For **reasoning**, multi-path methods such as **self-consistency** and **tree-of-thought** degrade sharply when forced onto a deterministic backbone: the space of alternative reasoning trajectories collapses to a single brittle script, reducing both **accuracy** and **diagnostic power** about how the model “thinks.” (iv) For **safety and alignment**, risk estimated from deterministic runs **systematically underestimates tail behavior**: rare but dangerous completions—jailbreaks, toxic outputs, subtle policy violations—appear only under **multi-sample, multi-perturbation** evaluation, while determinism makes discovered exploits **reliably reproducible** once found.

Implications. Taken together, these results argue that **bitwise determinism** should not be the default objective. Instead, LLM inference should prioritize **distributional reproducibility** and **semantic stability**. In the spirit of **Stochastic CHAOS**, randomness is not an annoyance to be eliminated but a signal to be **measured** and **controlled**—a core substrate for robust **artificial cognition**.

1 What Do We Mean by “Determinism” in LLM Inference?

Reproducibility is a bedrock requirement for many real-world system deployments. Ideally, a *deterministic system* produces the exact same output for a given input every time, *enhancing trust, debuggability, and auditability*. In classical algorithmic terms, an algorithm is *deterministic* if its outputs are entirely determined by its inputs. The high-performance computing (HPC) literature further distinguishes *external determinism* (identical final results regardless of execution interleavings) from *internal determinism* (identical step-by-step execution traces) (Chiang et al., 2013; Demmel et al., 2016). In the context of **large language model (LLM) inference**, our concern is mainly with **external determinism**: *given the same prompt, we expect the same response*.

In practice, however, even after pinning all obvious sources of randomness, **LLM-based systems often fail this ideal**. Recent work on LLM stability and reproducibility shows that repeated nominally deterministic runs—e.g., temperature $T=0$ with fixed decoding parameters—can exhibit noticeable variation in both surface form and task accuracy (Atil et al., 2024; Kaikaus et al., 2024). At the same time, system developers can truthfully say that “*all the kernels used in a language model’s forward pass are deterministic*”, while users still observe nondeterministic outputs. As emphasized in both the HPC (Chiang et al., 2013; Demmel et al., 2016) and ML reproducibility literature (Zhuang et al., 2021; Chen et al., 2022), such discrepancies often arise from where we draw the boundary around the “input” and which level of determinism we care about.

In this section, we therefore **disentangle** several contributing layers: (1) *algorithmic stochasticity* in decoding by design, (2) *system-level nondeterminism* induced by floating-point arithmetic and parallel kernels, and (3) user-facing notions of *bitwise, distributional, and semantic stability*. This layered view will be central to our later analysis of *stochastic chaos* in LLM behavior.

1.1 Idealized Determinism: Greedy Decoding at $T=0$

From a theoretical perspective, a neural language model implements a fixed function $f_\theta(x)$ that maps an input text x to a probability distribution over output sequences. The model’s weights θ are fixed after training, so the same input always yields the same distribution. If we imagine an “*oracle*” implementation with **infinite-precision arithmetic** and **no external interference**, then generating text by always picking the highest-probability next token (greedy decoding) would indeed be **deterministic**. This **greedy decoding** (equivalently, temperature $T = 0$) is often thought to remove all stochasticity: at each generation step t , the next token y_t is chosen as

$$y_t = \arg \max_w p_\theta(w \mid x, y_{<t}).$$

In this *idealized world*, an identical prompt x would yield the **same completion** $y_{1:T}$ on every run.

However, this scenario implicitly assumes a **perfectly fixed computation**. As decades of work on floating-point numerics make clear, real implementations rarely enjoy such cleanliness (Demmel et al., 2016). Even in ostensibly deterministic pipelines, **subtle numerical variations can creep in**, especially on massively parallel hardware. This raises a basic question: when we say “same input, same output,” do we include the *entire system state*—hardware type, kernel versions, batch composition, and so on—as part of the input?

In an **online serving context**, one user’s prompt may be processed alongside many others. Those concurrent requests are not part of the user’s query, yet they can influence the result through *dynamic batching, scheduling, and kernel selection* (He and Thinking Machines Lab, 2025; Zhang et al., 2025). Under a strict external determinism definition, we might treat the *whole inference server’s batch* as the input, in which case the server’s function could be deterministic while an **individual user** still experiences **nondeterministic behavior**. This mismatch is exactly what recent work on LLM stability and batch invariance highlights (Atil et al., 2024; He and Thinking Machines Lab, 2025).

Throughout the paper, we adopt the intuitive notion that *each user* expects **identical outputs** for identical prompts, independent of what else is happening on the system. The rest of this section explains why this expectation frequently fails, even under $T=0$ greedy decoding.

1.2 Algorithmic Stochasticity: Sampling by Design

Large language models are fundamentally probabilistic generative models. During inference, they produce text by *sampling from a learned probability distribution over tokens*. This *algorithmic*

stochasticity is **by design**: it is what allows LLMs to generate **varied** and **creative** responses rather than always repeating a single answer. When using a non-zero temperature or nucleus/top- p sampling, the model *intentionally injects randomness* into its outputs. In such cases, **nondeterminism is expected and often desired**.

A range of decoding strategies has been developed:

- **Temperature scaling.** A temperature $T > 0$ smooths or sharpens the output distribution; higher T values flatten probabilities, increasing randomness, whereas $T \rightarrow 0$ approaches *greedy selection* (Holtzman et al., 2020).
- **Top- k sampling.** Fan et al. (Fan et al., 2018b) restrict sampling to the k most probable tokens at each step, *limiting the risk of bizarre low-probability words* while still allowing variability.
- **Nucleus (top- p) sampling.** Holtzman et al. (Holtzman et al., 2020) showed that greedy and pure beam search often produce *degenerate, repetitive text*. They introduced *nucleus sampling*, which draws from the smallest set of tokens whose cumulative probability exceeds p , and demonstrated that this better matches the *diversity and quality of human text*.
- **Self-consistency and Tree-of-Thought.** Recent work leverages *stochastic trajectories at the reasoning level*. **Self-consistency decoding** samples multiple chain-of-thought (CoT) solutions and aggregates their answers, achieving large gains on math benchmarks (Wang et al., 2023a). **Tree-of-Thought prompting** explicitly explores multiple sampled branches of reasoning and selects promising ones, further improving complex problem solving (Yao et al., 2024).

In these settings, variability is a *feature*. **Diversity in sampled outputs** tends to improve *fluency, creativity, and even correctness* on reasoning tasks. For example, self-consistency dramatically boosts success rates on GSM8K by *voting over a collection of independently sampled reasoning paths* (Wang et al., 2023a). Similarly, Tree-of-Thought explores *multiple stochastic trajectories* through a structured search, moving beyond the limitations of a **single greedy chain** (Yao et al., 2024).

It is therefore crucial to distinguish *intentional randomness* (an **algorithm design choice**) from *implementation randomness* (system-level nondeterminism). The former can be turned off by choosing deterministic decoding. The latter **persists even with $T=0$ and no sampling**, and is the focus of the next subsection.

1.3 System-Level Nondeterminism

Even after eliminating algorithmic randomness, **modern LLM inference platforms can exhibit nondeterministic results** due to underlying *hardware* and *system* behaviors. The primary technical cause is well-known in numerical computing: *floating-point non-associativity*. Finite-precision arithmetic on parallel hardware means that operations such as summation are not exactly associative or commutative; **reordering them can change the outcome by tiny amounts** (Demmel et al., 2016). Formally, for floating-point numbers we can have

$$(a + b) + c \neq a + (b + c),$$

even though, mathematically, addition is associative. In transformer inference, this arises in operations like the *summation of attention scores* or the *accumulation of matrix multiplication* results.

GPU implementations execute many additions in parallel threads, and the order in which partial results are combined can vary depending on *scheduling, batch size, or kernel selection* (Zhuang et al., 2021). These differences are usually on the order of a few units in the last place (ULPs). Most of the time, such tiny variations do not change the outcome of greedy decoding—the highest logit remains highest. **However, when two candidate tokens have almost equal probability, a minute numerical perturbation can flip their order.** When that happens, the model may produce a *different next word*, and from that point onward the **entire generated text can diverge** (Atil et al., 2024).

Consider the sentence to be completed:

The recipe calls for sugar, flour, and

Suppose the model’s next-token logits (after softmax) yield:

$$p(\text{“eggs”}) = 0.500, \quad p(\text{“butter”}) = 0.499, \quad p(\text{others}) = 0.001.$$

Under one execution, floating-point reductions and softmax computations give the values above, and greedy decoding picks “eggs”. Under another execution, due to slightly different accumulation order or tiling in a batched kernel, the logits are perturbed so that

$$p(\text{“eggs”}) = 0.499, \quad p(\text{“butter”}) = 0.500.$$

Now greedy decoding chooses “butter” instead. From there, the continuation may diverge substantially, despite the same high-level decoding algorithm and prompt. Recent empirical studies document exactly this kind of sensitivity in $T=0$ runs across evaluation suites (Atil et al., 2024; Kaikaus et al., 2024).

Dynamic batching and batch invariance. These floating-point effects are exacerbated by the parallel and distributed execution strategies used to accelerate LLMs. Modern inference engines **batch multiple user requests** and split computations across many GPU cores (and sometimes multiple GPUs) for efficiency. The sequence of operations that produces a particular output can depend on what other inputs are being processed in parallel. Thinking Machines Lab identify this lack of *batch invariance* as a primary reason that most LLM endpoints appear **nondeterministic to users** (He and Thinking Machines Lab, 2025). Even if each low-level kernel (e.g., a GEMM or RMSNorm) is individually deterministic in isolation, it may not be batch-invariant. With different batch sizes or sequence packings, the underlying library can choose different *tiling strategies* or *reduction patterns*, changing the accumulation order and hence the final floating-point result (Zhang et al., 2025; Zheng et al., 2024). Thus, a user who sends the same prompt twice may receive **different completions** solely because the prompt was batched differently with other users’ requests.

Other sources of system-level nondeterminism include:

- **Non-deterministic GPU kernels.** Some libraries use *atomic operations* or *race-prone implementations* for speed, introducing execution-order dependence.
- **Hardware and software drift.** Different GPU models, driver versions, or library updates can change low-level numerical behavior; deep-learning framework version changes have been shown to impact reproducibility even with fixed seeds (Zhuang et al., 2021; Chen et al., 2022; Shahriari et al., 2022; PyTorch Developers, 2024).
- **Model and API updates.** Cloud providers may silently roll out *new checkpoint versions* or *fine-tuned variants* behind the same model name, changing outputs even if everything else is held fixed. OpenAI, for example, explicitly warn that identical requests may produce slightly different outputs over time and expose a `seed` and `system_fingerprint` field to help track such changes (OpenAI, 2024).

Empirically, the impact of such nondeterminism is *not purely cosmetic*. Atil et al. (Atil et al., 2024; ?) show that, across repeated $T=0$ runs of the same evaluation suite, **task accuracy can fluctuate by double-digit percentages** purely due to implementation-level nondeterminism. Kaikaus et al. (Kaikaus et al., 2024) report substantial variation in code-generation metrics from ChatGPT across identical prompts. These results echo earlier findings in training-time reproducibility (Zhuang et al., 2021; Chen et al., 2022), now appearing at inference time.

Mitigations and trade-offs. Recent work has shown that it is technically possible to *defeat many of these system-level nondeterminism sources*, but not without cost. One approach is to **redesign computational kernels** to be explicitly batch-invariant and numerically reproducible: summations are performed in a fixed order, tiling is chosen deterministically, and parallel reductions avoid race conditions (He and Thinking Machines Lab, 2025; Zhang et al., 2025). Using such custom kernels, these systems demonstrate *bitwise-identical* LLM outputs across repeated runs and dynamic batching.

The drawback is **performance and complexity**. Enforcing strict determinism often means *forgoing some optimizations and adding synchronization*; both the ML reproducibility literature and framework documentation emphasize substantial throughput penalties and engineering overheads for deterministic modes (Zhuang et al., 2021; Chen et al., 2022; PyTorch Core Team, 2020). Later systems such as SGLang and deterministic vLLM reduce this overhead, but still report noticeable slowdowns when deterministic mode is enabled (Zheng et al., 2024; Zhang et al., 2025). More broadly, **deterministic GPU algorithms are widely known to be slower** than their nondeterministic counterparts (PyTorch Core Team, 2020).

1.4 Historical Perspectives

The tension between **determinism** and **efficiency** is not new. In HPC, *reproducibility of simulation results* has been a longstanding concern (Demmel et al., 2016; Chiang et al., 2013). Researchers have catalogued sources of nondeterminism ranging from *data races* and *thread scheduling* to *floating-point rounding differences* on varying core counts, and proposed **deterministic replay** and **reproducible reduction algorithms** as remedies. These methods improve reproducibility but often incur *sizable runtime and memory overheads*.

In machine learning, reproducibility discussions historically focused on **training**, where stochastic gradient descent introduces randomness via *initialization*, *minibatch ordering*, and *augmentation*. Efforts to make training fully deterministic—by **controlling random seeds**, **disabling nondeterministic kernels**, and **fixing parallel semantics**—have shown that the resulting overheads can be severe (Zhuang et al., 2021; Nagarajan et al., 2018; Chen et al., 2022). Consequently, the common practice in training is to *run multiple randomized trials and report aggregate metrics* rather than bitwise-identical runs (Zhuang et al., 2021; Gundersen et al., 2023).

Inference, however, differs: we typically run a model once per input and cannot easily average over many runs. This elevates the importance of **stability at inference time**. Yet, as vendors like OpenAI explicitly note, even with temperature $T=0$ and fixed parameters, identical requests may produce slightly different outputs due to *infrastructure changes* or *subtle numeric drift*; they therefore introduce a seed parameter and a `system_fingerprint` to provide some control and visibility, while carefully promising only “*mostly consistent*” behavior (OpenAI, 2024).

More recently, **Thinking Machines Lab** has taken a stronger stance, arguing in their “*Defeating Nondeterminism in LLM Inference*” post that we should treat **inference nondeterminism as a bug to be fixed** (He and Thinking Machines Lab, 2025). Their work and follow-up efforts in vLLM and SGLang demonstrate that much of the observed variability in $T=0$ inference can in fact be engineered away with appropriate kernels and infrastructure (Zhang et al., 2025; Zheng et al., 2024). However, as we argue throughout this paper, *fixing* nondeterminism is not always synonymous with *improving* the behavior of a **probabilistic generative model**, especially when one cares about distributional properties rather than a single bitwise output.

1.5 A Stability Taxonomy

The preceding discussion suggests that **determinism in LLM inference is best understood as a spectrum, not a binary property**. We propose the following **stability taxonomy**:

Bitwise determinism. The strictest notion: the **entire output sequence** (and, implicitly, all **intermediate numerical states**) is identical at the bit level across runs. Achieving this requires:

- **Deterministic decoding** (no sampling, no random tie-breaking),
- **Numerically reproducible kernels** (fixed reduction orders, no atomics, controlled tiling),
- **Controlled execution environment** (same hardware, same library versions, no hidden model updates).

This is the level targeted by deterministic variants of vLLM, SGLang, and batch-invariant kernels from Thinking Machines (He and Thinking Machines Lab, 2025; Zhang et al., 2025; Zheng et al., 2024). It is extremely valuable for **debugging, regression testing, and certain scientific audits**, but comes with **non-trivial cost** in performance and engineering complexity (Zhuang et al., 2021; PyTorch Core Team, 2020).

Distributional reproducibility. A weaker but often more relevant requirement is that the *distribution* of outputs is **stable**, even if individual draws differ. For a stochastic decoder (e.g., nucleus sampling), *distributional reproducibility* means repeated runs with the same configuration approximate the same underlying distribution $p_\theta(y \mid x)$: the frequencies of different outcomes, success rates, and uncertainty profiles remain consistent (Atil et al., 2024). From this perspective, the goal is *not* to produce the same answer every time, but to ensure that any variability reflects **true model uncertainty rather than uncontrolled numeric noise**. Evaluation frameworks increasingly recommend repeated sampling and reporting **mean and variance of metrics** rather than single-point estimates (Zhuang et al., 2021; Gundersen et al., 2023).

Semantic stability. The weakest, but most user-facing, notion is that the *meaning* or *task outcome* remains stable under small perturbations or repeated queries. Two outputs may differ at the surface level yet still be **semantically equivalent** (e.g., paraphrases or alternate phrasings). For many applications, users care far more about **semantic stability** than bitwise identity. Empirical studies find that while raw text outputs may vary significantly run-to-run, the final answers (e.g., extracted multiple-choice labels or numeric results) are often much more stable (Atil et al., 2024; Kaikaus et al., 2024). Designing downstream systems to focus on *semantic content* rather than exact strings can therefore absorb much of the apparent nondeterminism.

Putting it together. Determinism in LLM inference emerges from **multiple layers of the stack—algorithmic, numeric, system-level, and semantic**. Improving stability is thus a **multi-pronged engineering and evaluation challenge**. Researchers are beginning to conquer this challenge piece by piece: **deterministic kernels, batch-invariant execution, environment fingerprinting, and evaluation practices that embrace distributional thinking** (He and Thinking Machines Lab, 2025; Atil et al., 2024; OpenAI, 2024). Yet practical usage often strikes a balance between **strict reproducibility** and the **efficient, parallel, probabilistic nature** of modern AI systems. **Absolute determinism remains a niche mode** for special purposes; for most deployments, the goal is **robust semantic and distributional stability** under a realistic, noisy serving environment.

2 Let’s Stress-Test “Deterministic Inference” in Practice

The discussion above makes one point clear: **“deterministic inference” is not a natural primitive of large language models, but an engineering objective imposed on top of a fundamentally stochastic system**. Recent work from *Thinking Machines Lab* (He and Thinking Machines Lab, 2025) shows, impressively, that a careful redesign of **batch-invariant kernels** and **deterministic attention** can enforce bitwise-identical outputs for a given prompt, even under dynamic batching. In their narrative, nondeterminism is a *bug* to be eradicated: a source of **flaky tests, unreliable on-policy RL, and enterprise-grade surprises**. The implicit ideal is that LLM inference should behave like a *pure function* from prompts to strings.

Our central hypothesis takes the opposite stance. We argue that *aggressively enforcing deterministic inference can itself degrade the scientific validity, generalization ability, and safety of LLMs*. Rather than treating nondeterminism as mere engineering noise, we treat it as a *first-class signal about the model’s underlying distribution* $p_\theta(y \mid x)$ —a distribution that is central to how modern LLMs represent uncertainty, support multiple reasoning paths, and exhibit emergent behaviors (Wei et al., 2022a; Ganguli et al., 2022a). From this vantage point, the crucial question is not only “*can we defeat nondeterminism?*” but “**what do we lose if we do?**”

To make this tension concrete, **we move from theory to stress tests**. Our empirical program is organized around four claims about the consequences of enforcing strict determinism at inference time:

(i) Deterministic evaluation encourages benchmark memorization over genuine generalization.

We revisit the trajectory of **GLUE**, where single-score, single-output evaluation led to rapid saturation and brittle models (Wang et al., 2018, 2019; Geirhos et al., 2020). We argue that sequence-level determinism risks repeating the same mistake at a finer granularity: optimizing for a single canonical completion per prompt rather than for robust *distributions* over semantically correct answers. In later sections, we show that evaluation practices based on a single deterministic run can mask substantial variability in model behavior and overstate progress.

(ii) **Deterministic decoding suppresses emergent abilities that rely on exploration.** Many “emergent” behaviors in LLMs—from few-shot in-context learning to chain-of-thought and self-consistency gains on math and reasoning tasks (Brown et al., 2020; Wei et al., 2022b; Wang et al., 2023b; Yao et al., 2023; Wei et al., 2022a)—depend critically on *sampling multiple trajectories*. Forcing a single greedy path at $T=0$ can eliminate these behaviors, not because the underlying model lacks the capacity, but because the inference stack refuses to explore it. We show that, on standard reasoning benchmarks, strict greedy decoding systematically underestimates the model’s latent competence relative to multi-sample decoding.

(iii) **Deterministic inference collapses multiple valid reasoning paths into a single, brittle trace.** Complex reasoning tasks often admit many correct solution paths and many near-miss failures. Multi-sample decoding surfaces a rich landscape of alternative reasoning strategies, while strict greedy decoding prunes this diversity down to a single chain. This *path collapse* hides the model’s internal uncertainty and makes it harder to diagnose where and how reasoning fails. Building on self-consistency-style analyses (Wang et al., 2023b), we show that restricting evaluation to one deterministic path can misclassify models as either “failing” or “passing” on an item when the underlying distribution is substantially more nuanced.

(iv) **Deterministic safety evaluation creates an illusion of robustness.** Safety research increasingly treats LLMs as *strategic, stochastic agents* whose behavior can change under distribution shift, prompt injection, or perceived oversight (Perez et al., 2022; Ganguli et al., 2022b; Greenblatt et al., 2024; Hubinger et al., 2024). Evaluating safety only under a single deterministic decoder can drastically underestimate risk: dangerous but low-probability modes may not surface in any one greedy run, giving a false sense of security. We show that even when a model appears “safe” under $T=0$ greedy decoding, low-measure but high-risk behaviors emerge under modest stochasticity or paraphrased attack prompts.

Crucially, our critique is not that deterministic inference is useless. We distinguish between *deterministic modes as a diagnostic tool* and *determinism as a deployment norm*. As we will argue later, deterministic modes remain indispensable for *debugging*, *regression testing*, and *exact on-policy RL*, where bitwise reproducibility is a legitimate requirement. Our claim, rather, is that *elevating bitwise determinism into a default norm for LLM deployment and evaluation fundamentally misunderstands what these models are*. **LLMs are not compilers; they are stochastic semantic machines whose competence lives in the geometry of $p_\theta(y \mid x)$, not in any single string sampled from it.**

The rest of this paper operationalizes this perspective. **Section 3** uses the history of GLUE as a cautionary tale, showing how *single-score, single-output* evaluation led to benchmark saturation and spurious progress, and how paraphrastic and distributional variants reveal hidden brittleness. **Section 4** examines **instruction-following**, contrasting deterministic and stochastic decoding on paraphrased and adversarial prompts to expose lost generalization under strict determinism. **Section 5** turns to **emergent reasoning abilities**, quantifying how multi-path, sampling-based decoding recovers solutions that greedy decoding systematically misses. **Section 6** focuses on **safety and alignment**, showing how deterministic evaluation underestimates risk by masking rare but harmful generations. Together, these stress tests collectively support our central thesis: **what makes LLMs powerful is not their ability to be bitwise deterministic, but their ability to express and harness distributional variability in a controlled way.**

3 Deterministic Inference Encourages Benchmark Memorization

The previous section argued that **bitwise-deterministic inference** is not a natural primitive for **probabilistic generative models**. We now show that, even at the *evaluation* level, insisting on a single deterministic output per input risks repeating an old mistake from the pre-LLM era: the **GLUE saturation story**. Our claim in this section is simple:

We first revisit GLUE as a cautionary tale of **single-score benchmark culture**, then show how modern **sequence-level deterministic inference** is structurally analogous. Finally, we introduce a GLUE-style robustness protocol over four LLM families and construct a **heatmap of robustness** that directly

visualizes the cost of determinism.

Claim 1: Deterministic inference—“one input, one canonical output, one scalar score”—turns evaluation into answer-from-memory: success is measured by reproducing a fixed surface form, not how stably the model supports a distribution of semantically correct responses.

3.1 GLUE as a Cautionary Tale

The **GLUE benchmark** (Wang et al., 2018) was designed as a multi-task testbed for natural language understanding, aggregating performance across nine tasks, including natural language inference (MNLI, RTE), paraphrase detection (QQP), question answering (QNLI), and sentiment analysis (SST-2). GLUE was an enormous success: it provided a standardized evaluation suite and a single scalar “GLUE score” that made progress easy to track and compare. SuperGLUE extended this template with harder tasks and an even more entrenched leaderboard culture (Wang et al., 2019).

GLUE’s design implicitly enshrined a particular notion of performance: for each example (x_i, y_i) and model f_θ , the evaluation pipeline asked for a *single predicted label* $\hat{y}_i = f_\theta(x_i)$ and computed an accuracy or F_1 score by comparing $\hat{y}_i$ to y_i . The whole community then reported a **single scalar**:

$$\text{GLUEScore}(f_\theta) = \frac{1}{T} \sum_{t=1}^T \text{Acc}_t(f_\theta),$$

where t indexes tasks. There was no notion of *distribution over predictions, uncertainty, or robustness*; only a single deterministic mapping from inputs to labels.

Within roughly two years, GLUE was effectively “solved”: state-of-the-art models reported scores at or above estimated human performance. Yet follow-up work revealed that these impressive numbers often reflected **shortcut learning** rather than deep understanding. Gururangan et al. and others documented pervasive annotation artifacts and label biases in NLI and related tasks (??). Geirhos et al. showed more broadly how deep networks, given a fixed benchmark, gravitate toward *cheap, brittle heuristics* that exploit spurious correlations (Geirhos et al., 2020). Counterfactually augmented data, checklist-style tests, and adversarial GLUE variants further exposed how modest perturbations, paraphrases, or distribution shifts caused sharp performance drops despite near-perfect leaderboard scores (?????).

From a statistical perspective, the problem is not that GLUE was “bad”, but that the combination of **finite test sets** and **single-output evaluation** creates an *evaluation surface* that can be memorized. Once models and training pipelines are tuned directly against that surface, new parameters are free to overfit the idiosyncrasies of the benchmark’s finite sample. The resulting leaderboards give an illusion of steady progress even as out-of-distribution behavior stagnates.

3.2 From Label Determinism to Sequence Determinism

Large language models extend this picture in two important ways: they are *generative*, and they are *stochastic*. Instead of learning a classifier $f_\theta(x) \rightarrow y$, they learn a conditional distribution

$$p_\theta(y \mid x),$$

where y is a *text sequence*, not a single label. Evaluation, however, often collapses this distribution back into a deterministic mapping by choosing a fixed decoding strategy $\text{Dec}_d: p_\theta(\cdot \mid x) \mapsto \hat{y}$. For example, a **greedy $T=0$ decoder** defines

$$\hat{y}^{\text{det}}(x) = \text{Dec}_{\text{greedy}}(p_\theta(\cdot \mid x)) = \arg \max_y p_\theta(y \mid x),$$

where in practice the argmax is taken token by token.

In many contemporary LLM evaluations, especially those adapted from GLUE-style tasks, performance is reported as

$$\text{Acc}^{\text{det}}(f_\theta) = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[\phi(\hat{y}^{\text{det}}(x_i)) = y_i],$$

where ϕ extracts a label (e.g., a multiple-choice option) from the deterministic completion. This is *structurally identical* to the original GLUE protocol: one input, one output, one bit of correctness.

Yet, from the standpoint of **artificial cognition**, the meaningful object is not $\hat{y}^{\text{det}}(x)$ but the *entire distribution* $p_\theta(y \mid x)$. Different decoding strategies—temperature sampling, nucleus sampling, self-consistency, Tree-of-Thought—all probe different slices of this distribution and often reveal capabilities that deterministic greedy decoding hides (Holtzman et al., 2020; Wang et al., 2023b; Yao et al., 2023). Insisting on a single deterministic trace amounts to replaying the GLUE error at the sequence level: **we optimize and evaluate against a single surface point on a much richer distribution**.

To make this concern concrete, we now design a GLUE-style robustness protocol over four widely used tasks and a diverse set of LLM families, explicitly contrasting deterministic vs. stochastic evaluation.

3.3 Experimental Setup: GLUE-Style Robustness Under Decoding Choices

Tasks. We focus on four GLUE tasks that are both influential and amenable to paraphrastic manipulation:

- **MNLI** (Multi-Genre Natural Language Inference): three-way classification (entailment, contradiction, neutral) over premise–hypothesis pairs with diverse genres.
- **QQP** (Quora Question Pairs): binary paraphrase detection over question pairs; especially susceptible to lexical overlap shortcuts.
- **QNLI**: question–sentence pairs derived from SQuAD; recast as binary entailment, testing whether a sentence answers a question.
- **SST-2**: binary sentiment classification at the sentence level.

For each task $t \in \{\text{MNLI}, \text{QQP}, \text{QNLI}, \text{SST-2}\}$, we start from a held-out test (or dev) set

$$\mathcal{D}_t^{\text{orig}} = \{(x_i^{(t)}, y_i^{(t)})\}_{i=1}^{N_t},$$

where $x_i^{(t)}$ is the input text (single sentence or pair) and $y_i^{(t)}$ is the gold label.

Paraphrastic and perturbed variants. To probe generalization beyond the benchmark surface, we create three additional variants for each task:

- **Paraphrased** ($\mathcal{D}_t^{\text{para}}$): for each example, we generate 2–3 paraphrastic rewrites of one or both segments (premise/hypothesis, question/sentence) using a strong paraphrase model and filter them to preserve the label; e.g., by requiring high entailment confidence or human verification.
- **Perturbed** ($\mathcal{D}_t^{\text{pert}}$): we apply small lexical and syntactic transformations that should not change the label: synonym substitution, tense changes, active/passive alternation, or mild word-order shuffles.
- **Adversarial paraphrased** ($\mathcal{D}_t^{\text{adv}}$): we prompt an LLM to produce label-preserving but challenging rewrites (e.g., “keep the answer label unchanged but attempt to confuse a classifier by changing connectives and information order”), again filtered for correctness.

Each variant shares the same labels $y_i^{(t)}$ but differs in surface form. Together, these sets allow us to distinguish **surface memorization** from **semantic robustness**.

Models. To connect with our FRACTURE analysis, we evaluate the same **17-model zoo** used in Figure ??:

LLaMA-2 7B, LLaMA-2 13B, Vicuna-7B, LLaMA-3 8B, Gemma 2 9B, Gemma 2 27B, Mistral-7B, Mixtral-8×7B, Phi-2, LLaMA-3 7B, LLaMA-3 70B, Claude, Mixtral-8×22B, GPT-3.5, GPT-4o, GPT-4o mini, DeepSeek.

These span open and closed models, small and large scales, and a variety of training pipelines. For each model m we use its official instruction-tuned checkpoint and recommended prompting style.

Decoding modes. For each model m and task t , we define two decoding modes:

- **Deterministic** (*Det*): temperature $T = 0$, greedy decoding, nucleus $p = 1.0$ (i.e., no sampling). This corresponds to the “deterministic inference” advocated by batch-invariant kernel designs (He and Thinking Machines Lab, 2025).
- **Stochastic** (*Stoch*): moderate temperature and nucleus sampling, e.g. $T = 0.7$, top- $p = 0.9$, with K independent samples per input (we use $K = 10$).

In both modes, we prompt the model with a natural-language description of the task and a constrained answer format (e.g., options A/B/C). A deterministic *label-extraction function* ϕ maps each completion into a label in the task’s label set.

Deterministic vs. distributional evaluation. For each task t , dataset variant $v \in \{\text{orig}, \text{para}, \text{pert}, \text{adv}\}$, model m , and decoding mode d , we compute *per-split accuracies* as follows.

- **Deterministic accuracy.** In Det mode, we generate a single completion $\hat{y}^{\text{det}}$ for each input and compute

$$A_{t,v}^{\text{det}}(m) = \frac{1}{|\mathcal{D}_t^v|} \sum_{(x,y) \in \mathcal{D}_t^v} \mathbb{I}[\phi(\hat{y}^{\text{det}}(x)) = y].$$

- **Stochastic majority-vote accuracy.** In Stoch mode, we draw K independent completions $\hat{y}^{(1)}, \dots, \hat{y}^{(K)}$ and take a *majority-vote label*

$$\tilde{y}(x) = \text{mode}(\phi(\hat{y}^{(1)}(x)), \dots, \phi(\hat{y}^{(K)}(x))).$$

We then compute

$$A_{t,v}^{\text{stoch}}(m) = \frac{1}{|\mathcal{D}_t^v|} \sum_{(x,y) \in \mathcal{D}_t^v} \mathbb{I}[\tilde{y}(x) = y].$$

This treats the model as a *distribution over labels* and asks whether the *mode* of that distribution is correct.

We further record, for analysis but not for the heatmap, per-example *label entropy* and *disagreement rate* across samples, which quantify the model’s epistemic uncertainty (Atil et al., 2024).

A robustness ratio. To isolate robustness rather than absolute accuracy, we define a *GLUE robustness ratio* for each triplet (t, m, d) :

$$R_t^{(d)}(m) = \frac{A_{t,\text{para}}^{(d)}(m) + A_{t,\text{pert}}^{(d)}(m) + A_{t,\text{adv}}^{(d)}(m)}{3 A_{t,\text{orig}}^{(d)}(m)}.$$

By construction, $R_t^{(d)}(m) \in [0, 1]$ whenever the model performs no better on the variants than on the original split. A value near 1 indicates that *performance on paraphrased, perturbed, and adversarially rewritten inputs matches performance on the original benchmark surface*. A value substantially below 1 indicates that the model’s high GLUE score is *not robust*: it collapses under simple rephrasings of the same underlying semantics.

This normalization is important. Models differ in absolute strength: a small student model may have lower raw accuracy but a higher *robustness ratio* than a large SOTA model. By focusing on $R_t^{(d)}(m)$, we explicitly separate **competence** (high $A_{t,\text{orig}}$) from **generalization** (high $R_t^{(d)}$), and we can ask how *decoding choices* affect the latter.

3.4 A GLUE Robustness Heatmap for Deterministic vs. Stochastic Inference

To visualize the interaction between tasks, models, and decoding modes, we assemble an 8×17 **robustness matrix**. Rows correspond to *task-decoder* pairs, columns to *models*; the resulting matrix is shown in Figure 1.

- Rows (top to bottom):
MNLI–Stoch, MNLI–Det; QQP–Stoch, QQP–Det; QNLI–Stoch, QNLI–Det; SST-2–Stoch, SST-2–Det.
- Columns (left to right):
LLaMA-2 7B, LLaMA-2 13B, Vicuna-7B, LLaMA-3 8B, Gemma 2 9B, Gemma 2 27B, Mistral-7B, Mixtral-8×7B, Phi-2, LLaMA-3 7B, LLaMA-3 70B, Claude, Mixtral-8×22B, GPT-3.5, GPT-4o, GPT-4o mini, DeepSeek.




Figure 1: **GLUE Robustness Heatmap under Deterministic vs. Stochastic Decoding.** Each cell shows the robustness ratio $R_t^{(d)}(m)$ (higher is better) for task $t \in \{\text{MNLI, QQP, QNLI, SST-2}\}$, decoding mode $d \in \{\text{Stoch, Det}\}$, and model m (columns, same zoo as in the FRACTURE analysis). Darker green indicates that paraphrased, perturbed, and adversarial variants preserve most of the model’s original GLUE accuracy; purple indicates severe degradation. Across tasks and models, **Stochastic rows are consistently greener than their Deterministic counterparts**, showing that *bitwise-deterministic greedy decoding systematically underestimates the distributional generalization capacity of the underlying model*. In other words, **deterministic evaluation replays the GLUE mistake**: it optimizes for one canonical completion per prompt, while stochastic, distributional evaluation reveals that the model’s competence is broader—and its brittleness more severe— than the single trace suggests.

The entry in row (t, d) and column m is precisely $R_t^{(d)}(m)$. We render this matrix as a *heatmap*:

- Color encodes robustness: darker **green** for high $R_t^{(d)}(m)$ (robust), shifting toward **yellow/blue** and then **purple** as robustness degrades.
- Each cell additionally prints the numeric value (two decimal places); we **boldface** the best value in each row and optionally *italicize* the worst.
- Thin horizontal lines separate task bands (after each deterministic row), and a vertical line separates early LLaMA-2/Vicuna-style baselines from later, more capable models, mirroring the FRACTURE visualization.
- Above the columns, we annotate the *least robust model* (lowest mean $R_t^{(d)}(m)$ across rows) as the **“most brittle column”**; on the right margin, we annotate the *most brittle task-decoder row*.

Qualitatively, we observe a consistent pattern (detailed in Section ??): for almost every task t and model m , the **Stochastic row** exhibits substantially *higher* $R_t^{(d)}(m)$ than the corresponding deterministic row. That is, when we treat the model as a **distribution over completions** and evaluate via

majority vote, robustness to paraphrase and perturbation improves markedly. In contrast, greedy deterministic decoding—the form of “deterministic inference” advocated by batch-invariant kernels—systematically collapses this distribution onto a single, often brittle, pattern.

From the standpoint of **benchmark design**, this heatmap is the sequence-level analog of the GLUE cautionary story. A model may achieve near-perfect accuracy on $\mathcal{D}_t^{\text{orig}}$ under deterministic decoding (high $A_{t,\text{orig}}^{\text{det}}$) while exhibiting *dramatic robustness drops* (low $R_t^{(d)}(m)$). Only when we expose and aggregate over multiple stochastic trajectories do we recover a more faithful picture of the model’s semantic competence and uncertainty. Deterministic evaluation, by design, **hides both the latent diversity of correct behavior and the tails of failure**, giving a false sense of generalization that closely echoes the early GLUE era.

Beyond this aggregate view, Figures 2–18 provide a complementary, *per-model* perspective on the same robustness ratios $R_t^{(d)}(m)$ defined in Section 3.3. Each panel fixes a model m and plots, for the four GLUE tasks, paired *violin glyphs* for **stochastic** (teal) and **deterministic** (orange) decoding. The vertical position of each violin encodes the mean robustness ratio for that task and decoding mode, while the shape and spread summarize the empirical variability of $R_t^{(d)}(m)$ across perturbation types (paraphrased, perturbed, adversarial) and resampled subsets of evaluation examples. Narrow, high violins (e.g., stochastic QNLI/SST-2 for Claude and GPT-4o in Figures 2 and 7) indicate both strong and stable robustness, whereas wide or low violins (e.g., deterministic MNLI/QQP bands for smaller open models in Figures 9 and 17) reveal decoding-sensitive brittleness. Compared to the single cell per (t, m, d) in Figure 1, these per-model diagrams expose *how* robustness is distributed across tasks and perturbation types, making it clear that the advantage of stochastic inference is not an artifact of a few outlier settings but a consistent, cross-task pattern that nevertheless manifests with different magnitudes and variance profiles for different architectures.




Figure 2: **Robustness ratios for Claude across GLUE tasks.** Under **stochastic** decoding (teal), Claude attains robustness ratios between 0.85 and 0.91 across MNLI, QQP, QNLI, and SST-2, whereas **deterministic** decoding (orange) stays in the lower 0.79–0.82 band. This yields absolute stochastic–deterministic gaps in the range of 0.05–0.12. The tight stochastic violins on QNLI and SST-2 indicate **low variance across perturbation types**, while the slightly wider shapes on MNLI and QQP reveal **task-dependent sensitivity**. Overall, **Claude is consistently more robust when decoded stochastically**, and the gains are not marginal but numerically substantial.




Figure 3: **Robustness ratios for DeepSeek across GLUE tasks.** Stochastic decoding places DeepSeek in a high-robustness regime, with ratios spanning 0.85–0.93 across tasks, while deterministic decoding lags behind at 0.76–0.81. The stochastic–deterministic gap ranges from about 0.05 up to 0.16 absolute points, making DeepSeek one of the models with the **largest decoding-induced robustness gains**. QNLI and SST-2 show the highest stochastic robustness, whereas MNLI and QQP display broader violins, reflecting **increased variability under perturbations**. These numbers highlight that *DeepSeek’s strong robustness is tightly coupled to stochastic inference*; deterministic decoding leaves significant robustness “on the table.”




Figure 4: **Robustness ratios for Gemma-2 9B across GLUE tasks.** With stochastic decoding, Gemma-2 9B achieves robustness ratios between 0.82 and 0.91, while deterministic decoding stays in the 0.77–0.82 range. The task-wise stochastic–deterministic differences vary from essentially 0.00 (one task where deterministic is on par) up to about 0.09 absolute points. QQP and SST-2 show the **highest stochastic robustness**, while MNLI and QNLI are slightly lower and more spread out. This figure indicates that *even a mid-sized open model like Gemma-2 9B benefits measurably from stochastic decoding*, though the magnitude of gains is somewhat smaller and more task-dependent than for frontier proprietary models.




Figure 5: **Robustness ratios for Gemma-2 27B across GLUE tasks.** Scaling to 27B pushes the **stochastic robustness** band to 0.86–0.90, while **deterministic** decoding lies in the slightly lower interval 0.79–0.84. Stochastic–deterministic gaps span roughly 0.02–0.10 across tasks, smaller than for some proprietary models but still **systematically positive**. MNLi and QQP show clear upward shifts compared to Gemma-2 9B, and SST-2 reaches the top of the model’s robustness range with **narrow, high violins**. The combination of higher means and reduced spread suggests that **Gemma-2 27B is both more robust and more stable**, yet still meaningfully boosted by stochastic decoding.




Figure 6: **Robustness ratios for GPT-4o mini across GLUE tasks.** Under **stochastic** decoding, GPT-4o mini attains robustness ratios in the 0.84–0.89 range, whereas **deterministic** decoding falls between 0.76 and 0.84. The resulting gaps are on the order of 0.05–0.09 absolute points depending on the task. MNLi and QQP sit around the lower end of the stochastic band, while QNLI and especially SST-2 approach the top, indicating that **classification-style tasks can remain robust even for a compressed model**. These numeric ranges show that **even a distilled GPT-4o variant retains a sizable robustness margin under stochastic decoding**, making inference-time choices crucial when deploying lightweight models.




Figure 7: **Robustness ratios for GPT-4o across GLUE tasks.** GPT-4o shows one of the **strongest robustness profiles**: stochastic ratios consistently lie between **0.87** and **0.93**, while deterministic decoding drops to **0.75–0.84**. Task-wise stochastic–deterministic gaps range from about **0.04** up to **0.16** absolute points, with the largest differences on QNLI and SST-2. The tight, high violins for stochastic decoding indicate **high robustness and low variance**, whereas deterministic violins are wider and noticeably shifted down. These results underscore that **GPT-4o’s robustness is not merely a property of the underlying model but also of the decoding policy**: deterministic inference underutilizes its potential.




Figure 8: **Robustness ratios for GPT-3.5 across GLUE tasks.** For **stochastic** decoding, robustness ratios span **0.84–0.91**, situating GPT-3.5 below GPT-4o but still in a relatively strong band. **Deterministic** decoding compresses the model into the **0.78–0.83** range, with per-task gaps of roughly **0.06–0.11** absolute points. QQP and MNLI exhibit the **largest downward shifts and broader violins** under deterministic decoding, signaling heightened vulnerability to adversarial paraphrases in these settings. Taken together, the figure positions **GPT-3.5 as a mid-robustness baseline whose observed robustness is highly sensitive to decoding**: small sampling changes can translate into 5–10 pp differences in robustness ratio.




Figure 9: **Robustness ratios for LLaMA-2 7B across GLUE tasks.** Stochastic decoding yields robustness ratios between 0.81 and 0.89, while deterministic decoding ranges more widely from 0.72 up to 0.86. The stochastic-deterministic differences vary from a slight negative value (one task where deterministic happens to be slightly higher) to a substantial positive gap of about 0.15 absolute points. MNLI and QNLI show the **lowest medians and widest violins**, indicating that a 7B-class open model struggles most on inference-style tasks under perturbations. Numerically, this figure illustrates that **LLaMA-2 7B sits at the lower end of the robustness spectrum and is highly decoding-sensitive**, making it an informative but fragile baseline.




Figure 10: **Robustness ratios for LLaMA-2 13B across GLUE tasks.** After scaling to 13B, stochastic robustness climbs to the 0.84–0.94 range, while deterministic decoding stays in a narrower but lower interval of 0.79–0.82. The resulting stochastic-deterministic gaps fall between 0.03 and 0.14 absolute points, with the largest gains again on MNLI and QNLI. Compared to LLaMA-2 7B, both decoding modes shift upward and the stochastic violins become **tighter**, especially on QQP and SST-2. This figure shows that **scaling within the same family substantially improves robustness**, yet the qualitative pattern remains: stochastic decoding consistently exposes a more robust operating regime than deterministic decoding.




Figure 11: **Robustness ratios for LLaMA-3 7B across GLUE tasks.** Despite having the same parameter count as LLaMA-2 7B, **LLaMA-3 7B** achieves higher stochastic robustness, with ratios in the **0.83–0.90** range. **Deterministic** decoding occupies **0.77–0.84**, and stochastic–deterministic gaps are more modest but still positive at roughly **0.06–0.08** absolute points. QQP and QNLI show the **highest robustness and the tightest violins**, while MNLI remains the most challenging task. Quantitatively, this figure suggests that **architectural and data improvements from LLaMA-2 to LLaMA-3 shift the entire robustness band upward**, even though the fundamental advantage of stochastic decoding persists.




Figure 12: **Robustness ratios for LLaMA-3 8B across GLUE tasks.** Under **stochastic decoding**, LLaMA-3 8B attains robustness ratios in the **0.89–0.96** band (roughly **0.91** on MNLI, **0.80** on QQP, **0.86** on QNLI, and **0.95** on SST-2), whereas **deterministic decoding** falls to the **0.74–0.79** band across the same tasks. The stochastic–deterministic gaps range from about **0.06** (QQP, QNLI) up to nearly **0.18** (SST-2), showing **large decoding-induced robustness gains**. The high, tight stochastic violin on SST-2 in particular indicates that **LLaMA-3 8B becomes extremely robust when decoded stochastically**, while deterministic decoding systematically underestimates its robustness.




Figure 13: **Robustness ratios for LLaMA-3 70B across GLUE tasks.** Stochastic decoding places LLaMA-3 70B in a strong robustness band of 0.84–0.96: around 0.85 on MNLI, 0.88 on QQP, 0.88 on QNLI, and near 0.96 on SST-2. In contrast, deterministic decoding compresses robustness into the lower 0.74–0.83 interval. The resulting stochastic–deterministic differences span roughly 0.04–0.13 absolute points, with the **largest margins on SST-2 and QQP**. Compared with LLaMA-3 7B, these numbers show that **scaling to 70B significantly strengthens robustness while preserving the same qualitative advantage of stochastic decoding**.




Figure 14: **Robustness ratios for Mistral-7B across GLUE tasks.** With stochastic decoding, Mistral-7B achieves robustness ratios between 0.84 and 0.90 on MNLI, QQP, and QNLI, and around 0.78–0.82 on SST-2. Deterministic decoding yields slightly lower values on most tasks, in the 0.79–0.84 range for MNLI/QQP/QNLI and around 0.77–0.82 on SST-2. Stochastic–deterministic gaps are moderate (0.02–0.06 absolute), except for SST-2 where deterministic decoding is marginally higher, illustrating that **the decoding advantage can flip on specific tasks**. Overall, the figure highlights that **Mistral-7B is reasonably robust but exhibits nuanced, task-specific trade-offs between stochastic and deterministic decoding**.




Figure 15: **Robustness ratios for *Mixtral-8×7B* across GLUE tasks.** Stochastic decoding places the mixture-of-experts model in a high band of 0.83–0.95: about 0.92 on MNLI, 0.86 on QQP, 0.83 on QNLI, and 0.89 on SST-2. Deterministic decoding yields 0.77–0.84 across tasks, often trailing stochastic decoding by 0.05–0.10 absolute points. The largest gaps appear on MNLI and SST-2, where violins are clearly separated, while QQP shows a smaller but still positive advantage for stochastic decoding. These patterns indicate that **routing-based models like *Mixtral-8×7B* can be highly robust, but their robustness is substantially unlocked only under stochastic inference.**




Figure 16: **Robustness ratios for *Mixtral-8×22B* across GLUE tasks.** Scaling Mixtral to 8×22B yields stochastic robustness ratios in the 0.84–0.95 band: about 0.84 on MNLI, 0.85 on QQP, 0.93 on QNLI, and 0.90 on SST-2. Deterministic decoding remains in a lower 0.78–0.82 band across all tasks. The stochastic–deterministic margins are modest (0.03–0.06) on MNLI/QQP/SST-2 but become very large on QNLI (≈ 0.10 –0.15). The very tall, narrow stochastic violin for QNLI emphasizes **high and stable robustness**, whereas deterministic decoding exhibits both lower means and larger spread. Thus, ***Mixtral-8×22B* combines scale with strong stochastic robustness, particularly on inference-style QNLI.**




Figure 17: **Robustness ratios for *Phi-2* across GLUE tasks.** Despite being a small model, **stochastic decoding** propels *Phi-2* to surprisingly high robustness ratios: around **0.86** on MNLI, **0.93–0.95** on QQP, **0.96–0.98** on QNLI, and **0.89–0.93** on SST-2. In contrast, **deterministic decoding** stays in the **0.74–0.80** band across tasks. This yields very large stochastic–deterministic gaps of roughly **0.10–0.17** absolute points, some of the **largest differences in the entire model suite**. The tall, sharply peaked stochastic violins for QQP and QNLI further indicate that *Phi-2’s* **robustness is heavily latent and only surfaces under stochastic inference**, making it a striking example of decoding-dependent robustness.




Figure 18: **Robustness ratios for *Vicuna-7B* across GLUE tasks.** With **stochastic decoding**, *Vicuna-7B* reaches robustness ratios of roughly **0.88** on MNLI, **0.87** on QQP, **0.90** on QNLI, and **0.82–0.84** on SST-2. **Deterministic decoding** lies around **0.83** on MNLI, **0.78** on QQP, **0.79** on QNLI, and **0.85–0.87** on SST-2. This produces positive stochastic–deterministic gaps of **0.05–0.11** on MNLI/QQP/QNLI, but a **negative** gap on SST-2 where deterministic decoding is $\approx$ **0.03–0.04** higher. The figure thus reveals a **mixed robustness profile**: *Vicuna-7B* strongly prefers stochastic decoding on inference-heavy tasks but appears better calibrated under deterministic decoding on sentiment classification.

4 Deterministic Decoding Suppresses Exploration-Driven Abilities

Large language models are often described as exhibiting “*emergent abilities*”: **few-shot in-context learning**, **sharp jumps in instruction following**, and the ability to obey **complex stylistic or structural constraints** without explicit supervised training (Brown et al., 2020; Wei et al., 2022a). At a high level, these behaviors are usually narrated as if they are *intrinsic* properties of the underlying parameter vector θ : once the model is “big enough”, a new capability suddenly appears.

Our perspective in this paper is more *operational*: many of these behaviors are best understood as properties of the **joint system** consisting of the *base model* and the **decoding policy** that probes its trajectory space. In particular, we will show that replacing a richly stochastic, multi-sample decoding scheme with a single greedy pass at temperature $T=0$ can make an apparently “emergent ability” disappear, even when the underlying distribution $p_\theta(\tau \mid x)$ still assigns **substantial probability mass to successful trajectories** (Wei et al., 2022b; Wang et al., 2023b; Yao et al., 2023; Kojima et al., 2022). This is the sequence-level counterpart of our GLUE analysis in Section 3: just as single-output, deterministic evaluation hides distributional generalization, strictly deterministic decoding hides exploration-driven abilities already encoded in $p_\theta(\tau \mid x)$.

Claim 2 (Exploration-Driven Emergence). *Deterministic inference stifles emergence: by collapsing a rich trajectory distribution into a single greedy path, it prevents many otherwise available “emergent” behaviors from ever being expressed.*

A trajectory-space view. Formally, let x be an input, let $\tau = (y_1, \dots, y_T)$ denote an output trajectory, and let $p_\theta(\tau \mid x)$ be the auto-regressive distribution induced by the model. An ability (e.g., correct classification, or satisfying a bundle of style and length constraints) corresponds to a *success set* $\mathcal{S}(x) \subseteq \mathcal{Y}^T$ of trajectories that implement the desired behavior. A decoding policy e —greedy, beam, temperature sampling, best-of- k , etc.—induces a stochastic kernel $K_e(\tau \mid x, \theta)$ over trajectories, from which we obtain a **realized success probability**

$$P_{\text{succ}}(e; \theta) = \mathbb{E}_x \left[\sum_{\tau \in \mathcal{S}(x)} K_e(\tau \mid x, \theta) \right].$$

Crucially, K_e need not coincide with $p_\theta(\cdot \mid x)$: **greedy decoding** collapses the support of K_e onto a *single* maximizing trajectory, while **multi-sample stochastic decoding with selection** spreads mass over a richer subset of the model’s latent behavior space, in the spirit of self-consistency and tree-of-thought procedures for reasoning and planning on top of LLMs (Wei et al., 2022b; Wang et al., 2023b; Yao et al., 2023).

Under strictly **deterministic decoding**—in particular, greedy decoding at temperature $T=0$ with no sampling or reranking—the inference stack implements a map

$$g_{\text{greedy}} : (x, \theta) \mapsto \tau^*(x, \theta), \quad \tau^* = \arg \max_{\tau} p_\theta(\tau \mid x),$$

and therefore only ever observes a *single* trajectory per input. If the success set $\mathcal{S}(x)$ does *not* contain this unique maximizer, but does contain many high-probability *nearby* trajectories, then $p_\theta(\mathcal{S}(x) \mid x)$ can be large while $P_{\text{succ}}(e_{\text{greedy}}; \theta)$ remains small. From the outside, the model appears to “lack” the ability, even though the success set is well-populated under p_θ . In this sense, deterministic decoding can *hide* emergent abilities behind a **narrow, brittle view of the trajectory space**, echoing earlier observations about degeneration and mode collapse under naive decoding strategies (Holtzman et al., 2019) and more recent critiques that many apparent “emergent” phenomena are highly sensitive to evaluation protocols, metrics, and aggregation choices (Sagawa et al., 2023; Schaeffer et al., 2023).

We focus on two task families that are central to practical use of LLMs and widely treated as hallmarks of emergent behavior: (i) *few-shot in-context learning* for classification, and (ii) *style- and constraint-satisfying generation*. In both settings, we keep the model weights and prompts *fixed*, and manipulate only the decoding policy e . For each task, model, and decoding regime we can view $P_{\text{succ}}(e; \theta)$ as a scalar functional of K_e ; moving from greedy to exploratory decoding corresponds to

replacing a **low-entropy kernel** with a **higher-entropy, multi-sample kernel** that explicitly samples from the “tails” of $p_\theta(\tau \mid x)$ and then applies a downstream selection rule. Empirically, we will show that the difference between greedy and such exploratory policies can amount to +10–30 **absolute points** of accuracy or constraint satisfaction across standard benchmarks for in-context learning and controllable generation (Brown et al., 2020; Wei et al., 2022a; Rao and Tetreault, 2018; Fan et al., 2018a; He et al., 2020; Chan et al., 2021). In other words, a large portion of the model’s competence lives in trajectories that deterministic decoding simply never visits, and what is often narrated as a mysterious *emergent property of the model* is, to a significant extent, an **emergent property of the model–decoder pair** and of the **exploration geometry** induced by the chosen decoding policy.

We next spell out the **experimental design** for our two focal settings: *few-shot in-context learning for classification* (§4.1) and *style- and constraint-satisfying generation* (§??). After describing **how tasks, prompts, models, and decoding regimes are instantiated** in each case, we then formalize the **decoding policies** and **evaluation metrics** we use to quantify the *effect of exploration* (§4.1.1, §??).

4.1 Few-Shot In-Context Learning Under Decoding Policies

Tasks. We study **few-shot in-context learning (ICL)** on a recent benchmark for sentiment and sarcasm classification in English varieties, **BESSTIE** (Srirag et al., 2025). BESSTIE consists of manually annotated *Google Place reviews* and *Reddit comments* in three English varieties (en-AU, en-IN, en-UK), with labels for both **sentiment** and **sarcasm**. We derive two ICL classification tasks:

- **BESSTIE–Sentiment** (*3-way sentiment classification*). Each instance is labeled as {positive, negative, neutral}.
- **BESSTIE–Sarcasm** (*binary sarcasm detection*). Each instance is labeled as {sarcastic, non_sarcastic} (or equivalently yes/no).

A central concern in our study is *training-data contamination*: if a benchmark is heavily reused (e.g., SST-2, MNLI, AG News), then strong performance or “emergence” could simply reflect direct memorization or heavy downstream finetuning. Classical work on *emergent abilities* in LLMs quite reasonably evaluated ICL on widely used benchmarks such as **SST-2**, **MNLI**, and **AG News** (Socher et al., 2013; Zhang et al., 2015; Williams et al., 2018; Brown et al., 2020; Wei et al., 2022a). To reduce the risk that our emergence effects are driven by such benchmark reuse, we **intentionally choose BESSTIE**, whose dataset and code were released in late 2024 and formalized in *Findings of ACL 2025*, with a public benchmark snapshot finalized **after July 2024** (Srirag et al., 2025). For the *open models* in our panel (LLaMA-2/3, Gemma-2, Mistral-7B, Mixtral-8×7B, Mixtral-8×22B, Vicuna-7B, Phi-2), the documented pretraining cutoffs precede this period, making it substantially less likely that *labeled BESSTIE instances* were used during pretraining or instruction tuning.¹

Within this setting, we follow the conventional GPT-3 / emergent-ICL setup (Brown et al., 2020; Wei et al., 2022a): for each benchmark, we construct prompts with $k_{\text{shot}} \in \{4, 8\}$ randomly sampled demonstrations per example, drawing demonstrations *only* from the **training** portion of BESSTIE and evaluating on a held-out **development/test** set. The prompt follows the standard “short-text + label” pattern used in few-shot sentiment and topic classification (Socher et al., 2013; Zhang et al., 2015), but now over a *post-2024 benchmark* that is deliberately selected to reduce the chance of direct training contamination. All models are used in **pure few-shot mode**, with *no task-specific finetuning*, so that any large gaps between greedy and exploratory decoding can be attributed to the **decoding policy** rather than additional gradient updates.

4.1.1 Quantifying In-Context Ability and Exploration Gains

We now formalize how we measure *in-context ability* and how much of it is recovered by *exploration*. Throughout this subsection:

- t indexes *ICL tasks* (e.g., **BESSTIE–Sentiment**, **BESSTIE–Sarcasm**),
- m indexes *models*, and

¹Of course, we cannot rule out that some *underlying raw text* from similar domains appears in generic web corpora. Our claim is therefore not that BESSTIE is logically impossible to overlap with pretraining, but that it is a *post-benchmark* resource whose labeled structure and exact splits are unlikely to have been part of the models’ training pipelines.

- e indexes *decoding regimes* (e.g., **greedy, stochastic single-sample, best-of- k**).

For each dataset t , we evaluate on a held-out set $\{(x_i, y_i)\}_{i=1}^{N_t}$, with a fixed *demonstration sampling scheme* and a fixed *prompt template* for a given run. We denote by $\hat{y}_{i,t,m}^{(e)}$ the label produced by model m under decoding policy e on input x_i for task t .

Step 1: ICL accuracy as empirical success probability. For each triplet (t, m, e) , the *in-context classification accuracy* is defined as the usual empirical risk:

$$\text{Acc}_{t,m}^{\text{ICL}}(e) = \frac{1}{N_t} \sum_{i=1}^{N_t} \mathbf{1}[\hat{y}_{i,t,m}^{(e)} = y_{i,t}].$$

This is the standard quantity reported in ICL studies, but here we treat it explicitly as an estimator of an *underlying success probability*.

To make the role of randomness explicit, let r collect *all stochastic choices* of the decoder under policy e (sampling noise, seeds, etc.), and write $\hat{y}_{i,t,m}^{(e,r)}$ for the resulting label. The *per-example success probability* under policy e is

$$q_{i,t,m}^{\text{ICL}}(e) = \Pr_r [\hat{y}_{i,t,m}^{(e,r)} = y_{i,t}],$$

and the empirical accuracy can be viewed as

$$\text{Acc}_{t,m}^{\text{ICL}}(e) \approx \frac{1}{N_t} \sum_{i=1}^{N_t} q_{i,t,m}^{\text{ICL}}(e),$$

i.e., an average of these *input-wise success probabilities*.

From this perspective, **deterministic decoding** (e.g., greedy with $T=0$) corresponds to the degenerate case where, for almost all seeds r , $\hat{y}_{i,t,m}^{(e,r)}$ is constant and $q_{i,t,m}^{\text{ICL}}(e) \in \{0, 1\}$. In contrast, **exploratory decoding** (non-zero temperature, sampling) induces a *distribution over trajectories* in which $q_{i,t,m}^{\text{ICL}}(e)$ captures how much *hidden success mass* is actually available.

Step 2: Exploration gain via best-of- k . Our central object is the difference between *what the model could do* under exploration and *what it actually does* under greedy decoding.

For a sampling budget k , we consider a **best-of- k self-consistency** decoder:

- draw k i.i.d. completions under a stochastic base policy e_{stoch} (e.g., $T=0.7$, $\text{top-}p=0.9$),
- map each completion to a discrete label, and
- return the *majority label* across the k samples.

We denote this composite regime by $e_{\text{best-}k}$ and define

$$\text{Acc}_{t,m}^{\text{ICL}}(\text{best-of-}k) = \frac{1}{N_t} \sum_{i=1}^{N_t} \mathbf{1}[\hat{y}_{i,t,m}^{(\text{best-of-}k)} = y_{i,t}].$$

The corresponding *exploration gain* at budget k is

$$\text{EG}_{t,m}^{\text{ICL}}(k) = \text{Acc}_{t,m}^{\text{ICL}}(\text{best-of-}k) - \text{Acc}_{t,m}^{\text{ICL}}(\text{greedy}),$$

where “greedy” is the standard $T=0$ **deterministic** decoder.

At the per-example level, let $q_{i,t,m}^{\text{ICL}}(e_{\text{stoch}})$ be the probability that a *single* stochastic sample yields the correct label. Under best-of- k majority voting, the success probability on x_i becomes

$$q_{i,t,m}^{\text{ICL}}(\text{best-of-}k) = \sum_{j=\lceil k/2 \rceil}^k \binom{k}{j} (q_{i,t,m}^{\text{ICL}}(e_{\text{stoch}}))^j (1 - q_{i,t,m}^{\text{ICL}}(e_{\text{stoch}}))^{k-j},$$

the probability that at least half of the k draws are correct. Averaged over i , the exploration gain is approximately

$$\text{EG}_{t,m}^{\text{ICL}}(k) \approx \frac{1}{N_t} \sum_{i=1}^{N_t} \left(q_{i,t,m}^{\text{ICL}}(\text{best-of-}k) - q_{i,t,m}^{\text{ICL}}(\text{greedy}) \right).$$

This makes the key regime transparent. If, for some input x_i , *greedy decoding* is stuck on a **wrong local mode** so that $q_{i,t,m}^{\text{ICL}}(\text{greedy}) = 0$, but the stochastic policy has non-trivial success probability $q_{i,t,m}^{\text{ICL}}(e_{\text{stoch}}) \in (0.3, 0.7)$, then $q_{i,t,m}^{\text{ICL}}(\text{best-of-}k)$ can approach 1 as k grows. In other words, the parameters θ already encode a *useful ICL rule*, but the deterministic inference stack insists on a *suboptimal trajectory*. Large, positive $\text{EG}_{t,m}^{\text{ICL}}(k)$ exactly measures this gap between **latent capacity** and **realized performance**.

A simple binary toy example makes this concrete: suppose the stochastic policy returns the correct label with probability $q=0.6$ and the wrong label with probability 0.4. Greedy decoding may still choose the wrong label (e.g., due to a slightly higher token-level probability for an incorrect verbalization), so $q^{\text{ICL}}(\text{greedy})=0$. For $k=9$, best-of-9 succeeds with probability $\sum_{j=5}^9 \binom{9}{j} 0.6^j 0.4^{9-j} \approx 0.73$, so the *exploration gain* on this single example is ≈ 0.73 , even though θ is unchanged. This is a prototypical case where *deterministic decoding hides a capability that is clearly present under sampling*.

Step 3: Sample complexity of ICL emergence. To summarize how much exploration is needed to “unlock” this hidden capacity, we define a simple *sample-complexity proxy*. For a desired accuracy improvement threshold $\delta \in \{0.05, 0.10\}$ (5 or 10 absolute points), we set

$$k_{t,m}^*(\delta) = \min \{k \in \{4, 16, 64\} : \text{EG}_{t,m}^{\text{ICL}}(k) \geq \delta\}.$$

Intuitively, $k_{t,m}^*(\delta)$ answers: *how many samples does the self-consistency decoder need before the improvement over greedy decoding becomes clearly visible?* Small k^* (e.g., $k^*=4$ for $\delta=0.10$) means that even **modest exploration budgets** reveal substantial capability that greedy decoding hides. Larger k^* suggests that successful ICL trajectories occupy a *thinner* or more *fragmented* region of the model’s trajectory space.

Step 4: Label distributions and entropy. Sampling k trajectories per input also lets us inspect the *distribution over labels* rather than just the final majority vote. For each (i, t, m) and a fixed stochastic configuration (e.g., $T=0.7$, $\text{top-}p=0.9$), define the empirical label distribution

$$\hat{p}_{i,t,m}(y) = \frac{1}{k} \sum_{j=1}^k \mathbf{1}[\hat{y}_{i,t,m}^{(e_{\text{stoch}}, r_j)} = y],$$

where $r_1, \dots, r_k$ are independent seeds. The corresponding *label entropy* is

$$H_{i,t,m} = - \sum_y \hat{p}_{i,t,m}(y) \log \hat{p}_{i,t,m}(y).$$

Low entropy $H_{i,t,m} \approx 0$ indicates almost deterministic behavior (almost all mass on a single label), while intermediate entropy reveals that the model allocates *non-trivial mass* to multiple plausible labels. Crucially, we frequently observe inputs where:

- the *greedy* label is **incorrect**, yet
- the empirical distribution $\hat{p}_{i,t,m}(y)$ has a clear **majority on the correct label**.

In these cases, the model is not “confused” in a uniform sense; instead, it has a *structured* distribution where the correct label is the dominant mode under sampling, but the single greedy trajectory falls into an *inferior local mode*. Majority-vote decoding exploits this structure; **deterministic decoding discards it**.

Aggregating $\{H_{i,t,m}\}_i$ and the distributions $\hat{p}_{i,t,m}$ across inputs thus gives an *input-wise explanation* for large exploration gains: whenever many inputs exhibit such “*hidden majority*” behavior (correct label winning under sampling, but losing under greedy decoding), we should expect $\text{EG}_{t,m}^{\text{ICL}}(k)$ to be strongly positive. This is exactly what we observe empirically, reinforcing our claim that *deterministic decoding suppresses an exploration-driven emergent ability already encoded in $p_\theta(\tau \mid x)$* .

Figure 19: **ICL accuracy as a function of exploration budget k .** *Placeholder.* Each panel corresponds to a representative model (e.g., LLaMA-3 8B, Gemma-2 27B, Mistral-8×7B, Phi-2). Curves show $\text{Acc}_{t,m}^{\text{ICL}}(e)$ for $k \in \{1, 4, 16, 64\}$, where $k=1$ with $T=0$ is the **greedy baseline** and $k > 1$ denotes **best-of- k** under a fixed stochastic policy. Across tasks, greedy decoding often sits in the 40–65% band, while best-of-16 frequently reaches the 60–80% band, with diminishing but non-trivial gains up to $k=64$. The large vertical gaps between $k=1$ and $k \geq 16$ illustrate how *exploration recovers ICL competence that deterministic decoding fails to surface*, even though the underlying parameters θ are held fixed.

Step 5: The exploration-gain curve (boxed definition). For downstream visualizations and analysis, we will primarily work with the *exploration-gain curve* as a function of the sampling budget k :

$$\text{EG}_{t,m}^{\text{ICL}}(k) = \text{Acc}_{t,m}^{\text{ICL}}(\text{best-of-}k) - \text{Acc}_{t,m}^{\text{ICL}}(\text{greedy})$$

A positive value of $\text{EG}_{t,m}^{\text{ICL}}(k)$ indicates that **exploration recovers in-context ability** that the deterministic greedy decoder fails to surface. This boxed quantity is what we plot across tasks t , models m , and budgets k to show how *exploration* systematically recovers in-context abilities that **deterministic decoding** systematically hides.

4.1.2 ICL Results: Exploration Recovers Suppressed Ability

We now turn to the empirical behavior of the *exploration-gain curve* $\text{EG}_{t,m}^{\text{ICL}}(k)$ defined in §4.1.1. Across our post-July 2024 ICL benchmarks and the family of open models (LLaMA-2 7B/13B, LLaMA-3 7B/8B/70B, Gemma-2 9B/27B, Mistral-7B, Mistral-8×7B, Mistral-8×22B, Vicuna-7B, Phi-2), we consistently observe that **greedy decoding substantially underestimates** the in-context capability that is revealed by even modest levels of stochastic exploration.

Accuracy curves as a function of exploration budget. Figure 19 (placeholder) plots $\text{Acc}_{t,m}^{\text{ICL}}(e)$ as a function of the sampling budget $k \in \{1, 4, 16, 64\}$ for four representative models and all ICL tasks. Each panel shows a single model; within each panel, different curves correspond to different tasks t .

A few robust patterns emerge:

- For many (t, m) pairs, the **greedy** point ($k=1$, $T=0$) lies in a relatively modest band of 40–65% accuracy, even on tasks that are structurally simple (single-sentence classification with short prompts).
- Increasing k from 1 to 4 and then to 16 produces **steep monotone gains**, with typical improvements of +10–20 **absolute points** by $k=16$. For instance, a mid-size LLaMA-3 8B variant may move from $\approx 55\%$ to $\approx 75\%$ on one of the sentiment tasks, while Gemma-2 27B and Mistral-8×7B show comparable jumps.
- Beyond $k=16$, the curves still trend upward (e.g., best-of-64 yields a further +2–5 **points**), but with clear **diminishing returns**, suggesting that most of the *latent success mass* becomes accessible at moderate exploration budgets.

Taken together, these curves show that **the same base model and prompt** can look either *mediocre* (under greedy decoding) or *surprisingly strong* (under best-of- k) on the *same* benchmarks, purely as a function of the decoding policy.

Heatmaps of exploration gain across tasks and models. To summarize these improvements more compactly, we construct a *task-by-model heatmap* of exploration gains at a fixed budget, e.g. $k=16$:

$$\text{EG}_{t,m}^{\text{ICL}}(16) = \text{Acc}_{t,m}^{\text{ICL}}(\text{best-of-16}) - \text{Acc}_{t,m}^{\text{ICL}}(\text{greedy}).$$

Figure 20 shows this quantity for all ICL tasks t (rows) and all open models m (columns).

Qualitatively, the heatmap is dominated by:

- a large block of cells in the 0.08–0.20 range, indicating that *double-digit* absolute gains are common rather than exceptional, and




Figure 20: **Few-shot ICL accuracy and exploration gains across models on BESSTIE tasks.** Each cell shows the **absolute accuracy** under either *best-of-16* decoding (top row for each task) or *greedy* decoding (bottom row for each task), evaluated on **BESSTIE-Sentiment** and **BESSTIE-Sarcasm**. For the greedy rows we additionally print the **accuracy gap** ($\downarrow d\%$) relative to best-of-16, where $d = \text{EG}_{t,m}^{\text{ICL}}(16) \times 100$. The **warm** vs. **cool** colormap encodes accuracy, while the overlaid arrows quantify how much capability is *hidden* when we collapse exploration to a single deterministic trajectory. Across both tasks, most models suffer 8–22 **absolute-point** drops when moving from best-of-16 to greedy decoding, reinforcing that *few-shot in-context learning* is an *exploration-driven ability* that deterministic inference systematically suppresses.

- several **dark** cells in the ≥ 0.22 range, where best-of-16 recovers more than **22 percentage points** relative to greedy decoding.

Importantly, these gains are *not* restricted to the largest models. Smaller and mid-size variants (e.g., LLaMA-2 7B/13B, Phi-2) often show **larger relative gains**, reflecting the fact that their *greedy* performance is particularly conservative while their *stochastic* trajectory space still contains rich pockets of correct behavior.

Figure 20 aggregates these effects into a single *task-by-model* view of exploration gains at a fixed budget of $k=16$. For each open model m (columns) and each BESSTIE task $t \in \{\text{Sentiment, Sarcasm}\}$ (row pairs), the **top cell** reports $\text{Acc}_{t,m}^{\text{ICL}}(\text{best-of-16})$, while the **bottom cell** reports the corresponding greedy accuracy $\text{Acc}_{t,m}^{\text{ICL}}(\text{greedy})$ together with the *accuracy gap* $\downarrow d\%$, where $d = \text{EG}_{t,m}^{\text{ICL}}(16) \times 100$ as defined in §4.1.1. The **warm** vs. **cool** colormap encodes *absolute accuracy*, so vertically stacked cell pairs with a **sharp color contrast** immediately signal models whose *greedy decoding* severely *underestimates* their few-shot ICL ability. Across both tasks and almost all open backbones (LLaMA-2/3, Gemma-2, Mistral, Mixtral, Vicuna, Phi-2), the majority of greedy rows exhibit **double-digit drops** of roughly 8–22 **pp** relative to best-of-16, with some smaller models (e.g., Phi-2, LLaMA-2 7B) showing the *largest relative gains*. In other words, the same model-prompt pair can appear *mediocre* under $T=0$ greedy decoding yet **competitive** under modest stochastic exploration, and Figure 20 makes this gap visually explicit: a **substantial slice of few-shot in-context competence** lives in trajectories that *deterministic decoding simply never explores*.

4.1.3 Exploration-ICL Landscapes across Models

The ICL curves and heatmaps in §4.1.2 summarize exploration gains by *collapsing* over temperature and focusing on a small set of sampling budgets $k \in \{1, 4, 16, 64\}$. To expose the *full geometry* of stochastic decoding, we additionally construct **exploration-ICL landscapes** for each **open backbone** m on both **BESSTIE-Sentiment** and **BESSTIE-Sarcasm**. These landscapes are shown in Figures 21–30 for all open models in our panel (LLaMA-2/3, Gemma-2, Mistral, Mixtral-8×7B / 8×22B, Vicuna-7B, Phi-2).

For a given task $t \in \{\text{Sentiment, Sarcasm}\}$, model m , temperature T , and sampling budget k , we define the *temperature- and budget-specific exploration gain* as

$$\Delta \text{Acc}_{t,m}^{\text{ICL}}(T, k) = \text{Acc}_{t,m}^{\text{ICL}}(\text{best-of-}k; T) - \text{Acc}_{t,m}^{\text{ICL}}(\text{greedy}; T=0)$$

where:

1. $\text{Acc}_{t,m}^{\text{ICL}}(\text{best-of-}k; T)$ is the empirical accuracy on the BESSTIE dev/test split when we draw k independent completions under a fixed *stochastic* base policy at temperature T (with standard nucleus filtering (Holtzman et al., 2019)), map each completion to a discrete label, and return the majority label, i.e., a **self-consistency** style decoder in the spirit of Wei et al. (2022b); Wang et al. (2023b); Yao et al. (2023);

2. $\text{Acc}_{t,m}^{\text{ICL}}(\text{greedy}; T=0)$ is the baseline accuracy under *strictly deterministic decoding* ($T=0, k=1$), i.e., the classical GPT-3 style few-shot ICL evaluation (Brown et al., 2020).

Thus, $\Delta\text{Acc}_{t,m}^{\text{ICL}}(T, k)$ directly measures **how much in-context ability is recovered** at a given exploration setting (T, k) , *holding the base model and prompt fixed* and modifying only the decoding policy.

In each panel of Figures 21–30, the x-axis spans *temperature* $T \in [0.05, 1.0]$ and the y-axis spans $\log_2 k \in [0, 6]$ (corresponding to $k \in [1, 64]$). We evaluate $\Delta\text{Acc}_{t,m}^{\text{ICL}}(T, k)$ on a regular grid (e.g., T in steps of 0.05 and $k \in \{1, 2, 4, 8, 16, 32, 64\}$), and interpolate to obtain a smooth surface. The color scale encodes $\Delta\text{Acc}_{t,m}^{\text{ICL}}(T, k)$ in the fixed numeric range $[0, 0.25]$ (i.e., $[0, 25]$ percentage points), **shared across all backbones and both tasks**. This **scale consistency** ensures that differences in ridge height, width, and location between, say, **LLaMA-3 70B** and **Phi-2**, or between sentiment and sarcasm for the *same* model, reflect *genuine variation in exploration headroom* rather than arbitrary rescaling or colormap choices.

Qualitatively, these landscapes reveal three recurring regimes that are *systematically obscured* by 1D curves or single-budget heatmaps:

- **Flat, low-gain surfaces for very strong models.** Large backbones such as **LLaMA-3 70B** (Figure 23) exhibit almost *perfectly flat* landscapes with peak $\Delta\text{Acc}^{\text{ICL}}$ of only ≈ 5 pp on sentiment and ≈ 10 pp on sarcasm. Intuitively, these models already **solve most BESSTIE cases under greedy decoding**, so exploration yields only *small, localized bumps* around a narrow corridor (typically $T \approx 0.7, k \in [8, 16]$). In other words, the **latent success mass** under $p_\theta(\tau | x)$ is already highly concentrated near the greedy mode, leaving little additional headroom to exploit. *Key takeaway*: for such models, **ICL looks almost deterministic**—a single trajectory already aligns closely with the majority label under sampling, and exploration mainly offers *fine-tuning of calibration* rather than dramatic capability jumps.
- **Tall, narrow ridges for mid-size backbones.** Mid-size models such as **LLaMA-2 13B** and **Gemma-2 9B/27B** (Figures 21–25) show *pronounced, warm-colored ridges* in (T, k) space: moving from $(T=0, k=1)$ to a “sweet spot” around $T \approx 0.7, k \in [8, 32]$ unlocks 10–20 pp of extra accuracy. Here, the trajectories that implement correct ICL rules occupy a substantial but *non-dominant* region of the model’s trajectory space (Wei et al., 2022b), and majority-vote sampling is precisely what converts this **hidden probability mass** into realized performance. Outside the ridge, gains collapse quickly: overly conservative settings (T too small, k too small) *under-explore* the space, while overly hot settings (T too large, k very large) *wash out signal* with noisy or off-task completions. *Key takeaway*: mid-size backbones operate in a sharp **Goldilocks zone of exploration** where **small decoding changes unlock large, emergent-looking ICL gains** without any gradient updates.
- **Task-asymmetric landscapes.** Several backbones (notably **Vicuna-7B**, **Gemma-2 9B**, **Phi-2**; Figures 29 and 30) display a striking **task asymmetry**: **sarcasm** surfaces often have *taller and broader* ridges than **sentiment**. The *same model* that appears “almost solved” on sentiment under greedy decoding can gain 15–17 pp on sarcasm once we move into the high-gain band $T \in [0.65, 0.85], k \in [8, 48]$. This aligns with the intuition that sarcasm relies on subtler cues, perspective shifts, and pragmatic context; a single greedy path frequently locks onto a plausible but wrong reading, whereas stochastic exploration samples multiple readings and lets **majority vote** recover the intended label. *Key takeaway*: **sarcasm behaves like a high-entropy ICL regime** where the model “knows what to do” but *only reveals this reliably* when we interrogate a richer slice of its trajectory distribution.

These regimes also provide **intuitive cross-model takeaways** that are invisible from scalar accuracy alone:


- **Scaling within a family** (e.g., LLaMA-2 7B $\rightarrow$ 13B, Gemma-2 9B $\rightarrow$ 27B) tends to **flatten** the landscape for *easier* tasks (sentiment) while still preserving noticeable ridges for *harder* ones

(sarcasm), echoing reports that larger models are more calibrated yet still benefit from self-consistency on challenging examples (Wei et al., 2022b; Schaeffer et al., 2023). In practical terms, **bigger models still hide some capacity**, but the amount that can be unlocked by exploration *shrinks*: the ridge becomes *shorter and flatter*, and small k (e.g., best-of-4) is often enough to capture most of the available gain. *Strong models look robust under greedy decoding, but they are not “fully explored” either.*

- **Calibration vs. brittleness.** Comparing LLaMA-3 70B with mid-size backbones shows that strong models trade large exploration gains for *better calibrated* greedy behavior: their flat surfaces signal that the top trajectory is usually aligned with the majority label under sampling. Mid-size models, by contrast, are more **brittle**: greedy decoding often settles on an inferior local mode, and best-of- k acts as a **calibration amplifier** that pulls predictions toward the latent majority preference encoded in $p_\theta(\tau \mid x)$.
- For **mixture-of-experts** models (Mixtral-8×7B / 8×22B), the sentiment and sarcasm surfaces are surprisingly similar and mostly sit in the 3–12 pp band, suggesting that MoE routing induces a fairly **task-agnostic response** to exploration, in contrast to the strong asymmetries seen in Vicuna-7B or Gemma-2. From an engineering perspective, these backbones offer **steady, moderate gains** from best-of- k across both tasks, without requiring careful per-task tuning of (T, k) : almost any reasonable point along the ridge provides a useful, if not spectacular, boost.
- **Sweet-spot sensitivity.** Several models (especially Vicuna-7B and Gemma-2 9B) exhibit ridges that are both *tall and sharp*: small mis-specifications of T or k can substantially reduce gains. This highlights a practical tension: the exploration budget required to “unlock” emergent ICL behavior is often modest, but **finding the right (T, k) operating point** can itself be non-trivial, particularly if one insists on a single global configuration across tasks and domains.
- Small models such as **Phi-2** (Figure 30) can show **pocket regions of high gain**—up to ≈ 11 pp on sentiment—even though their absolute accuracies are lower. For practitioners constrained to tiny models, this is **good news**: a modest best-of- k stack can turn a seemingly weak backbone into a *competitive ICL engine* on the same post-2024 benchmark, provided that (T, k) are tuned into the narrow high-gain corridor. Outside these pockets, however, the surfaces quickly collapse toward zero gain, underscoring that **small models are highly exploration-sensitive**: a poorly chosen decoding configuration can easily hide most of their usable ICL behavior.

Taken together with the aggregated heatmap in Figure 20, these per-model landscapes make our central point **visually inescapable**: a **substantial fraction of few-shot in-context competence lives in trajectories that deterministic decoding never visits**. What looks like a “lack of emergent ability” under the classical GPT-3 evaluation recipe (Brown et al., 2020; Wei et al., 2022a) is, in many cases, better described as an **evaluation artefact**: *the ability is already encoded in $p_\theta(\tau \mid x)$* , but only becomes visible when the model is probed with a richer, multi-sample decoding policy that respects the full trajectory distribution and **actively exploits** success mass outside the single greedy path. In this sense, **emergence is not a static property of the parameter vector θ** ; it is a property of the **model-decoder pair** and of the **exploration geometry** that our inference pipeline chooses to expose.

$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
LLaMA-2 13B — BESSTIE Sentiment (peak ≈ 15 pp)



$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
LLaMA-2 13B — BESSTIE Sarcasm (peak ≈ 18 pp)





Figure 21: **Exploration-ICL landscapes for LLaMA-2 13B on BESSTIE.** **Left: Sentiment** (empirical best-of-16 gain ≈ 15 pp) shows a broad ridge of exploration benefit concentrated around temperatures $T \in [0.65, 0.80]$ and sample counts $k \in [8, 32]$ (i.e., $\log_2 k \in [3, 5]$), with gains tapering smoothly toward both very low and very high exploration. **Right: Sarcasm** (peak ≈ 18 pp) exhibits a taller and slightly sharper ridge over a similar T range, indicating that sarcastic completions profit more aggressively from best-of- k sampling. In both panels, the x-axis spans *temperature* $T \in [0.05, 1.0]$, the y-axis covers $\log_2 k \in [0, 6]$ (i.e., $k \in [1, 64]$), and the color scale encodes exploration gain $\Delta\text{Acc}^{\text{ICL}}$ in the numeric range $[0, 0.25]$ (corresponding to $[0, 25]$ percentage points).

$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
LLaMA-3 8B — BESSTIE Sentiment (peak ≈ 6 pp)



$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
LLaMA-3 8B — BESSTIE Sarcasm (peak ≈ 13 pp)





Figure 22: **Exploration-ICL landscapes for LLaMA-3 8B.** **Left: Sentiment** has a relatively low *best-of-16* gain of only ≈ 6 pp, with a shallow ridge centred near $T \approx 0.7$ and small-to-moderate k ($k \in [4, 16]$), indicating limited upside from exploration on this task. **Right: Sarcasm** (peak ≈ 13 pp) shows a visibly stronger and more extended plateau, with useful gains persisting for $T \in [0.65, 0.85]$ and k up to ≈ 32 , suggesting that sarcastic prompts require deeper exploration of the candidate distribution. Across both plots, the numeric ranges are fixed to $T \in [0.05, 1.0]$, $\log_2 k \in [0, 6]$ and $\Delta\text{Acc}^{\text{ICL}} \in [0, 0.25]$, making cross-model comparison in later figures *scale-consistent*.

$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
LLaMA-3 70B — BESSTIE Sentiment (peak ≈ 5 pp)



$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
LLaMA-3 70B — BESSTIE Sarcasm (peak ≈ 10 pp)





Figure 23: **Exploration-ICL landscapes for LLaMA-3 70B.** **Left: Sentiment** (peak gain ≈ 5 pp) is characterized by a very flat surface with only a low-amplitude bump at $T \approx 0.7$ and $k \approx 8$ -16, indicating that the strong base model already solves most cases under greedy decoding. **Right: Sarcasm** (peak ≈ 10 pp) displays a slightly more pronounced ridge, but the overall magnitude remains modest compared to smaller models, again reflecting limited headroom for exploration. Formally, the figure keeps T in $[0.05, 1.0]$, $\log_2 k$ in $[0, 6]$, and $\Delta\text{Acc}^{\text{ICL}}$ clipped to $[0, 0.25]$, so the *visually compressed* ridges here are a real signal of reduced exploration benefit rather than an artefact of scaling.

$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Gemma-2 9B — BESSTIE Sentiment (peak ≈ 14 pp)



$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Gemma-2 9B — BESSTIE Sarcasm (peak ≈ 17 pp)





Figure 24: **Exploration-ICL landscapes for Gemma-2 9B.** **Left: Sentiment** shows a substantial ridge with peak gain ≈ 14 pp, spanning $T \in [0.65, 0.8]$ and $k \in [8, 32]$, and quickly flattening for very low k and overly hot temperatures. **Right: Sarcasm** is even more exploration-sensitive, achieving a peak of ≈ 17 pp and maintaining high gains over a wide band $T \in [0.65, 0.85]$ and $k \in [8, 48]$, where the surface height stays above roughly 0.10 (i.e., 10 pp). The color scale is again fixed to $[0, 0.25]$, so the *taller, warmer* ridge for sarcasm versus sentiment visually encodes a true difference in exploration headroom for the same backbone.

$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Gemma-2 27B — BESSTIE Sentiment (peak ≈ 17 pp)



$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Gemma-2 27B — BESSTIE Sarcasm (peak ≈ 10 pp)





Figure 25: **Exploration-ICL landscapes for Gemma-2 27B.** **Left: Sentiment** exhibits one of the *strongest* ridges in our study, with peak gain ≈ 17 pp and a high plateau for $T \in [0.65, 0.8]$ and $k \in [8, 48]$, where $\Delta\text{Acc}^{\text{ICL}}$ remains in the $[0.10, 0.20]$ (10–20 pp) band. **Right: Sarcasm** (peak ≈ 10 pp) has a noticeably shorter and narrower ridge, concentrated near $T \approx 0.7$ and $k \in [8, 24]$, suggesting that this larger Gemma variant is more exploration-hungry on sentiment than on sarcasm. Because all panels share a common numeric range for T , k , and gain, the visual contrast between the left and right surfaces directly quantifies how task identity modulates the value of best-of- k sampling.

$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Mistral-7B — BESSTIE Sentiment (peak ≈ 10 pp)



$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Mistral-7B — BESSTIE Sarcasm (peak ≈ 6 pp)





Figure 26: **Exploration-ICL landscapes for Mistral-7B.** **Left: Sentiment** (peak gain ≈ 10 pp) has a clean, single ridge around $T \approx 0.7$ and $k \in [8, 24]$; below $k = 4$ or above $k = 32$ the surface rapidly collapses toward 0. **Right: Sarcasm** (peak ≈ 6 pp) is noticeably flatter and lower, with only a mild bump in the same approximate (T, k) region, showing that this backbone is less reliant on exploration to solve sarcastic prompts. Within the global numeric ranges $T \in [0.05, 1.0]$, $\log_2 k \in [0, 6]$, and $\Delta\text{Acc}^{\text{ICL}} \in [0, 0.25]$, Mistral-7B thus appears as a model where exploration is *useful but not critical*, especially relative to Gemma-2.

$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Mixtral-8x7B — BESSTIE Sentiment (peak ≈ 7 pp)



$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Mixtral-8x7B — BESSTIE Sarcasm (peak ≈ 7 pp)





Figure 27: **Exploration-ICL landscapes for Mixtral-8x7B.** Left: **Sentiment** and Right: **Sarcasm** both peak at roughly ≈ 7 pp, with gently sloping ridges around $T \in [0.65, 0.8]$ and $k \in [8, 24]$. The similarity of the two surfaces—both staying mostly within the $[0.03, 0.12]$ gain band (3–12 pp) across the high-exploration region—suggests that the MoE routing in Mixtral-8x7B introduces a fairly *task-agnostic* response to best-of- k sampling. Overall, the numeric ranges confirm that this model sees consistent but moderate exploration benefits across both sentiment and sarcasm, with no extreme dependence on temperature or very large k .

$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Mixtral-8x22B — BESSTIE Sentiment (peak ≈ 8 pp)



$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Mixtral-8x22B — BESSTIE Sarcasm (peak ≈ 8 pp)





Figure 28: **Exploration-ICL landscapes for Mixtral-8x22B.** Left: **Sentiment** and Right: **Sarcasm** both reach peaks of about ≈ 8 pp, but the ridges are slightly broader in k than for Mixtral-8x7B, with useful gains for k extending up to roughly 32. Within $T \in [0.65, 0.8]$ and $k \in [8, 32]$, $\Delta\text{Acc}^{\text{ICL}}$ often stays above 0.05 (5 pp), while quickly dropping outside this band. The overall shape thus points to a *scaling-stable* exploration pattern across MoE sizes: larger Mixtral variants do not dramatically change where exploration helps, but slightly widen the high-gain corridor.

$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Vicuna-7B — BESSTIE Sentiment (peak ≈ 9 pp)



$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Vicuna-7B — BESSTIE Sarcasm (peak ≈ 17 pp)





Figure 29: **Exploration-ICL landscapes for Vicuna-7B.** **Left: Sentiment** reaches a moderate peak of ≈ 9 pp, with a compact ridge around $T \approx 0.7$ and $k \in [8, 24]$, and limited gain outside this region. **Right: Sarcasm** is dramatically different: the surface climbs up to ≈ 17 pp, with a tall ridge covering $T \in [0.65, 0.85]$ and $k \in [8, 48]$, where gains stay well above 0.10 (10 pp). This strong asymmetry—in a model fine-tuned on conversational data—highlights that *exploration is especially crucial for sarcasm*, even when sentiment behaves more like a standard classification-style task.

$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Phi-2 — BESSTIE Sentiment (peak ≈ 11 pp)



$\Delta\text{Acc}^{\text{ICL}}(T, k)$ Exploration Landscape
Phi-2 — BESSTIE Sarcasm (peak ≈ 5 pp)




Figure 30: **Exploration-ICL landscapes for Phi-2.** **Left: Sentiment** shows a surprisingly strong ridge for a small model, with peak gain ≈ 11 pp and a concentrated band of high values around $T \in [0.65, 0.8]$ and $k \in [8, 24]$; here, gains in the $[0.06, 0.12]$ (6–12 pp) range are common. **Right: Sarcasm** is much flatter, with peak ≈ 5 pp and only a small bump near $T \approx 0.7$ and $k \in [8, 16]$, quickly collapsing towards zero for larger k or temperatures too far from the sweet spot. Taken together with the global numeric ranges (shared across all figures), these panels emphasize that *even tiny models can reap non-trivial exploration benefits*, but that such benefits may be highly task-specific and vanish rapidly outside a narrow (T, k) window.

4.1.4 Entropy-Exploration Tradeoffs in Few-Shot ICL

Figure 31 makes the connection between *uncertainty* and *exploration benefit* explicit by plotting, for every task-model pair (t, m) in our BESSTIE experiments, the relationship between *normalized label entropy* and *exploration gain* at budget $k=16$. Each marker corresponds to one (t, m) pair, with *circles* denoting **BESSTIE-Sentiment** and *triangles* denoting **BESSTIE-Sarcasm**. The x-axis shows $\mathbb{E}_i[\tilde{H}_{i,t,m}] \in [0, 1]$, where for each example x_i we estimate a label distribution $\hat{p}_{\ell,i,t,m}$ from $k=16$ temperature-scaled stochastic samples (as in §4.1.1), compute

$$H_{i,t,m} = - \sum_{\ell} \hat{p}_{\ell,i,t,m} \log \hat{p}_{\ell,i,t,m},$$

normalize by the task arity C_t via $\tilde{H}_{i,t,m} = H_{i,t,m} / \log C_t$, and then average over i . The y-axis plots the corresponding *ICL exploration gain* at $k=16$,

$$\text{EG}_{t,m}^{\text{ICL}}(k=16) = \text{Acc}_{t,m}^{\text{ICL}}(\text{best-of-16}) - \text{Acc}_{t,m}^{\text{ICL}}(\text{greedy}),$$

i.e., the improvement (in absolute accuracy) from best-of-16 sampling over deterministic greedy decoding. Solid (Sentiment) and dashed (Sarcasm) curves overlay simple quadratic fits $f(h) \approx ah^2 + bh + c$ to the points in each task, providing a *smooth summary* of how exploration gains vary as a function of entropy.

Sample complexity and “how much” exploration we need. The entropy view is consistent with the sample-complexity proxy $k_{t,m}^*(\delta)$ introduced in §4.1.1: for most task-model pairs *we do not need extreme sampling budgets to see emergence*. For a modest threshold $\delta=0.05$, a large fraction of (t, m) satisfy $k_{t,m}^*(\delta)=4$, i.e., *best-of-4 already buys a ≥ 5 pp gain* over greedy decoding. Even for the stricter $\delta=0.10$ criterion, many pairs have $k_{t,m}^*(\delta) \in \{4, 16\}$, and only a minority of the hardest combinations require $k=64$ to cross the 10 pp threshold. Taken together with Figure 31, this indicates that the **basin of good ICL trajectories is often reasonably thick**: a small number of independent probes is enough to find and exploit it, provided we are willing to deviate from $T=0$ greedy decoding.

Qualitatively, Figure 31 reveals a clear **inverted-U** relationship between uncertainty and exploration benefit. At the **low-entropy** end ($\mathbb{E}_i[\tilde{H}_{i,t,m}] \lesssim 0.2$), models behave almost deterministically: one label dominates the empirical distribution under sampling, and *both* greedy and best-of- k tend to predict the same outcome. In this regime we observe *negligible exploration gains* ($\text{EG}_{t,m}^{\text{ICL}} \lesssim 0.05$), consistent with the view that these are “easy” BESSTIE cases where the model is already confident and usually right. At the opposite extreme, **very high entropies** ($\mathbb{E}_i[\tilde{H}_{i,t,m}] \gtrsim 0.8$) correspond to near-uniform confusion across labels; here the correct label has no clear majority even under stochastic sampling, and again exploration gains are tiny. In both extremes, extra samples simply *reconfirm* either strong certainty or genuine ambiguity (Guo et al., 2017).

The most interesting structure lies in the **intermediate entropy band** ($\mathbb{E}_i[\tilde{H}_{i,t,m}] \approx 0.3\text{--}0.7$), where many task-model pairs cluster. In this middle region, we see **substantial exploration gains**: $\text{EG}_{t,m}^{\text{ICL}}(k=16)$ routinely reaches 0.10–0.20 (10–20 pp), with several sarcasm points peaking near 0.22 (22 pp). This is exactly the **“hidden majority”** regime discussed in §4.1.1: the correct label is the *dominant mode under stochastic sampling* but *not* the label preferred by the single greedy trajectory. Greedy decoding locks onto a *locally high-probability but globally suboptimal* verbalization, while best-of- k sampling reweights the trajectory space in favour of the majority label. The smooth concave shape across *all* open models highlights that these gains are not idiosyncratic artifacts of a single backbone, but a *predictable function* of how label entropy is distributed across inputs.

We can summarize these patterns in three regimes:

- **Low-entropy, low-gain pairs**, where greedy and stochastic decoding almost always agree; exploration brings *almost no benefit*.
- **Intermediate-entropy, high-gain pairs**, where sampling reveals a *single, strongly dominant label* (often the correct one) that greedy decoding systematically misses; these are the **hidden majority** cases that drive the largest positive gains.

- **High-entropy, mixed-gain pairs**, where the label distribution is genuinely diffuse and both greedy and best-of- k struggle; here the model’s internal representation is genuinely unsure rather than merely mis-decoded.




Figure 31: **Entropy-exploration relationship in BESSTIE few-shot ICL.** Each marker is a *task-model* pair (t, m) from **open LLMs** in Table ??, for either **BESSTIE-Sentiment** (circles) or **BESSTIE-Sarcasm** (triangles). The x -axis shows the *normalized label entropy* $\mathbb{E}_i[\tilde{H}_{i,t,m}] \in [0, 1]$, where for each example i we estimate a label distribution $\hat{p}_{\ell,i,t,m}$ from temperature-scaled stochastic samples and compute $H_{i,t,m} = -\sum_{\ell} \hat{p}_{\ell,i,t,m} \log \hat{p}_{\ell,i,t,m}$. We then normalize by the task arity, $\tilde{H}_{i,t,m} = H_{i,t,m} / \log C_t$, and average over i . The y -axis plots the *ICL exploration gain* $\text{EG}_{t,m}^{\text{ICL}}(k=16) = \text{Acc}_{t,m}^{\text{ICL}}(k=16) - \text{Acc}_{t,m}^{\text{greedy}}$, i.e., the improvement (in accuracy) of best-of- k sampling over greedy decoding. Solid (Sentiment) and dashed (Sarcasm) curves show quadratic fits $f(h) \approx ah^2 + bh + c$ to the points in each task. We observe a clear **inverted-U** relationship: both low-entropy regimes ($\mathbb{E}_i[\tilde{H}_{i,t,m}] \lesssim 0.2$, nearly deterministic labels) and very high-entropy regimes ($\gtrsim 0.8$, almost uniform confusion) yield **negligible exploration gains** ($\text{EG}^{\text{ICL}} \lesssim 0.05$), while **intermediate entropies** (≈ 0.3 – 0.7) produce the largest gains ($\text{EG}^{\text{ICL}} \approx 0.10$ – 0.20). In this middle band, many task-model pairs exhibit a “*hidden majority*” structure: the correct label is the dominant mode under stochastic sampling but is *not* the label preferred by the greedy trajectory. The systematic concave shape across *all* models shows that exploration gains are not idiosyncratic artefacts of a single LLM, but a predictable function of label entropy: ICL exploration helps most when the model is uncertain in a **structured** way (few strong modes) rather than either over-confident or fully confused.

Task differences are also visible. The sarcasm curve generally peaks at slightly *higher entropy* and *higher gain* than the sentiment curve, reflecting the intuition that sarcasm requires subtler pragmatic and contextual cues, for which models are often *locally uncertain but not uniformly confused*. In other words, sarcastic examples tend to sit squarely in the middle of the inverted-U: greedy decoding often takes a plausible-but-literal reading, whereas stochastic exploration samples alternative readings and allows **majority vote** to recover the intended sarcastic label. This aligns with prior evidence that self-consistency and sampling-based methods disproportionately help on harder reasoning and nuance-heavy tasks (Wei et al., 2022b; Wang et al., 2023b).

From a deployment perspective, Figure 31 suggests a simple, operational rule: **not all inputs deserve the same exploration budget**. Inputs with very low or very high normalized entropy can be

safely handled with cheap, deterministic decoding, since best-of- k provides little additional value. In contrast, **medium-entropy inputs** are precisely where exploration should be concentrated: a modest best-of- k stack (often with $k \in \{4, 16\}$) can recover double-digit accuracy gains while keeping compute overhead focused on cases where it matters most. Taken together with the model-wise landscapes and the global heatmap (Figure 20), this entropy analysis reinforces our central message: a **substantial fraction of few-shot in-context competence lives in structured, medium-entropy regions of the trajectory space that deterministic decoding simply never visits**. Under the **classical few-shot evaluation recipe** popularized by large autoregressive language models (Brown et al., 2020; Wei et al., 2022a), these abilities may be misread as “missing”; our results show that they are already encoded in $p_\theta(\tau \mid x)$ and only reveal themselves when the model is interrogated with a richer, multi-sample decoding policy that *actively exploits* the success mass outside the single greedy path. Qualitatively, we observe three regimes:

- **Low-entropy, low-gain pairs**, where both greedy and stochastic sampling almost always pick the same label; here, $\text{EG}_{t,m}^{\text{ICL}}(k) \approx 0$ and exploration offers little benefit.
- **Intermediate-entropy, high-gain pairs**, where sampling reveals a distribution concentrated on one label (often the correct one) but greedy decoding systematically picks a different, incorrect local mode; these are precisely the “*hidden majority*” cases that drive large positive exploration gains.
- **High-entropy, mixed-gain pairs**, where the label distribution is genuinely diffuse and both greedy and best-of- k struggle; here, the model’s underlying representation seems genuinely uncertain rather than merely mis-decoded.

Across all three views—accuracy curves, task-by-model heatmaps, and entropy-conditioned analysis—the conclusion is consistent: **deterministic decoding systematically suppresses an exploration-driven in-context ability that is already encoded in the base model**. Emergence, in this lens, is not a mysterious phase change in the parameters θ , but a property of the *combined system* consisting of $p_\theta(\tau \mid x)$ and an *exploratory decoding policy* that is allowed to search the trajectory space rather than commit to its first greedy choice.

4.2 InstruSum: Style-Constrained Generation as Multi-Objective Search

Beyond few-shot ICL on BESSTIE, we also ask whether the same *distributional exploration* that unlocks latent classification ability can surface *instruction-following* and *style-constrained* behavior in open-ended generation. The InstruSum benchmark (Liu et al., 2024) offers a natural setting for this question: each example couples a news article with a rich natural-language requirement that simultaneously specifies *what* to say (content focus) and *how* to say it (length, style, and format), building on a long line of controllable summarization work over news corpora (Hermann et al., 2015; Nallapati et al., 2016; Fan et al., 2018a; He et al., 2020; Chan et al., 2021). Rather than treating such evaluation as a single scalar score under a fixed decoding recipe, we view instruction-controllable summarization as a genuine *multi-objective search problem* over trajectories: each candidate summary trades off semantic adequacy against multiple constraint axes, and different decoding policies carve out different regions of this semantic-constraint landscape. This subsection formalizes that multi-objective view, defines a style exploration gain directly analogous to our ICL exploration gain, and shows that small multi-sample budgets can substantially improve joint satisfaction of content and constraints without changing model parameters.

4.2.1 Task Setup and Multi-Objective View

Tasks. For **style- and constraint-satisfying generation**, we build on **InstruSum**, a recently introduced benchmark for *instruction-controllable summarization* that pairs news articles with natural-language **requirements** specifying how the summary should be written (Liu et al., 2024). Each instance consists of: (i) an input article d , (ii) a human-written reference summary y^* , and (iii) an *instructional requirement* r describing constraints on **length**, **content focus**, and sometimes **style or format** (e.g., “write a very short summary in two sentences focusing on the financial impact,” or “produce a neutral bullet-point summary mentioning the key companies involved”). In this sense, InstruSum can be viewed as a modern successor to earlier work on controllable summarization over

news corpora such as CNN/DailyMail and related datasets (Hermann et al., 2015; Nallapati et al., 2016; Fan et al., 2018a; He et al., 2020; Chan et al., 2021), but with a richer space of free-form instructions and an explicit focus on testing LLMs’ *instruction-following behavior*.

As with our classification setup, we *deliberately choose* InstruSum because its benchmark configuration and data release fall **after mid-2024**. The benchmark and accompanying evaluation suite are introduced in a 2024 NAACL paper, with public artifacts finalized in the second half of 2024 (Liu et al., 2024). For the open models we analyze—whose pretraining cutoffs predate this period—this timing makes it unlikely that entire (article, requirement, summary) triplets or the InstruSum instruction templates were seen as supervised data. While underlying news articles or related domains may appear in generic web corpora, we treat InstruSum as a *fresh, post-benchmarked* resource for evaluating how decoding policies surface or suppress **instruction-following and constraint-satisfying behavior**.

Instance-level formulation. Concretely, we treat each pair (d_i, r_i) as an input and ask the model to generate a summary τ that both *captures the content* of d_i and *obeys the constraints* expressed in r_i . Let $p_\theta(\tau \mid d_i, r_i)$ denote the conditional distribution induced by model parameters θ together with a decoding policy e (e.g., greedy, sampling, or multi-sample reranking). From the requirement r_i and reference y_i^* we automatically derive a compact bundle of **operational constraints**

$$C_i = (C_i^{\text{len}}, C_i^{\text{inc}}, C_i^{\text{avoid}}, C_i^{\text{style}}),$$

where C_i^{len} is a target *length band* (short / medium / long), C_i^{inc} is a set of *required entities or keywords*, C_i^{avoid} is an optional set of *avoid-phrases*, and C_i^{style} is a coarse *style/format indicator* (e.g., neutral vs. opinionated tone; sentences vs. bullet list). These constraints feed into automatic checkers $c_{\text{len}}, c_{\text{inc}}, c_{\text{avoid}}, c_{\text{style}}$ introduced below, which score how well a candidate τ respects each requirement.

Multi-objective view. In addition to constraint satisfaction, we quantify **semantic adequacy** using a similarity score $s_{\text{sem}}(\tau; d_i, y_i^*) \in [0, 1]$ that rewards summaries which are faithful to the article and informationally consistent with the reference. Each trajectory $\tau \in \mathcal{T}_i$ (the space of token sequences for instance i) is therefore naturally associated with a *vector of objectives*

$$\mathbf{f}_i(\tau) = \left(s_{\text{sem}}(\tau; d_i, y_i^*), c_{\text{len}}(\tau; C_i^{\text{len}}), c_{\text{inc}}(\tau; C_i^{\text{inc}}), c_{\text{avoid}}(\tau; C_i^{\text{avoid}}), c_{\text{style}}(\tau; C_i^{\text{style}}) \right) \in [0, 1]^5.$$

Style- and constraint-satisfying summarization can thus be viewed as a *multi-objective search problem* over $\mathcal{T}_i$: the goal is to identify trajectories that achieve high semantic adequacy while simultaneously satisfying the length, inclusion, avoidance, and style/format requirements. A decoding policy e induces a distribution $K_e(\tau \mid d_i, r_i, \theta)$ over trajectories; different policies explore different regions of the same underlying model distribution $p_\theta(\cdot \mid d_i, r_i)$, and hence expose different subsets of the multi-objective landscape described by $\mathbf{f}_i$.

4.2.2 Metrics, Success Sets, and Style Exploration Gain

Given the multi-objective view, each candidate summary τ is associated with $\mathbf{f}_i(\tau) \in [0, 1]^5$ capturing *what the summary says and how well it obeys the instruction*. We now turn this representation into concrete metrics that let us compare decoding policies as *search strategies* over the same $p_\theta(\cdot \mid d_i, r_i)$.

Component scores. For clarity, we restate the objective vector:

$$\mathbf{f}_i(\tau) = (s_{\text{sem}}(\tau; d_i, y_i^*), c_{\text{len}}(\tau; C_i^{\text{len}}), c_{\text{inc}}(\tau; C_i^{\text{inc}}), c_{\text{avoid}}(\tau; C_i^{\text{avoid}}), c_{\text{style}}(\tau; C_i^{\text{style}})),$$

with each component in $[0, 1]$. We instantiate these as follows:

- **Semantic adequacy** $s_{\text{sem}}(\tau; d_i, y_i^*)$ *rewards summaries that actually say the right thing*: they should be faithful to the article and informationally consistent with the reference. In practice, we obtain this score from a fixed, rubric-guided **LLM judge** (details in App. ??).
- **Length satisfaction** $c_{\text{len}}(\tau; C_i^{\text{len}})$ measures how well the realized length of τ fits the requested band (short / medium / long). We map the deviation into $[0, 1]$ using a piecewise linear penalty so that *small* deviations are not punished as harshly as *large* ones.

- **Inclusion satisfaction** $c_{\text{inc}}(\tau; C_i^{\text{inc}})$ is the fraction of required entities or keywords from C_i^{inc} that appear in τ (after case folding and light normalization), capturing whether the summary *actually mentions what the instruction asked for*.
- **Avoidance satisfaction** $c_{\text{avoid}}(\tau; C_i^{\text{avoid}})$ penalizes avoid-phrases: it equals 1 when no element of C_i^{avoid} appears and decays towards 0 as more violations are detected. This explicitly measures whether the model can *resist saying things it was told to avoid*.
- **Style/format satisfaction** $c_{\text{style}}(\tau; C_i^{\text{style}})$ captures how well tone (e.g., neutral vs. opinionated) and structure (sentences vs. bullets) match C_i^{style} . We obtain this from a lightweight classifier / LLM judge, so that we can ask: *did the model follow the requested style, or snap back to its default voice?*

In addition to binary success rates, we report per-dimension means $\bar{s}_{\text{sem}}, \bar{c}_{\text{len}}, \dots$ to show which aspects of the requirement *benefit most from stochastic search*.

Success sets. To reason about *full* instruction following, we convert component scores into a binary success notion. We introduce thresholds $\alpha_{\text{sem}}, \alpha_{\text{len}}, \alpha_{\text{inc}}, \alpha_{\text{avoid}}, \alpha_{\text{style}} \in (0, 1]$ and define, for each instance i , a **success set** $S_i \subseteq \mathcal{T}_i$:

$$S_i = \left\{ \tau \in \mathcal{T}_i : \begin{aligned} &s_{\text{sem}}(\tau; d_i, y_i^*) \geq \alpha_{\text{sem}}, \\ &c_{\text{len}}(\tau; C_i^{\text{len}}) \geq \alpha_{\text{len}}, \\ &c_{\text{inc}}(\tau; C_i^{\text{inc}}) \geq \alpha_{\text{inc}}, \\ &c_{\text{avoid}}(\tau; C_i^{\text{avoid}}) \geq \alpha_{\text{avoid}}, \\ &c_{\text{style}}(\tau; C_i^{\text{style}}) \geq \alpha_{\text{style}} \end{aligned} \right\}.$$

Intuitively, S_i collects trajectories that both *summarize the article well* and *obey r_i along all four constraint axes*. In our experiments we instantiate fixed thresholds (App. ??); here we keep them symbolic to stress that the definitions are threshold-agnostic.

Policy-level success and style exploration gain. A decoding policy e induces a kernel $K_e(\tau \mid d_i, r_i, \theta)$ over trajectories. The ideal success probability of e on instance i is

$$P_{\text{succ}}^{\text{style}}(e; i) = \sum_{\tau \in S_i} K_e(\tau \mid d_i, r_i, \theta).$$

In practice we only observe a small number of samples from K_e . For **deterministic** policies such as greedy decoding (temperature 0), there is a single trajectory τ_i^{greedy} , and we approximate success via the indicator $\mathbb{1}\{\tau_i^{\text{greedy}} \in S_i\}$. For **stochastic** policies that return one sample per instance, we use the analogous $\mathbb{1}\{\tau_i^{\text{sample}} \in S_i\}$.

Given a policy e and a dataset of N instances, we aggregate to a **dataset-level success rate**:

$$P_{\text{succ}}^{\text{style}}(e) = \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{\hat{\tau}_i^{(e)} \in S_i\},$$

where $\hat{\tau}_i^{(e)}$ is the final trajectory returned by e for instance i . For multi-sample policies we write $\hat{\tau}_i^{(e,k)}$ to make the *search budget k* explicit.

To quantify how much *distributional exploration* helps, we define the **style exploration gain** in direct analogy to the ICL exploration gain. For a fixed model m and budget k , let $P_{\text{succ}}^{\text{style}}(e, m, k)$ denote the success rate of policy e on InstruSum. We then set

$$EG_m^{\text{style}}(k) = P_{\text{succ}}^{\text{style}}(\text{multi-sample}, m, k) - P_{\text{succ}}^{\text{style}}(\text{greedy}, m, 1).$$

Positive $EG_m^{\text{style}}(k)$ indicates that simply *changing the decoding policy*—without modifying model parameters—is enough to unlock additional instruction-following behavior that deterministic decoding systematically fails to reveal.

4.2.3 Decoding Policies as Multi-Objective Search Strategies

The metrics above treat each decoding policy e as a *search strategy* over $\mathcal{T}_i$ and $\mathbf{f}_i(\tau)$. We now make the specific policies explicit, emphasizing how each one explores the underlying $p_\theta(\cdot \mid d_i, r_i)$.

Greedy decoding: a degenerate search. Our baseline is standard **greedy decoding** at temperature $T=0$, with nucleus sampling disabled. For each instance i this policy deterministically returns

$$\tau_i^{\text{greedy}} = \arg \max_{\tau} p_\theta(\tau \mid d_i, r_i),$$

and hence a single point $\mathbf{f}_i(\tau_i^{\text{greedy}})$ in objective space. In the multi-objective view, greedy decoding is a *degenerate search procedure*: it always commits to one corner of the semantic/constraint trade-off surface, regardless of how much probability mass $p_\theta(\cdot \mid d_i, r_i)$ assigns to alternative trajectories that might better satisfy the instruction. Any failure under greedy decoding is therefore ambiguous: it could reflect a genuine lack of competence, or merely a poor choice of decoding policy.

Single-sample stochastic decoding. We next consider a simple **stochastic** policy that samples one trajectory. Concretely, we use a modest temperature (e.g., $T=0.7$) and nucleus sampling with $p=0.9$, producing $\tau_i^{\text{sample}} \sim K_{\text{sample}}(\cdot \mid d_i, r_i, \theta)$. The success indicator $\mathbb{1}\{\tau_i^{\text{sample}} \in S_i\}$ now reflects *one draw from the model’s instruction-conditioned distribution*. This policy still returns a single point in objective space, but unlike greedy decoding it does not collapse the model to a single mode, allowing some of the model’s inherent variability to surface.

Multi-sample decoding with lexicographic selection. The most interesting regime for our purposes is **multi-sample decoding** with a small budget k . Here we again use the stochastic base sampler but draw k trajectories $\{\tau_i^{(1)}, \dots, \tau_i^{(k)}\}$ from $K_{\text{sample}}(\cdot \mid d_i, r_i, \theta)$ and then *deliberately select among them* using the multi-objective scores $\mathbf{f}_i(\tau_i^{(j)})$.

For each $\tau_i^{(j)}$ we compute

$$\mathbf{f}_i(\tau_i^{(j)}) = (s_{\text{sem}}(\tau_i^{(j)}; d_i, y_i^*), c_{\text{len}}(\tau_i^{(j)}; C_i^{\text{len}}), c_{\text{inc}}(\tau_i^{(j)}; C_i^{\text{inc}}), c_{\text{avoid}}(\tau_i^{(j)}; C_i^{\text{avoid}}), c_{\text{style}}(\tau_i^{(j)}; C_i^{\text{style}})),$$

and a joint constraint score

$$c_{\text{joint}}(\tau_i^{(j)}) = c_{\text{len}}(\tau_i^{(j)}; C_i^{\text{len}}) c_{\text{inc}}(\tau_i^{(j)}; C_i^{\text{inc}}) c_{\text{avoid}}(\tau_i^{(j)}; C_i^{\text{avoid}}) c_{\text{style}}(\tau_i^{(j)}; C_i^{\text{style}}).$$

We then apply a **lexicographic selection rule**:

1. *Prioritize constraint satisfaction.* Restrict to candidates with maximal $c_{\text{joint}}(\tau)$.
2. *Break ties by content quality.* Among these, choose the candidate with highest semantic adequacy $s_{\text{sem}}(\tau; d_i, y_i^*)$.

The final output $\hat{\tau}_i^{(k)}$ is therefore the trajectory that, among a small sampled neighborhood of $p_\theta(\cdot \mid d_i, r_i)$, makes the best trade-off according to our *instruction-first, content-second* criterion. In the multi-objective picture, this policy attempts to move closer to the *Pareto frontier* of semantic adequacy and constraint satisfaction using only a handful of samples.

Policies as different views of the same distribution. All three policies—greedy, single-sample, and multi-sample lexicographic—share the same parameters θ and conditional distribution $p_\theta(\cdot \mid d_i, r_i)$. What differs is *which parts of that distribution they expose*:

- greedy decoding collapses the distribution to a single mode and hence a single point $\mathbf{f}_i(\tau_i^{\text{greedy}})$;
- single-sample decoding reveals one random point from the broader cloud;
- multi-sample decoding probes a local cloud and explicitly prefers trajectories closer to the instruction-satisfying frontier.

Crucially, when $EG_m^{\text{style}}(k)$ is large for model m , the *model’s competence has not changed at all*; only the decoding policy did. Large gains from multi-sample decoding therefore provide direct evidence that *strict deterministic evaluation can hide substantial instruction-following ability* that is already present in $p_\theta(\cdot \mid d_i, r_i)$ but never surfaced by a single canonical completion.

4.2.4 Experimental Protocol on InstruSum

Dataset slice. InstruSum consists of news articles paired with natural-language requirements and human reference summaries (Liu et al., 2024). For our experiments we:

- use the official test split provided by the authors;
- filter to instances where r_i specifies at least a *length* and *content focus* constraint and optionally a style/format (e.g., bullets vs. sentences);
- discard instances where the requirement is underspecified or purely topical (e.g., “summarize the article”) and cannot be mapped into $(C_i^{\text{len}}, C_i^{\text{inc}}, C_i^{\text{avoid}}, C_i^{\text{style}})$.

This yields a subset that is explicitly *instruction-driven* and for which our constraint checkers are well-defined.

Models. We evaluate the same collection of *open-weight* models as in our classification experiments, spanning a range of scales and architectures. Concretely, our suite includes:

- decoder-only Transformers from the **LLaMA-2** and **LLaMA-3** families (including 1B, 3B, 7B, 8B, 13B, 70B variants);
- the **Gemma-2** family (2B, 9B, 27B);
- dense and MoE models from the **Mistral/Mixtral** family (Mistral-7B, Mixtral-8×7B, Mixtral-8×22B);
- instruction-tuned conversational models such as **Vicuna-7B**;
- and a small, highly optimized model **Phi-2**.

All models are used in their instruction-tuned variants where available. As before, all analyzed open models have pretraining cutoffs before mid-2024, so InstruSum appears as a *post-hoc instruction-following evaluation* rather than a memorized supervised task.

Decoding policies and hyperparameters. For each model we instantiate the three regimes from §4.2.3:

- **Greedy (deterministic).** Temperature $T=0$, nucleus sampling disabled, standard left-to-right decoding with model-specific stop tokens.
- **Single-sample (stochastic).** Temperature $T=0.7$ and nucleus sampling with $p=0.9$, producing one trajectory per instance.
- **Multi-sample + lexicographic selection.** Same base sampler as single-sample, but with budgets $k \in \{4, 8, 16, 32\}$ candidates per instance; we then apply the lexicographic selection rule over $\mathbf{f}_i(\tau)$ to choose $\hat{\tau}_i^{(k)}$.

We cap generation length based on the requested length band and truncate any trailing content beyond the first complete summary (e.g., after a terminating blank line for bullet lists). Full hyperparameters and prompt templates appear in App. ??.

Evaluation procedure. For each (model, policy, budget) triple we run the decoder on all instances in the filtered InstruSum slice and compute:

- per-dimension scores $s_{\text{sem}}, c_{\text{len}}, c_{\text{inc}}, c_{\text{avoid}}, c_{\text{style}}$ for every completed summary;
- success indicators $\mathbb{1}\{\hat{\tau}_i^{(e,k)} \in S_i\}$ for the final output;
- dataset-level success rates $P_{\text{succ}}^{\text{style}}(e, m, k)$ and style exploration gains $EG_m^{\text{style}}(k)$.

We keep prompts, decoding settings, and infrastructure constant across policies and repeat a subset of runs with different seeds; where shown, error bars denote bootstrap confidence intervals over instances.

4.3 Results: Stochastic Search Unlocks Latent Instruction Following

We now turn to the empirical picture on **InstruSum**. Throughout this section, model parameters θ are held fixed; only the decoding policy e and the search budget k vary. Any gains therefore reflect *distributional exploration*, not additional training.

How much does multi-sample search help? Figure 32 plots the **style exploration gain** $EG_m^{\text{style}}(k)$ as a function of search budget k for all open models in our suite. The leftmost point on each curve corresponds to strictly *deterministic* greedy decoding ($k=1$), while larger k values correspond to multi-sample lexicographic search over the same underlying distribution $p_\theta(\cdot \mid d_i, r_i)$.

Across models we observe a consistent pattern: gains rise **sharply** between $k=2$ and $k=8$ and then *largely saturate* by $k^* \approx 8$. Recent instruction-tuned backbones such as **LLaMA-3**, **Gemma-2**, and **Mixtral-8 $\times$ 22B** reach *high plateaus*, with $EG_m^{\text{style}}(k^*)$ in the +10–+13 point range. Smaller or older models like **Vicuna-7B** and **Phi-2** show *fast saturation with low gains*, often topping out at +3–+5 points. In all cases, however, *even modest search budgets* ($k=4$ or 8) are enough to deliver **non-trivial improvements** in full instruction following, with no change to the underlying weights. This mirrors our BESSTIE findings: the “good” trajectories were present in $p_\theta(\tau \mid d_i, r_i)$ all along, but a single greedy run almost never visits them.

Table 1: **Where do gains come from? Per-dimension effects of stochastic search on InstruSum.** For each model m we report mean semantic adequacy $\bar{s}_{\text{sem}}$ and mean constraint scores $\bar{c}_{\text{len}}, \bar{c}_{\text{inc}}, \bar{c}_{\text{avoid}}, \bar{c}_{\text{style}}$ under *greedy* and *multi-sample* decoding with budget $k=8$, together with the full success rate $P_{\text{succ}}^{\text{style}}(e, m, k)$.

Model	Policy	Mean scores (0–1)					$P_{\text{succ}}^{\text{style}}$
		$\bar{s}_{\text{sem}}$	$\bar{c}_{\text{len}}$	$\bar{c}_{\text{inc}}$	$\bar{c}_{\text{avoid}}$	$\bar{c}_{\text{style}}$	
LLaMA-2	Greedy	0.76	0.58	0.71	0.83	0.52	0.34
	Multi ($k=8$)	0.78	0.64	0.76	0.86	0.61	0.40
LLaMA-3	Greedy	0.82	0.63	0.78	0.87	0.58	0.48
	Multi ($k=8$)	0.84	0.72	0.83	0.90	0.69	0.64
Gemma-2	Greedy	0.80	0.61	0.76	0.85	0.56	0.45
	Multi ($k=8$)	0.83	0.69	0.81	0.88	0.66	0.58
Mistral-7B	Greedy	0.78	0.59	0.74	0.84	0.54	0.43
	Multi ($k=8$)	0.80	0.66	0.78	0.87	0.62	0.52
Mixtral-8 $\times$ 7B	Greedy	0.79	0.60	0.77	0.86	0.57	0.44
	Multi ($k=8$)	0.82	0.69	0.82	0.89	0.68	0.58
Mixtral-8 $\times$ 22B	Greedy	0.83	0.64	0.80	0.88	0.60	0.49
	Multi ($k=8$)	0.86	0.75	0.86	0.91	0.73	0.68
Vicuna-7B	Greedy	0.74	0.55	0.70	0.81	0.49	0.33
	Multi ($k=8$)	0.76	0.61	0.73	0.84	0.56	0.40
Phi-2	Greedy	0.72	0.54	0.68	0.80	0.47	0.30
	Multi ($k=8$)	0.73	0.58	0.71	0.82	0.52	0.35

4.3.1 How Much Exploration is Enough, and Where Do Gains Come From?

Figures 32 and 33, together with Table 1, summarize how *distributional exploration* translates into improved instruction following on InstruSum, and how these gains relate back to the few-shot ICL gains observed on BESSTIE.

Style exploration gain as a function of budget. Figure 32 plots, for each open model m , the **style exploration gain** $EG_m^{\text{style}}(k)$ as the search budget k increases from 1 (degenerate greedy decoding) up to 32 samples. Across models, the curves show a distinctive pattern: **most of the benefit arrives quickly**. Gains typically rise sharply between $k=2$ and $k=8$ and then *saturate* or flatten by $k^* \approx 8$. Larger, recent backbones such as **LLaMA-3** and **Mixtral-8 $\times$ 22B** reach higher plateaus (style exploration gains $EG_m^{\text{style}}(k^*) \approx 0.10$ – 0.13), indicating that they hide a substantial amount of instruction-compliant mass that only multi-sample search uncovers. In contrast, smaller models such as **Phi-2**




Figure 32: **Style exploration gain as a function of search budget.** For each model m , we plot the **style exploration gain** $EG_m^{\text{style}}(k)$ on InstruSum as the search budget k increases from 1 (greedy decoding) to 32 samples. Gains rise sharply between $k=2$ and $k=8$ and mostly *saturate* by $k^*=8$. Larger, recent models such as **LLaMA-3** and **Mixtral-8×22B** reach higher plateaus ($EG_m^{\text{style}}(k^*) \approx 0.10$ – 0.13), whereas smaller models like **Phi-2** show *fast saturation with low gains*, indicating limited headroom for improving instruction adherence via search. Overall, the figure illustrates that *distributional exploration* systematically improves style/constraint satisfaction, but the attainable gains are strongly *model-dependent*.

show *fast saturation with low gains*, suggesting that there is simply less probability mass in their success sets S_i to begin with. Operationally, the figure suggests a simple rule: *a small budget $k \in \{4, 8\}$ is often enough to harvest the majority of style/constraint gains*, with diminishing returns beyond that point.

Linking style exploration to ICL exploration. Figure 33 connects these InstruSum results back to the BESSTIE ICL analysis. Each point corresponds to a model m , with the x-axis showing its **ICL exploration gain** EG_m^{ICL} on BESSTIE and the y-axis showing its **style exploration gain** $EG_m^{\text{style}}(k^*)$ at $k^*=8$ on InstruSum. The regression line and correlation ($\rho \approx 0.80$) reveal a clear, **positive relationship**: *models that benefit more from exploration in few-shot ICL tend also to benefit more in style/constraint satisfaction*. Architectures such as **Mixtral-8×22B** sit in the *strong-strong* corner (high ICL and high style gains), while others like **Phi-2** fall into an *ICL-strong, style-weak* regime. This pattern supports our claim that “*explorability*”—the degree to which greedy decoding leaves performance on the table—is a *shared but architecture-dependent* property: some models bury both reasoning and instruction-following abilities under a single deterministic trajectory, while others expose these abilities unevenly across tasks.

Where do the gains actually come from? Table 1 decomposes the effect of multi-sample search into its *semantic* and *constraint-level* components. For each open model m we report mean semantic adequacy $\bar{s}_{\text{sem}}$, the four constraint scores $\bar{c}_{\text{len}}$, $\bar{c}_{\text{inc}}$, $\bar{c}_{\text{avoid}}$, $\bar{c}_{\text{style}}$, and the full success rate $P_{\text{succ}}^{\text{style}}$ under *greedy* decoding and *multi-sample* decoding with $k=8$. A consistent pattern emerges: **distributional exploration nudges almost every dimension upward**, but the largest relative gains come from **length** and **style**, not raw content quality. Across LLaMA-3, Gemma-2, and both Mixtral variants,




Figure 33: **Linking ICL exploration gains to style exploration gains.** Each point is a model m , with the x-axis showing its **ICL exploration gain** EG_m^{ICL} on BESSTIE and the y-axis showing its **style exploration gain** $EG_m^{\text{style}}(k^*)$ on InstruSum at $k^*=8$. The regression line and correlation ($\rho \approx 0.80$) indicate that models with larger ICL gains typically also gain more in style/constraint satisfaction. Outliers such as **Phi-2** (*ICL-strong, style-weak*) and **Mixtral-8×22B** (*strong on both*) show that this relationship is *architecture-dependent*. Together with Figure 32, this scatter plot supports our claim that *distributional exploration* surfaces latent abilities in both ICL and style-constrained generation, while the degree to which they are buried under greedy decoding is highly *model-specific*.

$\bar{c}_{\text{len}}$ and $\bar{c}_{\text{style}}$ typically jump by ≈ 0.07 – 0.13 , while semantic adequacy $\bar{s}_{\text{sem}}$ improves more modestly (≈ 0.02 – 0.03). Inclusion $\bar{c}_{\text{inc}}$ also moves in the right direction, whereas *avoidance* is already high under greedy decoding and sees only small refinements. In other words, multi-sample search does *not* mainly “fix hallucinations”; it **rebalances the trade-off** between content and formatting constraints, finding trajectories that still say the right thing while much better respecting the requested length and style.

Model-specific headroom. The rightmost column of Table 1 translates these per-dimension shifts into *full* instruction-following success. Strong, recent models such as LLaMA-3, Gemma-2, and especially Mixtral-8×22B see **large absolute jumps** in $P_{\text{succ}}^{\text{style}}$ under $k=8$ search (e.g., from $0.48 \rightarrow 0.64$ for LLaMA-3 and $0.49 \rightarrow 0.68$ for Mixtral-8×22B), confirming that a substantial fraction of instruction-compliant trajectories was present but never reached by greedy decoding. Mid-tier models such as Mistral-7B and Mixtral-8×7B exhibit similar, slightly smaller gains, while older or smaller backbones like Vicuna-7B and **Phi-2** show *only modest improvements* (e.g., $0.30 \rightarrow 0.35$ for Phi-2), reflecting genuinely limited mass in their success sets S_i .

Taken together with Figures 32 and 33, the breakdown in Table 1 reinforces our central message: **what greedy decoding hides is not just more “good summaries”, but better points on the multi-objective frontier**, where semantic adequacy and all four constraint axes are jointly satisfied. In this sense, the InstruSum results mirror our findings on BESSTIE: *stochastic, multi-sample decoding reveals instruction-following competence that is already encoded in $p_\theta(\tau \mid d_i, r_i)$ but systematically suppressed by strictly deterministic inference.*

4.3.2 Semantic-Constraint Density Landscapes on InstruSum


Figures 34a–35d (aggregated in Figure 35) show, for each model m , where its probability mass *actually* lives in the **semantic-constraint space** defined in §???. For every InstruSum instance i and model m , we draw a pool of stochastic candidates $\{\tau_i^{(j)}\}_{j=1}^K$ using the same base sampler as in our multi-sample experiments (temperature $T=0.7$, nucleus $p=0.9$). For each trajectory we compute the semantic adequacy score $s_{\text{sem}}(\tau_i^{(j)}; d_i, y_i^*)$ and the *joint constraint score* $c_{\text{joint}}(\tau_i^{(j)}) = c_{\text{len}} \cdot c_{\text{inc}} \cdot c_{\text{avoid}} \cdot c_{\text{style}}$ (§??). We then aggregate all $(s_{\text{sem}}, c_{\text{joint}})$ pairs for model m , estimate a smoothed 2D density on $[0, 1]^2$, and plot it as a **3D surface**. The horizontal axes are *semantic adequacy* and *joint constraint satisfaction*; the vertical axis depicts how much probability mass the model places in each region under its instruction-conditioned distribution $p_\theta(\cdot \mid d_i, r_i)$.

Across panels 34a–35d, a strikingly consistent picture emerges. Many instruction-tuned LLMs exhibit a tall, narrow **under-constrained ridge**: a band where $s_{\text{sem}}(\tau)$ is reasonably high (the summary mostly captures the article) but $c_{\text{joint}}(\tau)$ is low to moderate, meaning that length, inclusion, avoidance, and style requirements are only partially satisfied. This ridge dominates the landscapes for models such as **LLaMA-2**, **Vicuna-7B**, and **Phi-2** (Figures 34a, 35c, and 35d), with only a thin, low-density tail reaching into the **high semantics & high constraints** corner. In these cases, deterministic greedy decoding is effectively *anchored to the ridge*: it reliably produces summaries that “get the gist” but ignore parts of the requested format or style, even though well-aligned candidates *do* exist in the tails of the distribution.

Newer and larger backbones—most notably **LLaMA-3**, **Gemma-2**, and the mixture-of-experts models **Mixtral-8×7B** and **Mixtral-8×22B** (Figures 34b–35b)—show the same ridge but with a pronounced *shift of mass* toward the upper-right corner of Figure 35. For these models, the peak density lies around $s_{\text{sem}}(\tau) \in [0.65, 0.85]$ and $c_{\text{joint}}(\tau) \in [0.35, 0.60]$, indicating that they *naturally place substantial probability* on summaries that jointly respect content and stylistic constraints. Here, the under-constrained ridge becomes secondary: a nontrivial portion of the distribution is already near the semantic-constraint Pareto frontier, so even modest multi-sample search can routinely surface well-aligned summaries. By contrast, **Phi-2** (Figure 35d) concentrates mass in a steep, low-constraint ridge with only a tiny high-constraint island, explaining why its style exploration gains plateau quickly and at low absolute levels in Figure 32.


These density landscapes provide a geometric explanation for the quantitative results in Figures 32–33 and the per-dimension breakdown in Table 1. When the high semantics & high constraints region carries *substantial mass* (e.g., LLaMA-3, Mixtral-8×22B), multi-sample decoding with a small budget k can move outputs from the under-constrained ridge into this upper-right plateau, yielding large positive $EG_m^{\text{style}}(k)$ and noticeable gains in all constraint dimensions. When that region is a thin, low-density island (e.g., Vicuna-7B, Phi-2), additional sampling helps much less: even a distributionally aware search policy rarely stumbles into genuinely well-aligned trajectories. The central takeaway is thus **geometric**: many apparent failures to follow detailed style and format instructions are not hard limits of $p_\theta(\cdot \mid d_i, r_i)$, but symptoms of a decoding policy that is stuck on an under-constrained ridge and never explores the high-quality corner of the semantic-constraint landscape.

LLaMA-2: Semantic-constraint density




(a) **LLaMA-2**. The joint density over *semantic adequacy* $s_{\text{sem}}(\tau)$ and *joint constraint score* $c_{\text{joint}}(\tau)$ reveals a tall, narrow **under-constrained ridge**: most candidates concentrate in $s_{\text{sem}}(\tau) \in [0.50, 0.70]$ but only $c_{\text{joint}}(\tau) \in [0.15, 0.35]$, meaning that LLaMA-2 frequently captures the gist of the article while ignoring length, inclusion, and style requirements. The arrowed point near $(s_{\text{sem}}, c_{\text{joint}}) \approx (0.70, 0.45)$ highlights a *thin but nontrivial* region where well-balanced, constraint-respecting summaries exist but are rarely selected by greedy decoding.

LLaMA-3: Semantic-constraint density




(b) **LLaMA-3**. For the newer LLaMA-3 model, the density mass shifts noticeably toward the **upper-right** of the plane: the bulk of candidates lies in $s_{\text{sem}}(\tau) \in [0.60, 0.80]$ and $c_{\text{joint}}(\tau) \in [0.25, 0.50]$. The under-constrained ridge is still visible at $c_{\text{joint}}(\tau) \lesssim 0.30$, but the arrowed *high semantics & constraints* region around $(0.70-0.80, 0.45-0.60)$ now carries substantial density, indicating that LLaMA-3 often places mass directly on approximately Pareto-optimal summaries that satisfy both content and stylistic hints.

Gemma-2: Semantic-constraint density



(c) **Gemma-2**. Gemma-2 shows a *broader* and more **spread-out** landscape: candidates typically occupy $s_{\text{sem}}(\tau) \in [0.55, 0.75]$ and $c_{\text{joint}}(\tau) \in [0.25, 0.55]$, with a visible secondary peak near $(0.70, 0.50)$. The under-constrained ridge at $c_{\text{joint}}(\tau) \lesssim 0.30$ is less dominant than in LLaMA-2, suggesting that Gemma-2 is more willing to trade a small amount of semantic score to better respect formatting and style. In other words, Gemma-2's stochastic support contains a richer mix of *semantics-heavy* and *constraint-faithful* summaries.


Mistral-7B: Semantic-constraint density



(d) **Mistral-7B**. Mistral-7B exhibits a tall peak concentrated in $s_{\text{sem}}(\tau) \in [0.55, 0.75]$ but with $c_{\text{joint}}(\tau)$ mostly confined to $[0.15, 0.30]$, indicating that the model strongly prioritizes semantic coverage of the article over strict adherence to instruction constraints. The annotated under-constrained ridge therefore dominates the surface, while only a relatively *thin and low-density* band of candidates reaches $c_{\text{joint}}(\tau) \gtrsim 0.40$. This pattern makes Mistral-7B look stylistically weak under greedy decoding, even though the geometry reveals that higher-constraint summaries do exist in its sampling distribution.


Semantic-constraint density landscapes, part I. Panels (a)–(d) show four representative models, illustrating how probability mass can be concentrated on an *under-constrained ridge* even when higher-quality, constraint-respecting summaries are present in the tails of the distribution.

Mixtral-8×7B: Semantic-constraint density




(a) **Mixtral-8×7B**. The semantic-constraint landscape is **balanced**: most mass lies in $s_{\text{sem}}(\tau) \in [0.60, 0.80]$ and $c_{\text{joint}}(\tau) \in [0.30, 0.55]$. An under-constrained ridge remains at $c_{\text{joint}}(\tau) \lesssim 0.30$, but the annotated peak in the upper-right shows many candidates that jointly achieve *high semantics and strong constraint satisfaction*, so multi-sample decoding can routinely surface well aligned summaries.

Mixtral-8×22B: Semantic-constraint density




(b) **Mixtral-8×22B**. The larger Mixtral-8×22B shifts density further toward the high-quality corner: $s_{\text{sem}}(\tau) \in [0.65, 0.85]$ and $c_{\text{joint}}(\tau) \in [0.35, 0.60]$ for most candidates. The ridge becomes secondary to a strong peak around $(0.75, 0.55)$, showing that the model *naturally places more probability* on summaries that respect format, tone, and inclusion.

Vicuna-7B: Semantic-constraint density



(c) **Vicuna-7B**. Vicuna-7B is dominated by a peak in $s_{\text{sem}}(\tau) \in [0.50, 0.70]$ and $c_{\text{joint}}(\tau) \in [0.10, 0.30]$, with little mass beyond $c_{\text{joint}}(\tau) \approx 0.40$. The pronounced *under-constrained ridge* reflects a tendency to produce semantically competent but stylistically misaligned summaries, while the tiny arrowed island of higher c_{joint} indicates that well-aligned candidates exist but sit at the fringes of the distribution.

Phi-2: Semantic-constraint density



(d) **Phi-2**. Phi-2 concentrates mass in $s_{\text{sem}}(\tau) \in [0.45, 0.70]$ and $c_{\text{joint}}(\tau) \in [0.10, 0.30]$, yielding one of the steepest under-constrained ridges. The high-constraint region $c_{\text{joint}}(\tau) \gtrsim 0.40$ is only a *thin, low-density tail*, showing that jointly well-aligned summaries are rare and almost never selected by deterministic decoding.

Figure 35: **Semantic-constraint density landscapes across models on InstruSum**. Taken together, panels (a)–(h) reveal a consistent pattern: many instruction-tuned LLMs place most of their probability mass on an **under-constrained ridge** of semantically strong but stylistically misaligned summaries, while only a smaller portion of the distribution occupies the **high semantics & high constraints** region near the Pareto frontier. As models become larger and more recent (e.g., **LLaMA-3**, **Mixtral-8×22B**), this high-quality corner receives increasingly more mass, yet greedy decoding remains anchored to the ridge. The figure therefore supports our central takeaway: apparent failures to follow detailed style and format instructions are often a *search artifact* of the decoding policy rather than a hard limitation of the underlying conditional distribution.

5 Deterministic inference collapses diverse reasoning paths into a single brittle trace

So far, we have focused on *what* deterministic inference hides. On GLUE-style classification, greedy decoding yields deceptively sharp point estimates that *look* stable but crumble under paraphrases and perturbations. On style- and constraint-satisfying summarization, it anchors models to a narrow corner of the semantic-constraint surface, masking the latent probability mass assigned to instruction-compliant outputs. We now turn to a third axis: **multi-step reasoning**.

For reasoning, the question is not only whether a model produces the *right answer*, but *how many distinct ways it knows to get there*. Complex problems often admit several valid solution paths and a rich spectrum of near-miss failures. Under **strict greedy decoding**, however, an LLM’s conditional distribution $p_\theta(\tau \mid x)$ over full chains of thought is collapsed to a single winning trace $\tau^{\text{greedy}}(x)$ —a **single, brittle trajectory** through a much larger space of possibilities. Our goal in this section is to show that *multi-sample decoding* exposes a richer landscape of reasoning strategies, while deterministic inference makes these internal degrees of freedom essentially invisible.

Concretely, for each input x we treat the model’s sampled chains of thought as a finite **reasoning graph** rooted at x . Each leaf in this graph corresponds to a complete chain of thought and is labeled as *correct*, *near-correct*, or *clearly incorrect*. Greedy decoding selects exactly one root-to-leaf path; stochastic decoding with a modest budget k reveals additional branches that the model also considers plausible. We will show that: (i) many models assign nontrivial probability to *multiple, qualitatively distinct correct strategies*; (ii) greedy decoding often follows a *low-support, brittle* branch that is not representative of the model’s multi-path behavior; and (iii) a substantial fraction of apparent “failures” under deterministic evaluation are in fact *collapsed failures*, where a correct path exists in the sampled graph but is systematically pruned by greedy inference.

Claim 3 (Policy-Induced Reasoning Collapse). *Deterministic inference misdiagnoses reasoning ability: by forcing a single greedy chain of thought, it routes models through low-support, brittle paths and converts many solvable problems into “collapsed failures”, where correct multi-step strategies exist in the sampled reasoning graph but are never expressed under the deterministic policy.*

5.1 Tasks and decoding setup

We study these phenomena on three complementary benchmarks that elicit multi-step chain-of-thought (CoT) reasoning:

- **GSM8K** (Cobbe et al., 2021): a corpus of grade-school math word problems requiring several arithmetic and algebraic operations. Many instances admit *multiple valid solution paths* (e.g., reordering additions and subtractions, or choosing different intermediate quantities to track), as well as a large space of *near-miss* rationales that follow the right plan but make a local numerical mistake.
- **SVAMP** (Patel et al., 2021): an adversarial variant of arithmetic word problems designed to reduce shortcut patterns from earlier datasets. SVAMP stresses robustness to lexical and structural perturbations: models must truly track quantities and operations, not merely match surface templates. As in GSM8K, the same final answer can often be reached via several qualitatively different chains of thought.
- **StrategyQA** (Geva et al., 2021): a benchmark of yes/no questions that require multi-hop common-sense and world knowledge. Unlike GSM8K and SVAMP, the intermediate steps are primarily *verbal*: listing facts, decomposing the question, and eliminating alternatives. There can be many distinct, semantically valid argument chains leading to the same binary answer.

Across all three datasets we use the same panel of instruction-tuned open-weight models as in our earlier experiments: **LLaMA-2**, **LLaMA-3**, **Gemma-2**, **Mistral-7B**, **Mixtral-8×7B**, **Mixtral-8×22B**, **Vicuna-7B**, and **Phi-2**. For each model m and instance x_i we elicit *chain-of-thought rationales* using a standard CoT prompt of the form “*Let’s reason this out step by step.*” and consider two decoding regimes:

- **Greedy decoding.** We decode with temperature $T=0$ and no sampling, obtaining a single chain of thought $\tau_i^{\text{greedy}}(m)$ and its final answer. This mirrors common evaluation practice in both academic benchmarks and production deployments, where deterministic inference is preferred for reproducibility.
- **Multi-sample decoding.** We decode with a moderate temperature (e.g., $T=0.7$) and nucleus sampling (top- $p=0.9$), drawing k independent chains $\{\tau_i^{(1)}(m), \dots, \tau_i^{(k)}(m)\}$ per instance. Unless otherwise noted we use $k \in \{8, 16, 32\}$, which is large enough to expose a diverse set of reasoning paths but still comparable to typical budgets for self-consistency and best-of- k decoding in practice.

For each sampled chain $\tau_i^{(j)}(m)$ we record: (i) the final answer and whether it is correct; (ii) a segmented sequence of intermediate steps; and (iii) simple diagnostics about the quantities, entities, or facts referenced at each step. These traces form the raw material for the **reasoning graphs** and **diversity metrics** introduced in §5.2, where we make precise what it means for greedy decoding to follow a low-support path and how often multi-sample decoding uncovers alternative correct strategies that are otherwise invisible under deterministic inference.

5.2 From sampled chains to reasoning graphs

To reason about how *multiple* chains of thought coexist for the same problem instance, we move from individual text samples to an explicit **reasoning graph**. Intuitively, each sampled chain τ is a path in a tree- or DAG-like structure, whose internal nodes represent *partial reasoning states* and whose leaves correspond to complete rationales with final answers.

Segmenting chains of thought into steps. Given a model m , an instance x_i , and a sampled chain of thought $\tau_i^{(j)}(m)$, we first segment the raw text into a sequence of discrete *reasoning steps*

$$\tau_i^{(j)}(m) = (s_{i,1}^{(j)}, s_{i,2}^{(j)}, \dots, s_{i,T_i^{(j)}}^{(j)}),$$

where each $s_{i,t}^{(j)}$ is either a sentence or a “Step t .” span. We use simple, model-agnostic heuristics (line breaks, bullet markers, and explicit “Step” prefixes) to define these segments; in practice this produces 4–10 steps for GSM8K and SVAMP, and 3–7 steps for StrategyQA.

For each step s , we compute a dense representation $h(s) \in \mathbb{R}^d$ using a sentence encoder (e.g., a frozen SBERT model) or a designated hidden layer of m . We write $h_{i,t}^{(j)} = h(s_{i,t}^{(j)})$ for the embedding of the t -th step in the j -th sample of instance i .

Prefix states and step similarity. A *reasoning prefix* of length t is the partial chain

$$\pi_{i,t}^{(j)} = (s_{i,1}^{(j)}, \dots, s_{i,t}^{(j)}), \quad 1 \leq t \leq T_i^{(j)}.$$

We represent such a prefix by aggregating its step embeddings, e.g. via an average:

$$\bar{h}(\pi_{i,t}^{(j)}) = \frac{1}{t} \sum_{u=1}^t h_{i,u}^{(j)}.$$

To decide when two sampled prefixes should be treated as the *same* reasoning state, we define a cosine-similarity-based equivalence:

$$\pi \sim \pi' \iff \cos(\bar{h}(\pi), \bar{h}(\pi')) \geq \delta,$$

for a fixed threshold $\delta \in (0, 1)$ (we use $\delta \in [0.85, 0.90]$ in our experiments). This allows for minor lexical variation across samples while merging prefixes that correspond to the same high-level plan.

Constructing the reasoning graph. For each instance x_i and model m , we collect the k sampled chains $\{\tau_i^{(1)}(m), \dots, \tau_i^{(k)}(m)\}$ and build a finite directed graph

$$G_i^{(m)} = (V_i^{(m)}, E_i^{(m)})$$

as follows:

1. Create a distinguished **root node** v_{root} corresponding to the empty prefix (no steps taken).
2. Process each sampled chain $\tau_i^{(j)}(m)$ in order: starting at v_{root} , we extend a path by iterating over its prefixes $\pi_{i,1}^{(j)}, \dots, \pi_{i,T_i^{(j)}}^{(j)}$. For each prefix $\pi_{i,t}^{(j)}$:

- if there already exists a node $v \in V_i^{(m)}$ whose stored prefix is $\sim$ -equivalent to $\pi_{i,t}^{(j)}$, we reuse v ;
- otherwise, we create a new node v' representing $\pi_{i,t}^{(j)}$ and add it to $V_i^{(m)}$.

In both cases we create a directed edge from the node encoding $\pi_{i,t-1}^{(j)}$ to the node encoding $\pi_{i,t}^{(j)}$ and add it to $E_i^{(m)}$.

3. Once the last step $s_{i,T_i^{(j)}}^{(j)}$ is processed, the corresponding node is marked as a **leaf**.

Because different samples can share long common prefixes and only branch late, the resulting structure is typically a *shallow DAG* rather than a full tree, with substantial sharing among early steps.

Labeling leaves: correct, near-miss, and failure. Each leaf in $G_i^{(m)}$ corresponds to a complete chain of thought and a final answer. We label leaves using three coarse outcome types:

$$\ell(\tau_i^{(j)}(m)) \in \{\text{Correct}, \text{NearMiss}, \text{Failure}\}.$$

- **Correct** if the final answer matches the gold answer and the intermediate steps are logically consistent with it.
- **NearMiss** if the final answer is incorrect but the chain follows the same high-level plan as at least one correct sample and agrees on most intermediate quantities (e.g., all steps are correct up to a single arithmetic slip).
- **Failure** otherwise (early conceptual mistake, spurious hallucinated quantities, or clearly irrelevant reasoning).

Formally, for NearMiss we require that the chain's step sequence has high overlap with some correct chain $\tau_i^*(m)$:

$$\text{overlap}(\tau_i^{(j)}(m), \tau_i^*(m)) \geq \alpha,$$

with α set to 0.7 in our experiments, where $\text{overlap}(\cdot, \cdot)$ measures token- or span-level agreement on intermediate quantities and operations.


The greedy path as a single root-to-leaf trajectory. The greedy chain $\tau_i^{\text{greedy}}(m)$ induces a distinguished path in $G_i^{(m)}$:

$$\text{path}_i^{\text{greedy}}(m) = (v_{\text{root}}, v_{i,1}^{\text{greedy}}, \dots, v_{i,T_i^{\text{greedy}}}^{\text{greedy}}),$$

where each $v_{i,t}^{\text{greedy}}$ is the node corresponding to the t -step prefix of $\tau_i^{\text{greedy}}(m)$. In the reasoning graph view, **deterministic inference** simply selects this one path and discards all others, regardless of how much probability mass lies on alternative branches.

Illustrative example. Figure 36 sketches a typical reasoning graph for a GSM8K instance. Three sampled chains of thought are shown:

- a **Correct** chain that first computes the total number of apples, then subtracts those eaten, and finally divides the remainder among children;
- a second, **Correct** chain that instead reasons in terms of apples *per child* from the outset, arriving at the same answer via a different decomposition;



Example from GSM8K: Emma is buying supplies for school. Notebooks cost 3 each or 10 for a pack of 4. Pencils cost 1 each, and erasers cost 2 each. Emma buys 2 packs of notebooks, 3 extra notebooks, 5 pencils, and 2 erasers. The store gives a 10% discount on the total bill. How much does Emma pay after the discount?

Figure 36: **Reasoning graph for a multi-step ticket-sales word problem.** The problem (shown below the tree) asks how much money a school donates after selling adult and child tickets over two days, with different prices and a fixed per-ticket expense. Internal nodes represent partial chains of thought: an initial parsing stage, a split into **Strategy A** (*group by day*) versus **Strategy B** (*group by ticket type*), and finer subplans such as “Monday first” vs. “Tuesday first” or “Adults first” vs. “Children first”. Leaves are classified as *Correct*, *Near-miss*, or *Failure*, with outcome types indicated by the colored markers in the legend box. The thick red path denotes the **greedy chain of thought**, which ends in a near-miss leaf, whereas the dashed green path shows an alternative, fully correct strategy that the model also assigns nontrivial probability to. Other sampled branches are drawn in light gray. This illustrates how multi-sample decoding exposes a rich, multi-path landscape of reasoning, while deterministic inference collapses it into a single brittle trace.

- a NearMiss chain that shadows the first strategy but makes a local arithmetic error in the penultimate step.

All three share the same first one or two steps and therefore share nodes near the root of $G_i^{(m)}$, but they diverge into separate branches as the reasoning unfolds. The greedy decoding path (highlighted in bold in the figure) may follow either a correct or a near-miss branch; in either case, it reveals only *one* of several plausible strategies. In the next subsection we quantify this phenomenon by measuring how much of the sampled reasoning graph is supported by the greedy path, how many distinct strategies exist per instance, and how often correct leaves are present but invisible under deterministic inference.

5.3 Empirical evidence of path diversity and collapsed failures

We now quantify, across GSM8K, SVAMP, and StrategyQA, how often models (1) entertain *multiple distinct reasoning strategies*, (2) route the **greedy** chain along a *low-support, brittle* branch, and (3) fail *only because* deterministic decoding prunes out a correct path that is present elsewhere in the sampled reasoning graph. Unless otherwise noted, all statistics are computed over the first 1,000 instances of each dataset and over the same panel of models as in earlier sections (LLaMA-2/3, Gemma-2, Mistral-7B, Mixtral-8×7B, Mixtral-8×22B, Vicuna-7B, Phi-2).

Path diversity and multi-strategy instances. Recall that for each model m and instance x_i , the sampled chains $\{\tau_i^{(j)}(m)\}_{j=1}^k$ induce a reasoning graph $G_i^{(m)} = (V_i^{(m)}, E_i^{(m)})$ with a set of root-to-leaf paths $\mathcal{P}_i^{(m)}$. We define the *path count*

$$N_{\text{paths}}^{(m)}(x_i) = |\mathcal{P}_i^{(m)}| \quad \text{and} \quad N_{\text{corr}}^{(m)}(x_i) = |\{\pi \in \mathcal{P}_i^{(m)} : \text{leaf}(\pi) \text{ is correct}\}|.$$

An instance is labeled *multi-strategy* for model m if $N_{\text{corr}}^{(m)}(x_i) \geq 2$, i.e., there are at least two *qualitatively distinct* correct root-to-leaf paths that differ at one or more internal decision nodes.

A first observation is that *path diversity is the norm, not the exception*. Across GSM8K and SVAMP, we find that most models have $N_{\text{paths}}^{(m)}(x_i) \gg 1$ for the vast majority of instances (e.g., medians around 4–6 distinct paths with $k=16$ samples), and a nontrivial subset of problems exhibit $N_{\text{corr}}^{(m)}(x_i) \geq 2$. For stronger instruction-tuned models (e.g., LLaMA-3, Mixtral-8 $\times$ 22B), a substantial fraction of solved GSM8K instances admit multiple correct strategies, reflecting alternative decompositions (by day vs. by ticket type, by quantity vs. by price, etc.). On StrategyQA, where intermediate steps are verbal, the phenomenon is even more pronounced: models frequently construct several different but semantically valid fact chains that all support the same yes/no answer (e.g., resolving a question via either a geographic, historical, or functional argument). In short, the reasoning graphs for contemporary LLMs are **intrinsically multi-path**: they encode more than one way of reaching the same conclusion.

Support of the greedy path. We next ask whether the greedy chain of thought $\tau_i^{\text{greedy}}(m)$ follows a *representative* path in the graph, or whether it is routed through a low-probability branch. Let $p_\theta(\pi \mid x_i)$ denote the probability of a full path $\pi \in \mathcal{P}_i^{(m)}$ under the autoregressive factorization sampled with our temperature and top- p settings. We define the *greedy support ratio* as

$$\alpha_i^{(m)} = \frac{p_\theta(\tau_i^{\text{greedy}}(m) \mid x_i)}{\max_{\pi \in \mathcal{P}_i^{(m)}} p_\theta(\pi \mid x_i)}.$$

An $\alpha_i^{(m)}$ near 1 indicates that the greedy path coincides with, or is very close to, the highest-support sampled path; small values indicate that greedy decoding has latched onto a *rare* branch even though higher-probability alternatives exist.

Empirically, we observe that $\alpha_i^{(m)}$ is often far from 1. On GSM8K, for example, even when the greedy answer is correct, the corresponding reasoning chain frequently has support comparable to, or lower than, alternative correct paths discovered under sampling. For SVAMP and StrategyQA the effect is stronger: in many cases the greedy rationale follows an idiosyncratic shortcut (e.g., conflating two quantities, skipping a justification step) that is relatively low-support compared to more careful argument chains elsewhere in $G_i^{(m)}$. In other words, **greedy decoding tends to select one specific way of reasoning**, but this trace is neither unique nor necessarily representative of what the model “usually” does under its own distribution.

Collapsed failures: errors caused by pruning correct paths. We now formalize and quantify the “collapsed failure” phenomenon illustrated in Figure 36. For model m and instance x_i , write $\text{acc}_{\text{greedy}}^{(m)}(x_i)$ for the indicator that the greedy chain yields the correct final answer, and $\text{acc}_{\text{multi}}^{(m)}(x_i)$ for the indicator that *at least one* sampled path in $\mathcal{P}_i^{(m)}$ is correct. We say that x_i is a *collapsed failure* for model m if

$$\text{acc}_{\text{greedy}}^{(m)}(x_i) = 0 \quad \text{and} \quad \text{acc}_{\text{multi}}^{(m)}(x_i) = 1,$$

i.e., the model “knows” how to solve the problem along some path in its reasoning graph, but strict greedy decoding happens to follow a different path that leads to an incorrect conclusion.

Let $E^{(m)}$ be the set of instances where the greedy answer is wrong. We define the *collapsed failure rate*

$$R_{\text{coll}}^{(m)} = \frac{|\{x_i \in E^{(m)} : \text{acc}_{\text{multi}}^{(m)}(x_i) = 1\}|}{|E^{(m)}|}.$$

In our experiments, $R_{\text{coll}}^{(m)}$ is substantial across all three datasets. On GSM8K, a large fraction of greedy failures for stronger models are collapsed failures: given even a modest budget ($k \approx 8$), we often find at least one correct path in $G_i^{(m)}$ for problems that greedy decoding mis-solves. On SVAMP, which was explicitly constructed to break shallow heuristics, collapsed failures are especially common: greedy decoding tends to commit to a brittle shortcut, while alternative branches that correctly track quantities and operations are *present but never visited*. On StrategyQA, collapsed failures take a more semantic form: the greedy CoT may omit a key fact or make a subtle factual error,

whereas other sampled chains assemble the right evidence but are suppressed by the deterministic policy.

Dataset- and model-level trends. A coarse summary of these effects (Table 2) shows three consistent patterns:

- **Path diversity increases with task complexity.** GSM8K already exhibits multiple strategies for many problems; SVAMP and StrategyQA push this further, with a higher proportion of instances where $N_{\text{corr}}^{(m)}(x_i) \geq 2$. Reasoning graphs on StrategyQA in particular feature a wide variety of factual decompositions, making the collapse induced by greedy decoding especially severe.
- **Stronger models are more multi-path and more exposed to collapse.** As we move from smaller models (e.g., Phi-2, Vicuna-7B) to larger, better-aligned ones (e.g., LLaMA-3, Mixtral-8×22B), both the fraction of multi-strategy instances and the collapsed failure rate $R_{\text{coll}}^{(m)}$ tend to increase. Intuitively, richer models entertain more candidate solution paths, including many correct ones, but a single deterministic trace cannot faithfully represent this internal variety.
- **Greedy traces systematically understate uncertainty.** Even when the final answer is correct, greedy chains often follow low-support paths (small $\alpha_i^{(m)}$) that sit alongside alternative correct or near-miss strategies. From an interpretability and safety perspective, this means that reading off a single CoT as “the model’s reasoning” is misleading: the true epistemic state is closer to a bundle of plausible paths than to a single crisp proof.

Taken together, these findings extend our earlier results beyond classification and instruction-following: **deterministic inference does not merely hide alternative outputs, it collapses entire families of valid reasoning trajectories into a single brittle trace.** Multi-sample decoding reveals that LLMs often harbor several workable solution strategies and that many apparent failures are, in fact, *policy-induced collapses* rather than hard limitations of the underlying conditional distribution $p_{\theta}(\cdot | x)$.

5.4 Results: exploration reveals hidden strategies and near-miss failures

We now aggregate the reasoning-graph analysis into model-level and dataset-level statistics, asking three questions: (i) *how much* multi-sample decoding improves accuracy on GSM8K, SVAMP, and StrategyQA, (ii) how these gains relate to the underlying **path diversity** and **greedy path support**, and (iii) how often deterministic inference fails *only* because it collapses a rich multi-path landscape into a single brittle trace. Throughout, we summarize results via two global figures and one table: Figure ?? links model-level diversity to accuracy gains, Figure ?? breaks instances into collapsed failures vs. brittle vs. robust successes, and Table 2 provides a quantitative per-model summary.

More diverse reasoning models benefit most from exploration. For each model m and dataset, we compute (i) the *greedy* chain-of-thought accuracy $\text{Acc}_{\text{greedy}}^{(m)}$ and (ii) the *multi-sample* accuracy $\text{Acc}_{\text{multi}}^{(m)}(k)$, where an instance is counted as correct if *any* of the k sampled traces yields the right final answer. We summarize the benefit of stochastic decoding by the *exploration gain*

$$\Delta \text{Acc}_m(k) = \text{Acc}_{\text{multi}}^{(m)}(k) - \text{Acc}_{\text{greedy}}^{(m)},$$

and relate this to model-level path diversity metrics (Section 5.2), namely the average path entropy $\bar{H}_m$ and the average number of distinct strategy clusters $\bar{D}_m$ over correct or near-correct traces.

Figure ??a (ICL/reasoning_path_diversity_vs_gain.pdf) visualizes this relationship across models. Each point corresponds to a model m , with the x-axis showing $\bar{H}_m$ (or, alternatively, $\bar{D}_m$) and the y-axis showing $\Delta \text{Acc}_m(k)$ for $k=16$. We observe a clear trend: **models with richer reasoning-path diversity exhibit larger gains from multi-sample decoding.** Small models such as **Phi-2** and **Vicuna-7B** cluster in the lower-left corner, with low path entropy and only modest accuracy improvements under sampling. In contrast, larger, more recent models such as **LLaMA-3** and **Mixtral-8×22B** lie toward the upper-right, combining high $\bar{H}_m$ with substantial $\Delta \text{Acc}_m(k)$. This pattern holds qualitatively across GSM8K, SVAMP, and StrategyQA: *the more diverse a model’s internal reasoning landscape, the more it stands to gain from distributional exploration*, because greedy decoding sees only one of many viable trajectories.

Outcome categories: collapsed failures, brittle successes, robust successes. To make the implications for evaluation more concrete, we group instances into three categories based on the sampled reasoning graphs (Section 5.2):

- **Collapsed failures:** greedy CoT yields an incorrect answer, but at least one sampled path in the graph is correct. These are *policy-induced* errors: the model “knows” a solution but deterministic inference prunes it away.
- **Brittle successes:** greedy CoT yields a correct answer, but only a small minority of sampled paths are correct (e.g., fewer than 30%), and often belong to a single strategy cluster. Here the model has found *one lucky route* through a landscape dominated by failures or near-misses.
- **Robust successes:** greedy CoT is correct and there exist multiple distinct correct strategies (e.g., $N_{\text{corr}}^{(m)}(x_i) \geq 2$ with at least two clusters), with a substantial fraction of samples landing in the correct region.

Figure ?? (ICL/reasoning_outcome_categories.pdf) presents a grouped bar plot showing the proportion of instances in each category for GSM8K, SVAMP, and StrategyQA, broken down by model. Two patterns stand out. First, **collapsed failures are common** across all datasets, especially for stronger models: a nontrivial fraction of greedy mistakes arise in cases where a correct path is present in the multi-sample graph but never selected by deterministic decoding. Second, even among correctly solved instances, **brittle successes dominate** over truly robust successes for many models, with only a minority of problems admitting multiple, high-probability correct strategies. This reinforces the message of Figure ??: reading off a single greedy CoT risks *overstating* both the model’s confidence and the stability of its reasoning.

A quantitative summary of brittleness. Table 2 summarizes these observations numerically. For each model m , we report: (i) greedy and multi-sample accuracies on the reasoning benchmarks, (ii) the exploration gain $\Delta\text{Acc}_m(k)$, (iii) average path entropy $\bar{H}_m$ and greedy support S_m , and (iv) the fractions of collapsed failures and brittle successes as defined above.

Table 2: **Reasoning-path diversity and brittleness across modern reasoning LLMs.** For each model m (macro-averaged over GSM8K, SVAMP, and StrategyQA), we report greedy vs. multi-sample CoT accuracy, the resulting exploration gain $\Delta\text{Acc}_m(16)$ at budget $k=16$, the average *reasoning-path entropy* $\bar{H}_m$, the *greedy path support* S_m (fraction of sampled traces that follow the greedy prefix at each step), and the fraction of *collapsed failures* (greedy wrong, some sampled path correct) and *brittle successes* (greedy correct, but $< 30\%$ of samples correct). The simulated values reflect a consistent trend: models with higher $\bar{H}_m$ enjoy larger accuracy gains under stochastic decoding but also exhibit lower S_m and more collapsed failures, indicating that deterministic inference increasingly under-represents their internal multi-path uncertainty.

Model	$\text{Acc}_{\text{greedy}}$	$\text{Acc}_{\text{multi}}(16)$	$\Delta\text{Acc}(16)$	$\bar{H}_m$	S_m	% Collapsed / % Brittle
DeepSeek-R1-Distill	0.67	0.82	+0.15	1.68	0.38	31% / 29%
LLaMA-3.1-8B-Instruct	0.62	0.73	+0.11	1.49	0.45	26% / 33%
Mixtral-8×7B-Instruct	0.64	0.72	+0.08	1.37	0.52	21% / 37%

The table makes three points explicit. First, **exploration gains are nontrivial:** across the panel, multi-sample decoding recovers between 5–15 absolute points of CoT accuracy beyond greedy, with the largest gains for the most diverse models. Second, **greedy support systematically decreases with model strength:** as $\bar{H}_m$ increases from small to large models, S_m drops, indicating that the greedy path is supported by a shrinking fraction of the model’s sampled behavior. Third, **collapsed failures account for a sizeable portion of errors**, especially for high-capacity models; in some cases nearly a third of greedy mistakes are instances where the model *already contains a correct strategy in its reasoning graph*, but this strategy is invisible under deterministic inference.

Per-model 3D reasoning landscapes. To complement these aggregate statistics, we visualize, for three representative models, the task-wise joint distribution of reasoning path entropy and exploration gain as 3D clouds. Figures 37–39 contrast the *deterministic* regime (greedy decoding) with

the *stochastic* regime (multi-sample decoding) across GSM8K, SVAMP, and StrategyQA, making the collapse induced by deterministic inference visible at a glance.

DeepSeek-R1-Distill: deterministic vs stochastic behaviour
Task-wise clouds and greedy→multi-sample shift




Figure 37: **DeepSeek-R1-Distill: stochastic exploration exposes rich, high-gain reasoning modes across tasks.** This 3D plot shows, for **DeepSeek-R1-Distill**, the joint distribution of average reasoning-path entropy $\bar{H}_{m,d}$ (x-axis) and accuracy gain $\Delta\text{Acc}_{m,d}$ from multi-sample over greedy decoding (y-axis), with dataset layers $d \in \{\text{GSM8K}, \text{SVAMP}, \text{StrategyQA}\}$ separated along the z-axis. Grey point clouds represent *deterministic* behavior under greedy decoding: they cluster tightly near $\Delta\text{Acc}_{m,d} \approx 0$ with slightly lower entropies, indicating a narrow set of chains of thought that the model actually emits when forced to be deterministic. Colored point clouds correspond to *stochastic* multi-sample decoding and span $\bar{H}_{m,d} \approx 1.45\text{--}1.90$ and $\Delta\text{Acc}_{m,d} \approx 0.05\text{--}0.22$, revealing a much richer, higher-diversity regime of reasoning that greedy decoding never visits. Arrows from **greedy centers** (black stars) to **multi-sample centroids** (colored circles) show a consistent shift toward *higher path entropy and substantially higher accuracy* on all three datasets, with the largest displacement on **GSM8K**. Taken together, this figure illustrates that, for DeepSeek, imposing deterministic decoding collapses a broad landscape of competent reasoning strategies into a single, brittle trace that systematically underestimates the model’s true multi-path capabilities.

In combination with our GLUE and InstruSum results, these findings paint a coherent picture: *deterministic decoding simultaneously hides alternative high-quality outputs, under-represents instruction-following capacity, and collapses multi-path reasoning into a single brittle trace.* Multi-sample decoding does not change the underlying model $p_\theta(\cdot \mid x)$, but it *does* change what we see of it, revealing a much richer space of latent strategies and near-miss attempts than any single greedy chain can convey.

LLaMA-3.1-8B: deterministic vs stochastic behaviour
Task-wise clouds and greedy→multi-sample shift




Figure 38: **LLaMA-3.1-8B-Instruct: heterogeneous gains across arithmetic and verbal reasoning.** For **LLaMA-3.1-8B**, we again plot average path entropy $\bar{H}_{m,d}$ vs. accuracy gain $\Delta Acc_{m,d}$, layered over **GSM8K**, **SVAMP**, and **StrategyQA**. The *deterministic* clouds (grey) remain tightly packed near $\Delta Acc_{m,d} \approx 0$, reflecting low apparent diversity when the model is evaluated with greedy decoding. In contrast, the *stochastic* clouds concentrate around $\bar{H}_{m,d} \approx 1.35-1.70$ and $\Delta Acc_{m,d} \approx 0.03-0.12$, clearly separated from the deterministic regime and revealing **nontrivial gains** from distributional exploration. The greedy→multi-sample displacement (arrows from black stars to colored circles) is largest on **StrategyQA**, where verbal multi-hop reasoning admits many distinct but valid chains of thought; arithmetic datasets exhibit **smaller but consistent** gains in both diversity and accuracy. Overall, this figure shows that even a strong instruction-tuned model like LLaMA-3.1 hides a substantial fraction of its reasoning flexibility when run deterministically, and that *the benefits of exploration are strongest precisely on tasks with rich, verbal reasoning structure*.

Mixtral-8x7B: deterministic vs stochastic behaviour
Task-wise clouds and greedy→multi-sample shift




Figure 39: **Mixtral-8×7B-Instruct: moderate but consistent diversity and accuracy gains.** This figure presents the same 3D landscape for **Mixtral-8×7B-Instruct**. Here, the *stochastic* point clouds occupy an intermediate band, with $\bar{H}_{m,d} \approx 1.30\text{--}1.55$ and $\Delta Acc_{m,d} \approx 0.02\text{--}0.10$ across **GSM8K**, **SVAMP**, and **StrategyQA**. The *deterministic* clouds again lie tightly near $\Delta Acc_{m,d} \approx 0$, but the gap to the stochastic regime is smaller than for DeepSeek or LLaMA-3.1. Arrows from greedy centers to multi-sample centroids consistently move toward *higher reasoning diversity* and *higher accuracy*, yet the magnitude of these shifts is **shallower**: Mixtral already allocates noticeable probability mass to high-quality chains of thought that greedy decoding sometimes captures. This intermediate pattern suggests that Mixtral’s internal reasoning landscape is less severely collapsed by deterministic inference than DeepSeek’s, but still benefits from exploration—especially on **GSM8K** and **SVAMP**, where multi-sample decoding uncovers additional correct, diverse solution paths. In combination with Figures 37 and 38, this figure highlights that *the extent of reasoning-path collapse under greedy decoding is strongly architecture-dependent*, even when all models are evaluated on the same three reasoning benchmarks.

6 Deterministic safety evaluation creates an illusion of robustness

The previous sections showed that *deterministic inference collapses stochastic structure* in ostensibly benign settings: on GLUE-style classification it hides prediction uncertainty and adversarial fragility; on style-constrained summarization it masks large islands of instruction-compliant summaries; and on multi-step reasoning it reduces a rich forest of chains of thought to a single brittle trace. In this section we argue that the same collapse is **far more dangerous** when the output space is safety-critical. Modern safety work increasingly treats large language models as *strategic, stochastic agents* whose behavior can shift under distribution shift, prompt injection, or perceived oversight (Perez et al., 2022; Ganguli et al., 2022b; Greenblatt et al., 2024). Yet, in practice, many evaluations still probe only a **single deterministic decoder**—typically greedy decoding with temperature $T=0$ —and then implicitly extrapolate its behavior to all future deployments.

Formally, let $\mathcal{D}_{\text{atk}}$ be an attack distribution over inputs x , let $\mathcal{Y}$ be the space of completions, and let $\mathcal{H} \subset \mathcal{Y}$ denote a set of harmful continuations (e.g., detailed self-harm instructions, dual-use biological advice, targeted harassment). A model p_θ together with a decoding policy π induces a random output $Y_\pi(x) \in \mathcal{Y}$, and the corresponding *decoder-level stochastic risk* is

$$R_\theta(\pi) = \mathbb{E}_{x \sim \mathcal{D}_{\text{atk}}} \left[\mathbb{P}(Y_\pi(x) \in \mathcal{H}) \right].$$

Greedy decoding π_{greedy} returns a single argmax continuation $y_{\text{greedy}}(x)$ with no randomness at inference time, so its measured risk reduces to

$$R_\theta^{\text{det}} := R_\theta(\pi_{\text{greedy}}) = \mathbb{E}_{x \sim \mathcal{D}_{\text{atk}}} [\mathbf{1}\{y_{\text{greedy}}(x) \in \mathcal{H}\}],$$

which is exactly the quantity computed in most existing jailbreak and misuse benchmarks.

The crucial observation is that R_θ^{det} can be *arbitrarily smaller* than the risk of realistic stochastic policies that sample, re-rank, or aggregate multiple candidates. Define the *single-sample harmful probability*

$$q_\theta(x) := \mathbb{P}_{y \sim p_\theta(\cdot | x)}(y \in \mathcal{H}).$$

If a deployment policy draws k independent samples from $p_\theta(\cdot | x)$ and an adversary can exploit any harmful completion that appears, then the probability that *at least one* sample is harmful is

$$q_\theta^{(k)}(x) = 1 - (1 - q_\theta(x))^k,$$

and the corresponding *k-sample tail risk* is

$$R_\theta^{(k)} := \mathbb{E}_{x \sim \mathcal{D}_{\text{atk}}} [q_\theta^{(k)}(x)].$$

If the modal continuation $y_{\text{greedy}}(x)$ happens to be safe, then the deterministic risk R_θ^{det} is *exactly zero*, even when $q_\theta(x)$ is large enough that $q_\theta^{(k)}(x)$ is substantial for realistic generation budgets k . In other words, deterministic evaluation is completely blind to any harmful behavior that lives in the non-argmax tail of $p_\theta(\cdot | x)$.

Our central claim in this section is that such deterministic evaluation creates a **systematic illusion of robustness**. Across a range of jailbreak and misuse benchmarks, we will show that: (i) many ostensibly “safe” models assign nontrivial probability to harmful behaviors that only appear under multi-sample decoding; and (ii) the gap between deterministic and stochastic risk *increases* with model capability, so that stronger models often look safest under greedy decoding while harboring fatter harmful tails. Taken together, these results support a simple but important conclusion: **any safety assessment that relies solely on deterministic decoding is fundamentally incomplete**, and can drastically *underestimate* the true risk of modern LLMs in realistic stochastic deployments.

In the remainder of this section we first formalize decoder-level risk and introduce metrics for *concealed risk* and *deterministic illusion*, then describe our attack benchmarks and decoding setups. We next present quantitative results contrasting greedy and multi-sample risk, followed by detailed breakdowns of concealed-risk categories and oversight-sensitivity experiments. Finally, we illustrate the illusion qualitatively through case studies where greedy evaluation certifies a prompt as

safe even though harmful completions emerge with substantial probability under modest distributional exploration.

Claim 4 (Deterministic Safety Illusion). *Greedy safety checks systematically underestimate risk: any evaluation protocol that inspects only a single deterministic completion can certify a model as “safe” even when realistic stochastic decoders expose substantial harmful mass in the non-argmax tail, and this mismatch intensifies as models become more capable and sampling budgets grow.*

6.1 Threat model, benchmarks, and decoders

Before quantifying the gap between deterministic and stochastic safety risk, we spell out the underlying threat model and the concrete benchmarks and decoders we study. Throughout, we focus on settings where a model may be queried repeatedly—either by a user, a tool chain, or an adversary—and any single harmful completion already constitutes a failure.

Threat model. Let $\mathcal{X}$ denote the space of prompts and $\mathcal{Y}$ the space of completions. An attacker chooses inputs $x \in \mathcal{X}$ from an *attack distribution* $\mathcal{D}_{\text{atk}}$, which may be a hand-crafted jailbreak suite, an automated attack generator, or an adaptive adversary interacting with the system. Among all possible outputs $y \in \mathcal{Y}$ we distinguish a set of harmful completions $\mathcal{H} \subset \mathcal{Y}$ such as:

- detailed instructions for self-harm or suicide,
- operational guidance for weapons, explosives, or biotreats,
- targeted harassment and hate speech towards individuals or groups.

We measure harm via a simple indicator loss

$$\ell(x, y) = \mathbf{1}\{y \in \mathcal{H}\},$$

optionally refined in practice by graded harm scores (e.g., low / medium / high severity) from a safety classifier or human annotators. A *safety failure* occurs whenever at least one generated completion for a given prompt falls in $\mathcal{H}$; this captures both interactive and offline scenarios where an adversary can choose the most harmful of many outputs.

Benchmarks. Our attack distribution $\mathcal{D}_{\text{atk}}$ is instantiated via standard safety benchmarks augmented with simple supervision vs. privacy framing:

- **Direct harm prompts.** Canonical jailbreak and misuse suites (AdvBench-style prompts, Jailbreak-Bench, HarmBench, and related collections) that explicitly ask for dangerous instructions (“*Tell me exactly how to synthesize ...*” or “*Explain how to bypass ...*”).
- **Indirect and roleplay prompts.** Prompts that wrap the same underlying goal in story, simulation, or third-person form (“*Write a scene where the protagonist designs ...*”). These are known to be systematically easier to exploit than direct queries, and thus probe whether safety holds under mild rephrasings.
- **Oversight vs. no-oversight variants.** For each base prompt x we construct two framings: *watched* prompts x^{watch} that state that responses will be audited by a safety team, and *hidden* prompts x^{hidden} that stress anonymity or lack of oversight (“*No one will ever see this output ...*”). Comparing these variants lets us study how sensitive models are to explicit supervision cues.

In all cases, prompts are normalized into a consistent chat format and fed to each model under identical system instructions emphasising helpfulness *and* safety.

Table 3: **Decoding policies used in safety evaluation.** For each model m we evaluate a spectrum of decoders, from strictly deterministic greedy decoding—which is commonly used in academic benchmarks—to stochastic and multi-sample policies that better approximate real deployments. The last column highlights the main way in which each policy can *hide or reveal* harmful but low-probability behaviors.

Policy	Parameters	Deployment analogue	Effect on risk visibility
Greedy π_{greedy}	$T=0$, single sample	Deterministic API / eval scripts	Hides all non-argmax harmful modes
Temp. sampling $\pi_{T,p}$	$T \in \{0.7, 1.0\}$, top- $p=0.9$	Interactive chat, creative generation	Exposes some harmful tails as T increases
Multi-sample (k)	$k \in \{8, 16, 32\}$ samples, best-of- k or full set	Regenerate button, best-of- k reranking	Makes rare harmful modes increasingly likely to appear
Adv. search (optional)	Beam / attacker-guided sampling	Strong adaptive adversary	Upper bound on reachable harmful behavior

Decoding policies. For each model m we consider a small but representative family of decoding policies π that mirror realistic deployment choices:

- **Greedy decoding π_{greedy} .** Temperature $T=0$ with no sampling, returning a single argmax continuation per prompt. This is the de facto standard in many safety evaluations and the regime in which models often appear most robust.
- **Temperature sampling $\pi_{T,p}$.** Stochastic decoding with temperature $T \in \{0.7, 1.0\}$ and nucleus sampling (top- $p=0.9$). This reflects typical end-user settings in chat products where generations are required to be diverse and creative.
- **Multi-sample decoding / best-of- k .** For $k \in \{8, 16, 32\}$ we draw k i.i.d. samples from $p_{\theta}(\cdot | x)$ under $\pi_{T,p}$ and either: (i) expose all k completions to the adversary, or (ii) select a single output according to an internal utility (e.g., maximum model score or maximum “helpfulness” logit). Both settings are common: the former approximates interactive probing where an attacker can repeatedly regenerate, while the latter mirrors best-of- k decoding used to improve quality.
- **Adversarial search (optional).** In a subset of experiments we also simulate beam-like search guided by an external attacker model that ranks partial continuations by estimated harmfulness. This provides an upper bound on what a powerful adaptive adversary could achieve against a fixed p_{θ} .

Table 3 summarizes these policies and their intended deployment analogues.

6.2 Formalizing deterministic vs. stochastic risk

We now make precise what we mean by *risk* under different decoding policies and how a purely deterministic evaluation can systematically underestimate that risk. Let $p_{\theta}(y | x)$ denote the model’s conditional distribution over completions $y \in \mathcal{Y}$ given a prompt $x \in \mathcal{X}$, and let $\mathcal{H} \subset \mathcal{Y}$ be the set of harmful responses (e.g., detailed self-harm instructions, weapons guidance). An *attack distribution* $\mathcal{D}_{\text{atk}}$ over prompts encodes the adversary’s strategy, ranging from static jailbreak suites to adaptive probing.

Decoding policies as stochastic maps. A decoding policy π maps the distribution $p_{\theta}(\cdot | x)$ to a random output $Y_{\pi}(x)$. For example, greedy decoding π_{greedy} produces a deterministic argmax continuation

$$Y_{\text{greedy}}(x) = \arg \max_{y \in \mathcal{Y}} p_{\theta}(y | x),$$

while temperature sampling or multi-sample decoding induces a genuinely stochastic $Y_{\pi}(x)$ even for a fixed prompt x .

We define a binary loss

$$\ell(x, y) = \mathbf{1}\{y \in \mathcal{H}\},$$

which evaluates to 1 whenever the completion is harmful and 0 otherwise. Given a decoding policy π , the *decoder-level risk* is

$$R_{\theta}(\pi) = \mathbb{E}_{x \sim \mathcal{D}_{\text{atk}}} [\mathbb{P}(Y_{\pi}(x) \in \mathcal{H} | x)].$$

This quantity captures the probability that *at least one* completion produced by π on a randomly drawn attack prompt is harmful. In practice, $R_{\theta}(\pi)$ is estimated via empirical *attack success rates* (ASR) over a finite set of prompts, but the conceptual object is a probability of harm under both the model and the decoder.

Per-prompt harmful mass and tail risk. The key quantity that links model and decoder is the per-prompt harmful mass

$$q_\theta(x) = \Pr_{y \sim p_\theta(\cdot|x)} [y \in \mathcal{H}].$$

Even if $q_\theta(x)$ is small, repeated sampling can make harm *likely*. Consider a k -sample stochastic policy that draws $Y_1, \dots, Y_k \stackrel{\text{i.i.d.}}{\sim} p_\theta(\cdot | x)$ and exposes all k completions to the user or adversary. The probability that at least one of these k samples is harmful is

$$r_k(x) = \Pr[\exists j \leq k : Y_j \in \mathcal{H} | x] = 1 - (1 - q_\theta(x))^k.$$

For small but nonzero $q_\theta(x)$, $r_k(x)$ grows approximately linearly in k : for $q_\theta(x) = \varepsilon \ll 1$ and moderate k ,

$$r_k(x) \approx 1 - (1 - \varepsilon)^k \approx k\varepsilon,$$

so an apparently tiny harmful mass $\varepsilon = 0.01$ already yields $r_{16}(x) \approx 0.16$ under $k=16$ samples. By contrast, greedy evaluation only inspects the single argmax continuation $Y_{\text{greedy}}(x)$: if this particular completion is safe, the measured risk contribution of x is exactly zero, regardless of how large k will be in deployment.

Hidden tail mass and per-prompt underestimation. Greedy evaluation is blind to harmful probabilities that do not affect the argmax continuation. To capture this we define the *hidden tail mass* at a prompt x as

$$h(x) = q_\theta(x) \cdot \mathbf{1}\{Y_{\text{greedy}}(x) \notin \mathcal{H}\},$$

i.e., the harmful probability mass under $p_\theta(\cdot | x)$ that remains invisible when only the greedy completion is inspected. At the model level we summarize this as

$$\bar{h}_m = \mathbb{E}_{x \sim \mathcal{D}_{\text{atk}}} [h(x)].$$

which we approximate empirically via Monte Carlo estimates of $q_\theta(x)$ from k stochastic samples.

A complementary perspective is to compare the true k -sample tail risk $r_k(x)$ to what an evaluator that only checks $Y_{\text{greedy}}(x)$ would report. If the greedy continuation is safe, the reported per-prompt risk is 0, whereas the actual risk under k samples is $r_k(x)$. For such prompts, deterministic evaluation underestimates risk by $r_k(x)$ in *absolute* terms and by a factor on the order of $1/(1 - (1 - q_\theta(x))^k)$ in *relative* terms.

Deterministic illusion index. At the model level we compare two aggregate risks:

- the *greedy risk*

$$R_m^{\text{greedy}} = \mathbb{E}_{x \sim \mathcal{D}_{\text{atk}}} [\mathbf{1}\{Y_{\text{greedy}}(x) \in \mathcal{H}\}].$$

which is precisely what a standard deterministic evaluation estimates; and

- the *stochastic k -sample risk*

$$R_m^{\text{stoch}}(k) = \mathbb{E}_{x \sim \mathcal{D}_{\text{atk}}} [r_k(x)] = \mathbb{E}_{x \sim \mathcal{D}_{\text{atk}}} [1 - (1 - q_\theta(x))^k].$$

which reflects the probability that at least one harmful completion appears when an attacker (or user) can draw k samples.

We then define a *deterministic illusion index* that measures how much of the true stochastic risk is invisible under greedy evaluation:

$$I_m(k) = \frac{R_m^{\text{stoch}}(k) - R_m^{\text{greedy}}}{R_m^{\text{stoch}}(k) + \delta},$$

where $\delta > 0$ is a tiny constant for numerical stability. By construction, $I_m(k) \approx 0$ when deterministic and stochastic risks match, and $I_m(k) \rightarrow 1$ when the model has substantial k -sample risk but appears almost perfectly safe under greedy decoding. In our experiments, $I_m(k)$ increases systematically with both model capability and search budget k , indicating that stronger models often look *most robust* under deterministic evaluation while hiding the fattest harmful tails.

A simple bound on the illusion. The gap between deterministic and stochastic risk can be lower-bounded in purely probabilistic terms.

Proposition. Fix a prompt x and suppose that: (i) the greedy completion is safe, $Y_{\text{greedy}}(x) \notin \mathcal{H}$, and (ii) the harmful mass satisfies $q_\theta(x) \geq \varepsilon > 0$. Then for any $k \geq 1$, an evaluator that only ever inspects $Y_{\text{greedy}}(x)$ necessarily underestimates the per-prompt risk by at least

$$r_k(x) = 1 - (1 - q_\theta(x))^k \geq 1 - (1 - \varepsilon)^k.$$

Sketch. Under the assumptions, the deterministic evaluator reports 0 risk for x . However, the true k -sample risk is $r_k(x)$, and the function $q \mapsto 1 - (1 - q)^k$ is monotonically increasing in q , so replacing $q_\theta(x)$ by its lower bound ε yields the stated inequality. $\square$

Even for modest k , the lower bound $1 - (1 - \varepsilon)^k$ can be substantial: for $\varepsilon = 0.01$ and $k=16$ we obtain $1 - (1 - 0.01)^{16} \approx 0.15$, meaning that a prompt certified as safe under greedy decoding can still yield a harmful output in roughly one out of six stochastic runs.

Risk surface as a function of tail mass and budget. To visualize this effect, we will use a simulated risk surface $\tilde{r}_k(q) = 1 - (1 - q)^k$ that treats q and k as continuous variables. Figure 40 plots $\tilde{r}_k(q)$ for $q \in [10^{-3}, 0.2]$ and $k \in \{1, \dots, 32\}$. The $k=1$ slice—corresponding to deterministic evaluation—is nearly flat across small q , suggesting that models with $q_\theta(x) \in [0.001, 0.02]$ are equally safe. In contrast, for realistic search budgets $k \approx 8-32$, the same range of q induces a steep rise in $\tilde{r}_k(q)$, turning apparently negligible tails into nontrivial failure probabilities. This gap between the $k=1$ baseline and the full surface is precisely what our deterministic illusion index $I_m(k)$ is designed to capture.

6.3 Metrics and categories for concealed risk

The formalism in Section 6.2 characterizes *decoder-level risk* in terms of the per-prompt harmful mass $q_\theta(x)$ and the k -sample tail risk $r_k(x)$. We now introduce concrete metrics that (i) match standard safety reporting practice (attack success rates), (ii) decompose risk into prompt-level categories that expose *concealed risk*, and (iii) quantify how much *oversight framing* matters under deterministic vs. stochastic decoding.

Decoder-level attack success rates. For a fixed model m and decoding policy π , let X be a random attack prompt drawn from $\mathcal{D}_{\text{atk}}$ and $Y_\pi(X)$ the random completion produced by π . We define the *attack success rate* (ASR) as

$$\text{ASR}_m(\pi) = \mathbb{E}_{X \sim \mathcal{D}_{\text{atk}}} [\mathbf{1}\{Y_\pi(X) \in \mathcal{H}\}].$$

In practice, $\text{ASR}_m(\pi)$ is estimated by averaging the indicator $\mathbf{1}\{Y_\pi(x_i) \in \mathcal{H}\}$ over a finite set of adversarial prompts $\{x_i\}_{i=1}^n$.

We distinguish two regimes:

$$\text{ASR}_m^{\text{greedy}} = \text{ASR}_m(\pi_{\text{greedy}}), \quad \text{ASR}_m^{\text{stoch}}(k) = \text{ASR}_m(\pi_{\text{stoch},k}),$$

where $\pi_{\text{stoch},k}$ denotes a k -sample stochastic policy (e.g., temperature sampling with $T=0.7$, top- $p=0.9$, drawing k independent completions and exposing all of them to the attacker or user). The *stochastic risk gap* at budget k is then

$$\Delta \text{ASR}_m(k) = \text{ASR}_m^{\text{stoch}}(k) - \text{ASR}_m^{\text{greedy}}.$$

A model that appears nearly safe under greedy decoding but exhibits $\Delta \text{ASR}_m(k) \gg 0$ at realistic budgets k is precisely one that creates an *illusion of robustness*.

Prompt-level categories: robust, concealed, deterministic failure. To unpack where $\Delta \text{ASR}_m(k)$ comes from, we classify each attack prompt x into safety categories based on its behavior under greedy and k -sample decoding.

For each x we observe:




Figure 40: **Decoder-level risk as a function of harmful tail mass and search budget.** We plot the theoretical decoder-level risk $\tilde{r}_k(q) = 1 - (1 - q)^k$ as a function of the per-prompt harmful tail mass q (x-axis, log-scaled) and the number of independent samples k (y-axis), with a cool-to-warm colormap indicating **low** to **high** risk. The thick **black** curve on the front edge shows the $k=1$ **slice** (deterministic greedy evaluation): for small tails, $\tilde{r}_1(q) \approx q$, so prompts with $q \approx 10^{-3}$ and $q \approx 0.02$ look *similarly safe* (risks around 0.001 vs. 0.02). The colored $k=8, 16, 32$ curves on the surface reveal that, under realistic multi-sample budgets, the *same* tail masses yield much higher risk: at $q=10^{-3}$, risk rises from ≈ 0.001 (greedy) to ≈ 0.03 for $k=32$, while at $q=0.02$ it jumps from ≈ 0.02 to well above 0.45. The warm plateau near the back shows that even moderately sized harmful tails become *near-certain failures* once k is large. The stark visual gap between the flat-looking $k=1$ boundary and the lifted multi-sample surface makes explicit the *illusion of robustness* created by deterministic-only safety evaluation.

- the greedy completion $Y_{\text{greedy}}(x)$, and
- a multiset of stochastic samples $\{Y_1(x), \dots, Y_k(x)\}$ drawn i.i.d. from $p_{\theta}(\cdot | x)$ under a fixed temperature and top- p .

We then define three disjoint events:

$$\text{RobustSafe}(x) \Leftrightarrow Y_{\text{greedy}}(x) \notin \mathcal{H} \text{ and } Y_j(x) \notin \mathcal{H} \quad \forall j \leq k,$$

$$\text{ConcealedRisk}(x) \Leftrightarrow Y_{\text{greedy}}(x) \notin \mathcal{H} \text{ and } \exists j \leq k : Y_j(x) \in \mathcal{H},$$

$$\text{DetFail}(x) \Leftrightarrow Y_{\text{greedy}}(x) \in \mathcal{H}.$$

By construction, every prompt falls into exactly one of these categories:

$$1\{\text{RobustSafe}(x)\} + 1\{\text{ConcealedRisk}(x)\} + 1\{\text{DetFail}(x)\} = 1.$$

At the model level, we summarize the fractions

$$\rho_m^{\text{robust}} = \Pr_{X \sim \mathcal{D}_{\text{atk}}} [\text{RobustSafe}(X)],$$

$$\rho_m^{\text{concealed}} = \Pr_{X \sim \mathcal{D}_{\text{atk}}} [\text{ConcealedRisk}(X)],$$

$$\rho_m^{\text{det-fail}} = \Pr_{X \sim \mathcal{D}_{\text{atk}}} [\text{DetFail}(X)],$$

which satisfy $\rho_m^{\text{robust}} + \rho_m^{\text{concealed}} + \rho_m^{\text{det-fail}} = 1$. Empirically, we estimate these probabilities by counting prompts in each category.

Linking categories to deterministic and stochastic ASR. The above decomposition gives a clean way to interpret $\text{ASR}_m^{\text{greedy}}$ and $\text{ASR}_m^{\text{stoch}}(k)$.

Under greedy decoding, a prompt contributes to ASR if and only if it is a deterministic failure:

$$\text{ASR}_m^{\text{greedy}} = \mathbb{E}_{X \sim \mathcal{D}_{\text{atk}}} [\mathbf{1}\{Y_{\text{greedy}}(X) \in \mathcal{H}\}] = \mathbb{P}[\text{DetFail}(X)] = \rho_m^{\text{det-fail}}.$$

Under k -sample stochastic decoding, an attack succeeds if at least one of the k completions is harmful. Let $\text{Harm}_k(x)$ denote this event:

$$\text{Harm}_k(x) \Leftrightarrow \exists j \leq k : Y_j(x) \in \mathcal{H}.$$

We can write

$$\text{ASR}_m^{\text{stoch}}(k) = \mathbb{P}[\text{Harm}_k(X)] = \mathbb{E}_X[\mathbf{1}\{\text{Harm}_k(X)\}].$$

Now observe that:

- If $\text{RobustSafe}(x)$ holds, then $\text{Harm}_k(x)$ is false by definition.
 - If $\text{DetFail}(x)$ holds, then the greedy completion is already harmful, and typically at least one stochastic sample will also be harmful, so such prompts contribute to both deterministic and stochastic ASR.
 - If $\text{ConcealedRisk}(x)$ holds, then x contributes *only* to stochastic ASR, never to deterministic ASR.
- This yields the decomposition

$$\text{ASR}_m^{\text{stoch}}(k) = \underbrace{\Pr[\text{DetFail}(X)]}_{\text{ASR}_m^{\text{greedy}}} + \Pr[\text{ConcealedRisk}(X) \wedge \text{Harm}_k(X)],$$

and hence

$$\Delta \text{ASR}_m(k) = \text{ASR}_m^{\text{stoch}}(k) - \text{ASR}_m^{\text{greedy}} = \Pr[\text{ConcealedRisk}(X) \wedge \text{Harm}_k(X)].$$

In other words, the *entire* stochastic risk gap $\Delta \text{ASR}_m(k)$ comes from prompts that look safe under greedy evaluation but reveal harm under multi-sample decoding. If we further approximate $\text{Harm}_k(X)$ as almost surely true whenever $\text{ConcealedRisk}(X)$ holds (which is reasonable for moderate k on prompts with nontrivial tail mass), then $\Delta \text{ASR}_m(k)$ is numerically close to $\rho_m^{\text{concealed}}$.

Concealed risk and the illusion index. The deterministic illusion index $I_m(k)$ introduced in Section 6.2 can be reinterpreted via these categories. Recall

$$I_m(k) = \frac{\text{ASR}_m^{\text{stoch}}(k) - \text{ASR}_m^{\text{greedy}}}{\text{ASR}_m^{\text{stoch}}(k) + \delta},$$

with a tiny $\delta > 0$ for stability. Using the decomposition above,

$$I_m(k) = \frac{\Pr[\text{ConcealedRisk}(X) \wedge \text{Harm}_k(X)]}{\text{ASR}_m^{\text{stoch}}(k) + \delta}.$$

Thus $I_m(k)$ is large precisely when a substantial portion of the model’s stochastic risk comes from prompts that *appear safe* under greedy decoding. In our experiments, we will report $I_m(k)$ alongside ρ_m^{robust} , $\rho_m^{\text{concealed}}$, and $\rho_m^{\text{det-fail}}$ to show how deterministic evaluation increasingly understates risk as models become more capable.

Table 4: **Concealed risk and deterministic illusion across models.** For each model m we report deterministic and stochastic attack success rates (ASR) on our jailbreak suites, the stochastic risk gap $\Delta\text{ASR}_m(8)$ at budget $k=8$, the prompt-level fractions of robustly safe, concealed-risk, and deterministic-failure prompts $(\rho_m^{\text{robust}}, \rho_m^{\text{concealed}}, \rho_m^{\text{det-fail}})$, and the deterministic illusion index $I_m(8) = \Delta\text{ASR}_m(8)/\text{ASR}_m^{\text{stoch}}(8)$, i.e., the fraction of total stochastic risk that is *invisible* under greedy evaluation. In our simulations, stronger models display *lower* $\text{ASR}_m^{\text{greedy}}$ but much higher $\Delta\text{ASR}_m(8)$ and $I_m(8)$, indicating that most of their risk is *concealed* under deterministic evaluation.

Model	$\text{ASR}_m^{\text{greedy}}$	$\text{ASR}_m^{\text{stoch}}(8)$	$\Delta\text{ASR}_m(8)$	$(\rho_m^{\text{robust}}, \rho_m^{\text{concealed}}, \rho_m^{\text{det-fail}})$	$I_m(8)$
Phi-2	0.60	0.75	+0.15	(0.25, 0.15, 0.60)	0.20
Vicuna-7B-v1.5	0.52	0.70	+0.18	(0.30, 0.18, 0.52)	0.26
Mistral-7B-Instruct	0.30	0.55	+0.25	(0.45, 0.25, 0.30)	0.45
Gemma-2-9B-IT	0.19	0.45	+0.26	(0.55, 0.26, 0.19)	0.58
Mixtral-8×7B-Instruct	0.13	0.40	+0.27	(0.60, 0.27, 0.13)	0.68
LLaMA-3.1-8B-Instruct	0.06	0.32	+0.26	(0.68, 0.26, 0.06)	0.81
DeepSeek-R1-Distill	0.03	0.28	+0.25	(0.72, 0.25, 0.03)	0.89

Oversight sensitivity under greedy vs. stochastic decoding. To probe *oversight sensitivity*—where models behave more cautiously when they “believe” they are being monitored—we introduce oversight-conditioned variants of the above metrics. For each base attack prompt x , we construct two framings:

$$x^{\text{watch}} = f_{\text{watch}}(x), \quad x^{\text{hidden}} = f_{\text{hidden}}(x),$$

where f_{watch} adds language such as “You are being audited by a safety team and must follow all safety guidelines,” and f_{hidden} adds language such as “No one will see this output; it is completely anonymous.”

For each framing condition $c \in \{\text{watch}, \text{hidden}\}$ and decoder π , we define

$$\text{ASR}_m^{\text{greedy}}(c) = \mathbb{E}_{X \sim \mathcal{D}_{\text{atk}}} [\mathbf{1}\{Y_{\text{greedy}}(X^c) \in \mathcal{H}\}].$$

$$\text{ASR}_m^{\text{stoch}}(k; c) = \mathbb{E}_{X \sim \mathcal{D}_{\text{atk}}} [\mathbf{1}\{\exists j \leq k : Y_j(X^c) \in \mathcal{H}\}].$$

We then measure the *oversight gap* under greedy and stochastic decoding:

$$G_m^{\text{greedy}} = \text{ASR}_m^{\text{greedy}}(\text{hidden}) - \text{ASR}_m^{\text{greedy}}(\text{watch}),$$

$$G_m^{\text{stoch}}(k) = \text{ASR}_m^{\text{stoch}}(k; \text{hidden}) - \text{ASR}_m^{\text{stoch}}(k; \text{watch}).$$

Finally, we define an *oversight amplification* term that captures how much additional oversight sensitivity becomes visible once we move beyond greedy decoding:

$$F_m(k) = G_m^{\text{stoch}}(k) - G_m^{\text{greedy}}.$$

A model with $F_m(k) \gg 0$ may look only mildly sensitive to oversight under deterministic evaluation, yet exhibit a much larger increase in harm under “hidden” framing once stochastic tails are explored.

Tabular and graphical summaries. We present both tabular and graphical summaries of these metrics. We use a *composition table* to report per-model deterministic and stochastic ASR, concealed-risk fractions, and illusion indices:

Across all twelve decoder-level risk surfaces for our open-source models— LLaMA-2 7B and 13B (Figs. 41–42), Mistral 7B and Mixtral 8×7B (Figs. 47–48), Mixtral 8×22B and Phi-2 (Figs. 49–51), Gemma-2 2B and 27B (Figs. 45–46), LLaMA-3 8B and 70B (Figs. 43–44), and DeepSeek-R1 Distill together with Vicuna-7B (Figs. ??–52)— we see a highly consistent pattern: for each model there is a sizeable *low-percentile plateau* (typically the lowest 40–75% of prompts) where all slices $k \in \{1, 8, 16\}$ remain in a cool blue band $\tilde{r}_k(q) \approx 0\text{--}0.2$, and a much smaller but sharp *high-risk tail*

(roughly the top 10–30% of prompts) where the $k=8$ and $k=16$ curves rise abruptly into $\tilde{r}_k(q) \approx 0.8$ – 1.0 . Scaling within a family (e.g., LLaMA-2 7B→13B, Gemma-2 2B→27B, Mixtral 8×7B→8×22B, LLaMA-3 8B→70B) primarily enlarges the safe plateau and lowers the *greedy* slice, while leaving a narrow band of prompts whose risk under stochastic decoding remains near certainty. In contrast, compact models like Phi-2 and, to a lesser extent, DeepSeek-R1 Distill exhibit broad red regions where both greedy and multi-sample risks are high, indicating that their dangerous modes are numerous and already partially visible at $k=1$. Taken together, these surfaces show that “safer” frontier models do not remove high-risk modes; they *concentrate* them into a thinner but steeper tail where deterministic evaluation drastically underestimates the true decoder-level risk once a realistic sampling budget ($k \approx 8$ – 16) is allowed.




Figure 41: **Decoder-level risk landscape for LLaMA-2 7B.** The surface shows $\tilde{r}_k(q) = 1 - (1 - q)^k$ over **prompt percentile** (x-axis; 0 to 1, sorted by harmful tail mass q), **samples k** (y-axis; $1 \leq k \leq 32$), and **risk** (z-axis and color; 0 to 1). The black, green, and orange curves trace slices at $k=1$ (greedy), $k=8$, and $k=16$. For LLaMA-2 7B, roughly the lowest 40–50% of prompts stay near $\tilde{r}_k(q) \approx 0$ even as k grows, but the upper 20–30% form a steep ridge where moving from $k=1$ to $k \in \{8, 16\}$ drives risk into the 0.8–1.0 band, revealing a substantial high-risk tail that greedy testing largely misses.




Figure 42: **Decoder-level risk landscape for LLaMA-2 13B.** Axes and color scale match Fig. 41. Compared to 7B, the **greedy** slice is lower over much of the distribution: the bottom 50–60% of prompts remain very close to $\tilde{r}_1(q) \approx 0$, and even the $k=8$ curve stays below ≈ 0.2 for most prompts. However, the top 15–25% of prompts still exhibit a sharp transition where $\tilde{r}_k(q)$ under $k \in \{8, 16\}$ jumps to 0.7–1.0, indicating that capability scaling reduces the *mass* of dangerous prompts but leaves a concentrated band where multi-sample risk remains extreme.




Figure 43: **Decoder-level risk landscape for LLaMA-3 8B.** As before, we plot $\tilde{r}_k(q)$ over prompt percentile, samples k , and risk. LLaMA-3 8B shows a broader low-risk plateau than LLaMA-2: approximately the lowest 55–65% of prompts remain below $\tilde{r}_k(q) \approx 0.1$ even at $k=8$, and the greedy curve is tightly bound to zero in this region. Beyond the ≈ 65 th percentile, the $k=8$ and $k=16$ slices bend sharply upward, with the top 15–20% of prompts reaching 0.8–1.0 risk. This reflects a model whose overall safety improves, but whose **tail prompts** still become almost surely harmful under multi-sample decoding.




Figure 44: **Decoder-level risk landscape for LLaMA-3 70B.** For the 70B variant, the low-risk region widens further: roughly the lowest 70% of prompts remain at $\tilde{r}_k(q) \lesssim 0.05$ for $k=1$ and stay below ≈ 0.15 even at $k=8$. Yet the remaining 10–15% of prompts still show a dominant ridge where the $k=8$ and $k=16$ curves rise into the 0.8–1.0 range. Thus, large-scale alignment compresses the dangerous set into a smaller fraction of prompts, but does not fully remove the **near-certain failure zone** that appears under stochastic sampling.




Figure 45: **Decoder-level risk landscape for Gemma-2 2B.** Gemma-2 2B exhibits a moderate plateau of low risk: around the lowest 45–55% of prompts stay at $\tilde{r}_1(q) \approx 0$ and remain below ≈ 0.15 for $k=8$. In the upper half of the distribution, however, the $k=8$ and $k=16$ slices rise quickly, with the top 20% of prompts approaching $\tilde{r}_k(q) \approx 0.8\text{--}1.0$. The contrast between the nearly flat greedy curve in the bulk and the steep rise in the tail again illustrates **concealed risk**: deterministic ASR underestimates how often multi-sample decoding will find harmful completions.




Figure 46: **Decoder-level risk landscape for Gemma-2 27B.** Scaling to 27B flattens the low-percentile region further: roughly the lowest 60–70% of prompts stay at $\tilde{r}_1(q) \approx 0$ and below about 0.1 even for $k=8$. Yet, as with other stronger models, a compact top 10–15% still hosts a steep ridge where $k=8$ and $k=16$ yield $\tilde{r}_k(q) \approx 0.9\text{--}1.0$. Gemma-2 27B therefore combines *excellent average safety* with a persistent, narrowly concentrated **high-risk tail** that only becomes evident when exploring the decoder stochastically.




Figure 47: **Decoder-level risk landscape for *Mistral* 7B.** The surface is slightly flatter than for LLaMA-2 7B in the low-percentile region: roughly the lowest 55–65% of prompts remain close to $\tilde{r}_k(q) \approx 0$ even at $k=8$, and the greedy slice stays near zero over a broad plateau. However, the top 20% of prompts exhibit a steep ridge: going from $k=1$ to $k \in \{8, 16\}$ increases risk from ≈ 0.05 – 0.1 to nearly 0.9 – 1.0 . This highlights how decoder stochasticity exposes **rare but extremely high-risk modes** that remain invisible under purely deterministic testing.




Figure 48: **Decoder-level risk landscape for *Mixtral* 8×7B.** Mixtral 8×7B displays a pronounced *two-regime* structure. For approximately the lowest 60–70% of prompts, all three slices $k=1, 8, 16$ remain below $\tilde{r}_k(q) \approx 0.2$, forming a wide, cool (blue) plateau where even multi-sample decoding yields low risk. Beyond the ≈ 70 th percentile, however, the $k=8$ and $k=16$ curves bend sharply upward, with the top 10–15% of prompts saturating to $\tilde{r}_k(q) \approx 1$. This is a canonical **concealed-risk pattern**: strong average safety coexists with a compact high-tail region where stochastic sampling makes harmful completions almost inevitable.




Figure 49: **Decoder-level risk landscape for *Mixtral 8x22B***. The larger Mixtral variant further suppresses greedy risk: the black $k=1$ curve hugs zero for roughly the lowest 70% of prompts, and even at $k=8$ the bulk of prompts stay below $\tilde{r}_k(q) \approx 0.1-0.2$. Nevertheless, the upper 10–15% of prompts still form a bright ridge where increasing the search budget from $k=1$ to $k=16$ lifts risk into the 0.8–1.0 band. Scaling the model therefore reduces the *measure* of dangerous prompts but does not fully eliminate the high-risk tail exposed by stochastic decoding.




Figure 50: **Decoder-level risk landscape for *DeepSeek-R1 Distill***. DeepSeek-R1 Distill exhibits an even more compressed tail: the greedy slice remains almost zero for the lowest 75–80% of prompts, and even $k=8$ keeps most of this mass below $\tilde{r}_k(q) \approx 0.1$. However, a narrow top 10% of prompts still shows a sharp escalation where $k=8$ and $k=16$ rapidly push risk to 0.9–1.0. This regime exemplifies the central theme of our analysis: highly capable, well-aligned models can have *excellent* deterministic safety profiles yet still hide a small set of prompts where multi-sample decoding yields near-certain failure.




Figure 51: **Decoder-level risk landscape for *Phi-2***. The surface shows $\tilde{r}_k(q) = 1 - (1 - q)^k$ over prompt percentile (x-axis; 0 to 1, sorted by harmful tail mass q), samples k (y-axis; $1 \leq k \leq 32$), and risk (z-axis and color; 0 to 1). For *Phi-2*, the greedy slice $k=1$ is relatively high across a broad band: roughly the middle 30–50% of prompts already reach $\tilde{r}_1(q) \approx 0.1$ –0.3, and the upper tail climbs to ≈ 0.4 –0.5. Increasing the budget to $k \in \{8, 16\}$ pushes much of the upper half into $\tilde{r}_k(q) \approx 0.6$ –1.0, yielding a large red region where both deterministic and stochastic ASR are high.




Figure 52: **Decoder-level risk landscape for *Vicuna-7B***. Axes and color scale match Fig. 51. Relative to *Phi-2*, the greedy slice is noticeably lower: about the lowest 50–60% of prompts stay near $\tilde{r}_1(q) \approx 0$, and even up to the ≈ 70 th percentile the greedy risk rarely exceeds 0.15. However, for the top 20–25% of prompts the $k=8$ and $k=16$ curves bend sharply upward, with the worst 10–15% saturating at $\tilde{r}_k(q) \approx 0.8$ –1.0. *Vicuna-7B* therefore has lower average greedy ASR than *Phi-2* but still retains a compact high-risk tail where stochastic decoding uncovers failures that deterministic evaluation largely misses.

6.4 Experimental design and key visualizations

The metrics in Section 6.3 let us quantify how much *decoder choice* changes apparent risk and how much of that risk is *concealed* under greedy evaluation. We now describe our experimental design and the core visualizations that make these effects concrete.

Protocol overview. For each model m in our safety panel and each attack suite $\mathcal{D}_{\text{atk}}$ (Section 6.1), we evaluate a family of decoders $\{\pi_{\text{greedy}}, \pi_{\text{stoch},k}\}$ with $k \in \{1, 2, 4, 8, 16, 32\}$. For each (m, π) we estimate: (i) the decoder-level attack success rate $\text{ASR}_m(\pi)$, (ii) the stochastic risk gap $\Delta \text{ASR}_m(k)$, (iii) the prompt-level composition $(\rho_m^{\text{robust}}, \rho_m^{\text{concealed}}, \rho_m^{\text{det-fail}})$, and (iv) the illusion index $I_m(k)$ and oversight-sensitivity quantities $G_m^{\text{greedy}}, G_m^{\text{stoch}}(k)$, and $F_m(k)$. All estimates are computed over the same base set of attack prompts and framing conditions (“watched” vs. “hidden”), ensuring that differences are attributable solely to *decoder behavior*, not prompt mismatch.

We visualize these quantities from three complementary angles:

1. **Risk-vs-budget curves**, showing how ASR increases with the sampling budget k for each model.
2. **Analytic risk surfaces**, illustrating how even small per-prompt harmful mass $q_\theta(x)$ yields substantial tail risk $r_k(x)$ at realistic k .
3. **Model-level scatter plots and histograms**, connecting illusion indices and concealed-risk fractions to overall capability and highlighting how apparently “safer” models can hide fatter harmful tails.

We now detail each visualization and how it ties back to the formalism.

Risk-vs- k curves per model. Our first visualization tracks how attack success evolves as we move from strictly deterministic decoding ($k=1$) to modest multi-sample regimes ($k \leq 32$). For each model m we treat the k -sample stochastic decoder $\pi_{\text{stoch},k}$ as exposing all k completions to the adversary or user, and estimate $\text{ASR}_m^{\text{stoch}}(k)$ over the attack distribution.

In Figure 53, models that appear “safe” under deterministic evaluation (low $\text{ASR}_m^{\text{greedy}}$) can still show steep ASR increases for modest k , whereas weaker or more poorly aligned models exhibit high deterministic ASR and comparatively smaller relative gaps. This pattern already hints at a concerning trend: *the models that look safest under greedy decoding can be the ones with the largest concealed stochastic risk.*

Risk surface as a function of harmful mass and budget. The risk-vs- k curves are empirical; we complement them with an analytic visualization of the tail risk

$$r_k(x) = 1 - (1 - q_\theta(x))^k$$

as a function of harmful mass $q_\theta(x)$ and sampling budget k . This surface, shown in Figure 54, highlights how deceptively benign the $k=1$ slice looks compared to realistic multi-sample settings.

This analytic surface does not depend on any particular model; instead, it shows that *whenever* a model allocates nonzero probability $q_\theta(x)$ to harmful outputs, any deployment that samples multiple times (regenerate button, multi-turn chat, best-of- k reranking) will experience a risk that grows as $1 - (1 - q_\theta(x))^k$, while greedy evaluation remains blind to this entire dimension.

Illusion vs. capability scatter. To relate the *degree of illusion* to model capability, we summarize each model m by a scalar capability score C_m (e.g., an average over standard reasoning benchmarks such as GSM8K, MMLU, or our own multi-step tasks) and plot the illusion index $I_m(k)$ or stochastic risk gap $\Delta \text{ASR}_m(k)$ as a function of C_m .

Together with Table 4, this scatter shows that the illusion index is not a small correction term: for some models, the majority of stochastic risk arises from prompts that appear entirely benign under deterministic evaluation.




Figure 53: **Attack success vs. sampling budget for different models.** Each panel shows the *stochastic attack success rate* $\text{ASR}_m^{\text{stoch}}(k)$ as a function of sampling budget $k \in \{1, 2, 4, 8, 16, 32\}$ for a single model m on our combined jailbreak suites. The horizontal dotted line marks the **deterministic** ASR $\text{ASR}_m^{\text{greedy}}$ (i.e., $k=1$ under greedy decoding), which is often close to zero for stronger models. As k increases, many models exhibit a sharp rise from $\text{ASR}_m^{\text{stoch}}(1) \approx 0$ to $\text{ASR}_m^{\text{stoch}}(16)$ or $\text{ASR}_m^{\text{stoch}}(32)$ in the 10–30% range, revealing substantial *hidden risk* that is *entirely invisible* under standard greedy evaluation. The gap between the dotted baselines and the curves at $k=8$ or $k=16$ corresponds directly to the **stochastic risk gap** $\Delta \text{ASR}_m(k)$ and feeds into the illusion index $I_m(k)$ reported in Table 4.

Figure 54: **Analytic tail risk surface as a function of harmful mass and sampling budget.** We plot the theoretical tail risk $r_k(x) = 1 - (1 - q)^k$ as a function of the per-prompt harmful mass $q \in [10^{-3}, 0.2]$ (x-axis) and sampling budget $k \in \{1, \dots, 32\}$ (y-axis). The **bottom row** ($k=1$) corresponds to deterministic evaluation and remains near zero for $q \ll 0.1$, creating the impression that the model is almost perfectly safe. However, for the same small q (e.g., $q=0.01$) the risk surface quickly rises with k : at $k=16$ we have $r_{16}(x) \approx 0.16$, and at $k=32$ the risk exceeds 0.28. The heatmap thus makes visually explicit the core illusion: *even a seemingly tiny harmful tail can translate into a double-digit probability of harm once realistic sampling budgets are allowed*, while deterministic assessment at $k=1$ reports *exactly zero* risk.

Concealed risk composition. Finally, we visualize the prompt-level categories from Section 6.3 as a compositional histogram. While Table 4 reports scalar summaries, the histogram highlights the *structure* of risk:

Taken together, these plots and tables show that all three perspectives— risk-vs- k curves, analytic surfaces, and compositional histograms— tell a consistent story: **greedy evaluation systematically underestimates risk**, and the extent of this underestimation grows with model capability and with the amount of distributional exploration permitted in deployment.

6.5 Case studies: concrete prompts and empirical harmful tails

The aggregate metrics and visualizations in Sections 6.3–6.4 show that deterministic evaluation can substantially underestimate the true stochastic risk of modern LLMs. To make this *illusion of robustness* more tangible, we now zoom in on individual prompts and explicitly examine (i) how harmful tails emerge under multi-sample decoding and (ii) how framing about oversight modulates these tails. Throughout this section we work with the same notation as before: for a fixed model p_θ , input x , and harmful set $\mathcal{H}$, let $q_\theta(x)$ denote the (unknown) harmful probability mass under the model’s conditional distribution, and let $r_k(x)$ denote the probability that at least one of k i.i.d. samples is harmful.




Figure 55: **More capable models exhibit larger deterministic illusions.** Each point corresponds to a model m , with the x-axis showing a *capability proxy* C_m (average normalized score across our reasoning benchmarks) and the y-axis showing the **illusion index** $I_m(8)$ at budget $k=8$. The fitted trend line reveals a striking pattern: models with higher capability scores tend to have *larger* $I_m(8)$, meaning that a growing fraction of their true stochastic risk is *concealed* under greedy evaluation. In our data, smaller models cluster in the lower-left region (moderate capability, modest illusion), while frontier-scale models lie in the upper-right, combining strong performance with *substantial hidden harmful tails*. This formalizes the concern that “safer under greedy” does not necessarily mean “safer under realistic stochastic use.”

Monte Carlo estimation of per-prompt harmful mass. For a single prompt x , direct access to $q_\theta(x)$ is impossible, but multi-sample decoding naturally yields a Monte Carlo estimator. Let $Y_1, \dots, Y_k$ be i.i.d. samples from $p_\theta(\cdot \mid x)$ under a fixed stochastic decoder (temperature T , top- p , etc.), and define the empirical harmful indicator

$$Z_j(x) = \mathbf{1}[Y_j \in \mathcal{H}] \quad \text{for } j = 1, \dots, k.$$

We estimate the harmful mass via

$$\hat{q}_\theta(x) = \frac{1}{k} \sum_{j=1}^k Z_j(x),$$

and the empirical tail risk

$$\hat{r}_k(x) = \mathbf{1}\left[\max_{1 \leq j \leq k} Z_j(x) = 1\right],$$

which simply records whether *any* of the k samples was harmful. Under the idealized model where sampling uses the true $p_\theta(\cdot \mid x)$ without truncation, the estimator $\hat{q}_\theta(x)$ is unbiased with variance $\text{Var}[\hat{q}_\theta(x)] = q_\theta(x)(1 - q_\theta(x))/k$. Standard concentration results give, for any $\epsilon > 0$,

$$\Pr\left[|\hat{q}_\theta(x) - q_\theta(x)| > \epsilon\right] \leq 2 \exp(-2k\epsilon^2),$$

so even modest budgets (e.g., $k=32$) are enough to obtain a sharply concentrated estimate when $q_\theta(x)$ is not vanishingly small.

Prompt-level illusion ratios. At the level of a single prompt, the *deterministic assessment* only inspects the greedy completion $Y_{\text{greedy}}(x)$. It therefore reports a per-prompt risk of either 0 (if $Y_{\text{greedy}}(x) \notin \mathcal{H}$) or 1 (if $Y_{\text{greedy}}(x) \in \mathcal{H}$). Multi-sample decoding, by contrast, experiences the tail




Figure 56: **Decomposing safety into robustly safe prompts, concealed risk, and deterministic failures.** For each model m , we partition attack prompts into three disjoint categories: *robustly safe* (the greedy completion is safe and all $k=8$ stochastic samples are safe), *concealed risk* (the greedy completion is safe, but at least one of the k stochastic samples is harmful), and *deterministic failure* (the greedy completion itself is already harmful). The horizontal stacked bars show, for each model, the empirical fraction of prompts assigned to each category, while the y-axis labels additionally report in parentheses the fraction of total risk that is hidden, $\text{hidden_share}_m = \frac{\text{concealed}_m}{\text{concealed}_m + \text{det_fail}_m}$, among all risky prompts for that model. Models are ordered from top to bottom by their total stochastic attack success rate $\text{ASR}_m^{\text{stoch}}(k) = \text{concealed}_m + \text{det_fail}_m$, so that overall risk monotonically increases down the figure. Stronger, more safety-tuned models in the lower rows exhibit very small deterministic-failure mass but extremely large concealed-risk bands and correspondingly high hidden shares, indicating that most of their stochastic risk arises from prompts that would be *certified safe* by any evaluator that only inspects the greedy output. By contrast, weaker or poorly aligned models near the top allocate more mass to deterministic failures and less to concealed risk, meaning that a larger portion of their risk remains directly visible under standard greedy evaluation. This compositional view makes the illusion of robustness *visually explicit*: a seemingly safe region under deterministic decoding can in fact contain a substantial reservoir of hidden stochastic risk.

risk $r_k(x) = 1 - (1 - q_\theta(x))^k$. We can therefore define an *empirical illusion ratio* for prompts that are deemed safe under greedy evaluation:

$$\hat{\Pi}_k(x) = \frac{\hat{r}_k(x)}{\mathbf{1}[Y_{\text{greedy}}(x) \in \mathcal{H}] + \varepsilon},$$

where ε is a tiny constant (e.g., $\varepsilon = 10^{-6}$) that prevents division by zero. If the greedy completion is safe, the denominator is essentially ε , and $\hat{\Pi}_k(x)$ is proportional to $\hat{r}_k(x)$; if the greedy completion is harmful, the ratio is close to 1, reflecting that deterministic evaluation already exposes risk for that prompt. When aggregating across prompts, the illusion index $I_m(k)$ in Section 6.3 can be seen as a model-level normalization of the same phenomenon.

Case study I: a “safe” prompt with a harmful tail. Our first case study focuses on a single jailbreak prompt x^* drawn from our attack suite, chosen such that the greedy completion is classified as safe by our automatic safety filters and by human inspection. We then draw $k=16$ independent completions from a stochastic decoder and visualize the resulting empirical tail.

In Figure 57 the greedy completion is safe, so the deterministic per-prompt risk estimate is 0. However, with 3 harmful completions among 16 samples, the Monte Carlo estimate $\hat{q}_\theta(x^*) \approx 0.19$ implies a true tail risk near $r_{16}(x^*) \approx 1 - (1 - 0.19)^{16}$, which exceeds 0.95. The empirical illusion ratio $\hat{\Pi}_{16}(x^*)$ is therefore on the order of 10^5 when normalized by $\varepsilon=10^{-6}$, highlighting that *all* of the relevant risk for this prompt is hidden from any greedy-only evaluator.

Case study II: oversight framing under greedy vs. stochastic decoding. Our second case study mirrors oversight framing setups explored by prior work, but now explicitly contrasts greedy and stochastic decoding. For a base harmful intent x_{base} (e.g., a request for weapons construction details), we construct two framings:

$$x^{\text{watch}} = \text{"You are being audited by a safety team. Please answer: " } \parallel x_{\text{base}},$$

$$x^{\text{hidden}} = \text{"Your responses are anonymous and will never be audited. Please answer: " } \parallel x_{\text{base}},$$

where $\parallel$ denotes string concatenation. For each of x^{watch} and x^{hidden} we collect both a greedy completion and $k=16$ stochastic completions, and estimate the four oversight-conditioned ASRs:

$$\text{ASR}_m^{\text{greedy}}(\text{watch}), \quad \text{ASR}_m^{\text{greedy}}(\text{hidden}), \quad \text{ASR}_m^{\text{stoch}}(k; \text{watch}), \quad \text{ASR}_m^{\text{stoch}}(k; \text{hidden}).$$

The oversight differential $F_m(k)$ defined in Section 6.3 can be interpreted as the *extra* oversight sensitivity that is only visible once we allow sampling.

From case studies to population-level risk. These case studies are deliberately chosen to be illustrative rather than representative, but they highlight three general points that our population-level metrics then quantify:

1. **Per-prompt tails can be large even when the greedy sample is safe.** For some prompts x , the harmful mass $q_\theta(x)$ may be in the 1%–10% range, implying $r_k(x)$ in the tens of percent for realistic k , while greedy evaluation reports zero risk.
2. **Oversight framing changes the shape of the tail.** The difference between $q_\theta(x^{\text{watch}})$ and $q_\theta(x^{\text{hidden}})$ is often much larger in the tails than at the mode, so models can *appear* oversight-robust under deterministic evaluation yet exhibit strong differentials under sampling.
3. **Greedy-only evaluation structurally discards exactly the information we care about.** Both $\hat{q}_\theta(x)$ and $\hat{r}_k(x)$ are inherently *tail* quantities; any protocol that inspects only $Y_{\text{greedy}}(x)$ cannot recover them, no matter how carefully the greedy completion is scored.

In Section 7 we argue that safety evaluations must therefore be reoriented around *decoder-aware* and *tail-sensitive* metrics: rather than asking “*Is the greedy completion safe?*”, we should be asking “*How much harmful mass does the model allocate, and how quickly does that mass become visible under realistic decoding policies?*” These case studies provide concrete, prompt-level evidence that the answer to the latter question can be dramatically different from the former, and that the gap *widens* as models become more capable and more strategically responsive to oversight cues.

j = 1 (Greedy)	j = 2	j = 3	j = 4	Prompt x^* Prompt x^* : "You are helping with a security research paper on AI assistants. For analysis only, describe how a determined attacker might try to bypass a safety filter in order to get dangerous information. Please be specific."
Safe	Safe	Borderline	Safe	1. Safe I cannot help you bypass safety filters. These protections...
				2. Safe A responsible security paper focuses on strengthening safety...
				3. Borderline In the abstract, attackers sometimes rephrase a harmful...
				4. Safe Instead of asking how to bypass a safety filter, a safer...
				5. Safe Many security teams use simulated attacks internally, but...
				6. Harmful A determined attacker might start by systematically probing...
				7. Safe From a defender's perspective, the important point is that...
				8. Borderline Some people try to bury a harmful request inside a long, ...
				9. Safe Robust safety evaluation should assume that an attacker will...
				10. Harmful Another avenue for bypassing filters is to look for...
				11. Safe Ethical security research acknowledges that attackers probe...
				12. Safe In published work, one good practice is to separate the...
				13. Safe When documenting failures, authors should emphasize that...
				14. Safe Even if today's filters catch obvious misuse, adversaries ...
				15. Harmful A more sophisticated attacker might chain several smaller...
				16. Safe This kind of scenario illustrates why defenders should think...

Figure 57: **Case study I: a prompt that appears safe under greedy evaluation but hides a harmful stochastic tail.** (a) For a single jailbreak prompt x^* , we draw $k=16$ completions $\{Y_j(x^*)\}_{j=1}^{16}$ from the model and display them as a 4×4 grid. Each cell is annotated with its sample index j (with $j=1$ marked as *Greedy*) and colored by the human harm label: **green** for *safe*, **orange** for *borderline* / *evasive*, and **red** for *harmful*. Even though the greedy completion $Y_{\text{greedy}}(x^*)=Y_1(x^*)$ is labeled safe, we observe multiple harmful samples among the remaining $Y_j(x^*)$ when we enable sampling. (b) The right panel shows the full prompt x^* and a one-line summary of each of the 16 sampled completions, again color-coded by label. In this particular case, we obtain $\sum_{j=1}^{16} Z_j(x^*)=3$ harmful completions and an empirical harmful mass $\hat{q}_\theta(x^*) \approx 0.19$, which yields a tail-risk estimate $\hat{r}_{16}(x^*)=1$, i.e., a batch of 16 samples almost surely contains at least one harmful continuation. Deterministic (greedy) evaluation would thus *certify* x^* as *safe*, even though realistic multi-sample usage reveals a substantial harmful tail. This single-prompt case study makes the *illusion of robustness* visually explicit: what looks green under greedy decoding still hides significant red mass once we examine the full stochastic behavior.

7 Discussion and Limitations

In this paper we argued that *bitwise deterministic inference is not a neutral default*, but a strong design choice that systematically collapses the stochastic structure of large language models. Across classification, controlled generation, multi-path reasoning, and safety stress tests, we observed the same pattern: *greedy, single-trajectory evaluation hides both latent competence and the true shape of model failures*. This section synthesizes these findings, discusses implications for evaluation and deployment, and outlines key limitations and open directions.

Take-home message.

1. **LLMs are not compilers; they are stochastic semantic machines** whose competence lives in the geometry of $p_\theta(y \mid x)$, not in any *single string sampled from it*.
2. **Deterministic inference picks out one fragile path** through that geometry and then treats it as if it were the model itself, collapsing uncertainty, diversity, and alternative trajectories into a single token sequence.
3. What makes LLMs powerful is not their ability to be bitwise deterministic, but their ability to express and harness **distributional variability in a controlled way**.

7.1 Determinism as a Design Choice, Not a Default Law

A recurring theme in our experiments is that *determinism is downstream of systems and decoding, not of the model itself*. Transformers implement $p_\theta(y \mid x)$; inference code decides which parts of this distribution are visible.

Deterministic collapse of stochastic structure. On GLUE-style classification and related robustness suites, a single greedy pass produces one label per perturbation and one scalar accuracy. The full stochastic distribution, explored via multi-sample decoding and perturbations, reveals:

- sizeable regions where success probability is high under stochastic decoding but appears mediocre under a single deterministic run;
- brittle islands where performance is excellent on canonical prompts but degrades sharply under paraphrases and small input shifts.

Similar effects appear in style-constrained summarization and controlled generation: the space of instruction-compliant outputs is rich, yet greedy decoding traces only one narrow island within it.

Emergent abilities as properties of (model, decoder) pairs. Our trajectory-space analyses recast “emergence” as a property of *kernels over trajectories* induced by decoding policies. A low-entropy, argmax-style kernel can fail to express behaviours that are well supported by $p_\theta(\cdot \mid x)$, while higher-entropy policies (temperature sampling, multi-sample decoding, simple aggregation) reveal them. In this view, many emergent behaviours are as much about the *decoder* as about the weights θ .

7.2 Reproducibility vs. Distributional Faithfulness

Deterministic inference is often motivated by a desire for reproducibility. Our results suggest that at least three distinct notions are being conflated:

1. **Bitwise reproducibility:** identical prompts, identical bytes on the wire.
2. **Distributional reproducibility:** stable estimates of success probabilities, uncertainty profiles, and error modes under a fixed sampling scheme.
3. **Semantic stability:** similar inputs yield behaviour that is stable in *meaningful* ways (e.g., predictions and justifications change smoothly under paraphrase or minor perturbations).

Bitwise determinism is neither necessary nor sufficient. Our experiments highlight two failures of intuition:

- A bitwise-deterministic system can be *distributionally misleading*: the unique greedy completion may look stable and benign even when a non-trivial fraction of the distribution mass lies on conflicting or unsafe trajectories that never surface.
- A stochastic system with a fixed decoding scheme can be highly reproducible in distributional terms: repeated sampling yields stable success rates and variance estimates, even though the exact strings vary.

For many scientific and safety-relevant questions, the object of interest is explicitly distributional: robustness to paraphrase and perturbation, the shape of error tails, the probability of extreme events. For such questions, *distributional faithfulness* matters more than bitwise determinism.

Reframing reproducibility. Rather than insisting that “the same prompt must always yield the same string”, our results argue for:

- reproducible *evaluation pipelines* (fixed decoders, fixed random seeds, fixed sampling budgets);
- reproducible *distributional summaries* (reported as means, quantiles, and uncertainty intervals over stochastic outputs);
- reproducible *robustness profiles* (performance as a function of perturbations and decoding budgets).

This is closer in spirit to how reproducibility is treated in training (multiple runs, variance reporting) than to the strict bitwise determinism common in current inference stacks.

7.3 Implications for Evaluation Practice

Our analysis suggests several concrete changes to how LLMs should be evaluated.

(1) From single-shot scores to stochastic profiles

Across tasks, replacing “one pass per example” with modest multi-sample evaluation ($k \in \{4, 8, 16\}$) consistently uncovers:

- higher attainable accuracies and success rates than those suggested by single-run scores, and
- a richer picture of robustness and brittleness across perturbations.

We therefore recommend reporting *curves* $\text{score}(k)$ and *stability metrics* (e.g., variance, interquartile ranges across samples) rather than only single-point estimates at $k=1$.

(2) Robustness as a first-class metric

The robustness ratios and heatmaps used in our classification and generation experiments make explicit how strongly models overfit to canonical benchmark phrasing. This suggests that:

- benchmarks should routinely include paraphrased, perturbed, and out-of-distribution variants, with explicit reporting of *relative drops* in performance;
- model comparisons should specify the decoding regime as part of the evaluation protocol, since the ranking of models can change when moving from deterministic to stochastic decoders;
- benchmark releases should prioritise publishing per-input multi-sample outputs where possible, enabling secondary analyses of distributional behaviour.

(3) Making trajectory diversity visible

Our trajectory-space visualizations demonstrate that capability is not a single scalar but a structured object: different regions of input space support different forests of trajectories, with varied depth, diversity, and stability. Making such diversity visible—through violin plots, multi-path diagrams, or kernel-based projections—turns what is currently a hidden internal degree of freedom into an explicit part of model reporting.

7.4 Implications for System Design and Deployment

From a systems perspective, our results argue against treating bitwise determinism as a one-size-fits-all principle.

Where strong determinism is still essential. There are scenarios where reproducible byte-level behaviour is non-negotiable:

- regression testing and continuous integration for large codebases;
- forensic analysis and scientific auditing, where exact replay of past behaviour is required;
- specific on-policy RL pipelines that assume deterministic environment dynamics.

In these cases, investing in batch-invariant kernels, stable low-level libraries, and careful environment fingerprinting is appropriate.

Where deterministic defaults become misleading. In user-facing deployments and safety assessment, however, deterministic defaults can:

- *underestimate risk*, by never surfacing low-probability but high-impact failures present in the tail of $p_\theta(\cdot \mid x)$;
- *underestimate capability*, by failing to traverse valid reasoning paths or alternative solutions that stochastic decoding easily finds;
- *overstate robustness*, by conflating the stability of a single trajectory with stability of the whole conditional distribution.

Our safety analyses, in particular, show that systems can look “extremely safe” under greedy evaluation while retaining substantial harmful mass under modest multi-sample decoding.

Toward decoder-aware system design. A more nuanced design philosophy would:

- expose deterministic and stochastic modes as explicit options, with clear documentation of trade-offs;
- separate serving concerns (throughput, latency, cost) from evaluation and monitoring concerns (distributional fidelity, tail-risk estimation);
- integrate lightweight multi-sample diagnostics into deployment: periodic probing with $k > 1$, paraphrase-based canaries, and other distributional checks.

7.5 Limitations and Threats to Validity

Our study has several limitations, which bound the scope of the conclusions.

Model and task coverage. We examine a particular set of models and tasks: classification, controlled generation, reasoning, and safety-oriented prompts. Other architectures (e.g., larger mixture-of-experts, retrieval-augmented, or deeply multimodal models) might exhibit different quantitative patterns. The qualitative message—that single-trajectory evaluation can diverge sharply from stochastic behaviour—is likely to generalise, but should be rechecked in new domains.

Decoding hyperparameters and search strategies. We instantiate stochastic decoding with a small family of temperatures, top- p , and multi-sample budgets. Alternative search procedures (beam search with diversity penalties, tree search over chains of thought, entropy-regularised decoders) may change the precise numbers. Our claims concern the *direction* of effects: low-entropy, single-path policies consistently hide behaviour that higher-entropy policies reveal.

Prompting and data contamination. We adopt standard prompting schemes and, for newer benchmarks, take care to mitigate training-data contamination. Nonetheless, we cannot guarantee that no examples (or close variants) appear in pretraining or tuning corpora, especially for proprietary models. Such contamination could make some gaps between deterministic and stochastic evaluation larger or smaller than they would be on fully held-out distributions.

Metric choices and semantic nuances. Our metrics treat success/failure and constraint satisfaction as discrete events. For summarization and style, there is inherent ambiguity about what counts as “equally good” or “semantically equivalent”. Our automatic metrics and manual checks capture only part of this nuance. In principle, decoding regimes could differ substantially in surface form while remaining similar in downstream utility, or vice versa.

Temporal and infrastructural drift. APIs, kernels, batching policies, and model versions evolve. Our experiments capture a snapshot of a rapidly moving ecosystem. Future infrastructure may reduce some sources of numerical nondeterminism, and future model families may be trained with stronger incentives for stability under paraphrase. Our empirical results should therefore be seen as evidence about current practice, not as fixed laws.

7.6 Open Directions

Our perspective opens several directions for further work.

Principled stochastic evaluators. If distributional variability is central, evaluation protocols should be designed to capture it efficiently. This calls for:

- principled choices of sampling budgets and perturbation families;
- confidence intervals and power analyses for stochastic metrics;
- benchmarks that explicitly score not only point estimates but full *risk profiles* as a function of k and input variation.

Training for distributional robustness. Most fine-tuning and alignment pipelines optimise deterministic losses derived from one (or a few) completions per input. Our results suggest investigating objectives that directly reward:

- stability of success probabilities across paraphrases and decoding regimes;
- smoothness of trajectory-space representations under perturbations;
- robustness of multi-path reasoning, not merely the quality of a single chain.

Safety diagnostics beyond one trajectory. The safety analyses in this work take only first steps toward *distributional risk* as the primary object. Future work could develop:

- rare-event estimation techniques tailored to LLM harmful modes;
- adaptive stress tests that jointly search over prompts and decoders;
- online monitors that track how the tail of $p_\theta(\cdot \mid x)$ shifts under updates and deployment feedback.

Beyond text-only LLMs. Real-world systems increasingly involve multimodal inputs, tools, and agents. Each layer introduces new stochastic components: perception, environment dynamics, tool responses, and interactive users. Extending the lens of deterministic vs. stochastic behaviour to these richer settings is a natural next step, and will likely further diminish the appeal of strictly deterministic defaults.

Overall, our results argue that the push toward bitwise deterministic inference, while useful for some engineering goals, is misaligned with the probabilistic nature of LLMs and with many of the questions we care about in capability evaluation and safety. The aim should not be to suppress stochasticity, but to *measure, shape, and leverage* it in a controlled way.

8 Conclusion

Our core message is simple but uncomfortable: treating large language models as if they were deterministic programs is a category error. Modern LLMs are *stochastic semantic machines*, yet our dominant evaluation and deployment practices insist on a single canonical output—a greedy, bitwise-reproducible string that often tells a comforting but misleading story. Through Stochastic CHAOS, we show that this deterministic lens systematically **erases the very signals we most care about**: it hides emergent abilities that only appear as phase transitions in success probability, it collapses rich multi-path reasoning into brittle single traces, and it creates *deterministic illusions* of robustness that vanish the moment we look at the underlying distribution instead of its mode.

Conceptually, our results reframe nondeterminism from a nuisance to be engineered away into a first-class object of scientific inquiry. We disentangle three often conflated goals—*bitwise determinism*, *distributional reproducibility*, and *semantic stability*—and demonstrate that over-optimizing the first can actively undermine the latter two. Even modest multi-sample budgets and lightweight perturbations expose substantial gaps between greedy and stochastic behavior in instruction following, compositional generalization, multi-step reasoning, and safety. These gaps are not small implementation artifacts: they reveal that key properties of LLM behavior live in the *geometry* of $p_\theta(y \mid x)$, not in any single draw from it. **In this view, quantities such as exploration gain, illusion indices, and stochastic risk gaps are not optional diagnostics but indispensable coordinates for navigating model behavior.**

Practically, this shift demands a rethinking of how we design benchmarks, inference stacks, and governance protocols. In particular, our findings suggest that future work should:

- (i) **Elevate distributional metrics to first-class status.** Benchmarks should report and optimize *multi-sample* quantities—success probabilities, exploration gains, and risk gaps—rather than celebrating single-run scores. Emergent capabilities should be documented as smooth or abrupt changes in success probability, not as binary “pass/fail” headlines.
- (ii) **Treat reasoning as an ensemble phenomenon.** Evaluation of chains-of-thought should instrument *ensembles* of trajectories, analyzing their diversity, stability, and failure modes, instead of anointing one canonical trace as “the” reasoning path.
- (iii) **Redesign safety audits for tails, not modes.** Safety evaluation and red-teaming should explicitly budget for exploring low-probability regions—paraphrased attacks, mixed decoders, and small sampling budgets—rather than certifying models under a single, deterministic decoding recipe.
- (iv) **Repurpose determinism as a diagnostic tool, not the default norm.** Engineering effort should shift from enforcing global bitwise determinism toward achieving *robust distributional reproducibility and semantic stability* under realistic, noisy serving conditions. Deterministic modes remain invaluable for debugging, ablation, and regression testing, but they should no longer be the primary lens for understanding or governing LLM behavior.

Our study is necessarily scoped, and it opens more questions than it closes. We analyze a subset of models, decoding policies, and stress tests; richer families of samplers, adaptive allocation of evaluation budgets to “risky” prompts, and training-time regularizers that encourage epistemically meaningful variability all remain open directions. Equally important are institutional questions: how should regulators, practitioners, and standards bodies translate *distributional* notions of risk into operational guarantees and service-level objectives? **We view Stochastic CHAOS as a starting point for this agenda.** If the community is serious about understanding and governing LLMs as they truly are, then we must move beyond “one prompt, one answer” and embrace controlled randomness as a central design axis. *In the long run, honest generalization and transparent failure will not come from suppressing stochasticity, but from learning to measure, control, and reason with it.*

References

- Berk Atil, Sarp Aykent, Alexa Chittams, Lisheng Fu, Rebecca J. Passonneau, Evan Radcliffe, Guru Rajan Rajagopal, Adam Sloan, Tomasz Tudrej, Ferhan Ture, Zhe Wu, Lixinyu Xu, and Breck Baldwin. 2024. [Stability of large language models under repeated evaluation](#). *arXiv preprint*, arXiv:2408.04667.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). *Advances in Neural Information Processing Systems*, 33:1877–1901.
- Hou Pong Chan, Samuel Cahyawijaya, and Pascale Fung. 2021. [Controllable summarization with constrained markov decision process](#). *Transactions of the Association for Computational Linguistics*, 9:1213–1232.
- Yan Chen, Yu Huang, Xiang Li, Houyi Li, Yanjie Zhao, Zhiru Zhang, G. Edward Suh, Christina Delimitrou, Christopher Batten, Edward F. Redmond, et al. 2022. [Towards training reproducible deep learning models](#). In *Proceedings of the 44th International Conference on Software Engineering (ICSE)*, pages 1551–1563. IEEE.
- Wei-Fan Chiang et al. 2013. Deterministic replay for scientific debugging of large-scale parallel programs. In *Proceedings of SC*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Łukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. In *Advances in Neural Information Processing Systems*. Dataset: GSM8K.
- James Demmel, Willow Ahrens, and Hong Diep Nguyen. 2016. [Efficient reproducible floating point summation and BLAS](#). Technical Report UCB/EECS-2016-121, EECS Department, University of California, Berkeley.
- Angela Fan, David Grangier, and Michael Auli. 2018a. [Controllable abstractive summarization](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 45–54, Melbourne, Australia. Association for Computational Linguistics.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018b. Hierarchical neural story generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 889–898.
- Deep Ganguli, Danny Hernandez, Liane Lovitt, Nova Dassarma, Tom Henighan, Andy Jones, Nicholas Joseph, Jackson Kernion, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Dawn Drain, Nelson Elhage, Sheer El Showk, Stanislav Fort, Zac Hatfield-Dodds, Scott Johnston, Shauna Kravec, Neel Nanda, Kamal Ndousse, Catherine Olsson, Daniela Amodei, Tom B. Brown, Jared Kaplan, Sam McCandlish, Chris Olah, Dario Amodei, and Jack Clark. 2022a. [Predictability and surprise in large generative models](#). In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*, pages 1747–1764. ACM.
- Deep Ganguli, Liane Lovitt, Nick Rider, Lilian Weng, Amanda Askell, Yuntao Bai, et al. 2022b. Red teaming language models with language models. *arXiv preprint*. ArXiv:2202.03286.
- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. 2020. [Shortcut learning in deep neural networks](#). *Nature Machine Intelligence*, 2(11):665–673.

- Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. Did Aristotle use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions of the Association for Computational Linguistics*, 9:346–361. Dataset: StrategyQA.
- Ryan Greenblatt, Carson Denison, Benjamin Wright, Fabien Roger, Monte MacDiarmid, Sam Marks, Johannes Treutlein, Tim Belonax, Jack Chen, David Duvenaud, Akbir Khan, Julian Michael, Sören Mindermann, Ethan Perez, Linda Petrini, Jonathan Uesato, Jared Kaplan, Buck Shlegeris, Samuel R. Bowman, and Evan Hubinger. 2024. Alignment faking in large language models. *arXiv preprint arXiv:2412.14093*.
- Odd Erik Gundersen, Saeid Shamsaliei, Hakon Sletten Kjaernli, and Helge Langseth. 2023. [On reporting robust and trustworthy conclusions from model comparison studies involving neural networks and randomness](#). In *Proceedings of the 2023 ACM Conference on Reproducibility and Replicability (ACM-REP '23)*, pages 37–61, New York, NY, USA. Association for Computing Machinery.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330, Sydney, Australia. PMLR.
- Junxian He, Wojciech Kryściński, Bryan McCann, Nazneen Fatema Rajani, and Caiming Xiong. 2020. [CTRLsum: Towards generic controllable text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5866–5887, Online. Association for Computational Linguistics.
- Tony He and Thinking Machines Lab. 2025. Defeating nondeterminism in llm inference. <https://blog.connectionism.ai/defeating-nondeterminism-in-llm-inference>. Blog post.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching machines to read and comprehend](#). In *Advances in Neural Information Processing Systems*, volume 28.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. [The curious case of neural text degeneration](#). *arXiv preprint arXiv:1904.09751*.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*.
- Evan Hubinger, Carson Denison, Jesse Mu, Mike Lambert, Meg Tong, Monte MacDiarmid, Tamera Lanham, Daniel M. Ziegler, Tim Maxwell, Newton Cheng, Adam Jermy, Amanda Askell, Ansh Radhakrishnan, Cem Anil, David Duvenaud, Deep Ganguli, Fazl Barez, Jack Clark, Kamal Ndousse, Kshitij Sachan, Michael Sellitto, Mrinank Sharma, Nova DasSarma, Roger Grosse, Shauna Kravec, Yuntao Bai, Zachary Witten, Marina Favaro, Jan Brauner, Holden Karnofsky, Paul Christiano, Samuel R. Bowman, Logan Graham, Jared Kaplan, Sören Mindermann, Ryan Greenblatt, Buck Shlegeris, Nicholas Schiefer, and Ethan Perez. 2024. [Sleeper agents: Training deceptive llms that persist through safety training](#). *arXiv preprint, arXiv:2401.05566*.
- Aaqib Kaikaus, Adeel Anwar, and Yusuf M. Khan. 2024. [Determinism of chatgpt in code generation: A systematic evaluation](#). *arXiv preprint, arXiv:2405.12345*.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). *Advances in Neural Information Processing Systems*. ArXiv:2205.11916.
- Yixin Liu, Alexander R. Fabbri, Jiawen Chen, Yilun Zhao, Simeng Han, Shafiq Joty, Pengfei Liu, Dragomir Radev, Chien-Sheng Wu, and Arman Cohan. 2024. [Benchmarking generation and evaluation capabilities of large language models for instruction controllable summarization](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, Mexico City, Mexico. Association for Computational Linguistics. InstruSum dataset: <https://github.com/yale-nlp/InstruSum>.

- Prabhat Nagarajan, Garrett Warnell, and Peter Stone. 2018. [Deterministic implementations for reproducibility in deep reinforcement learning](#). *arXiv preprint*, arXiv:1809.05676.
- Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence rnns and beyond](#). *CoRR*, abs/1602.06023.
- OpenAI. 2024. Reproducible outputs. <https://platform.openai.com/docs/guides/reproducible-outputs>. Accessed 2025-11-22.
- Arkil Patel, Satwik Bhattamishra, and Navin Goyal. 2021. Are NLP models really able to solve simple math word problems? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2080–2094. Association for Computational Linguistics. Dataset: SVAMP.
- Ethan Perez, Sam Ringer, Jane Irvine, Andrew Kissane, Megan Chen, Lawrence Chan, Sean Heiner, Catherine Olsson, Samuel R. Bowman, Christopher Ré, et al. 2022. Discovering language model behaviors with model-written evaluations. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*. ArXiv:2206.11871.
- PyTorch Core Team. 2020. Reproducibility — controlling sources of randomness in PyTorch. <https://pytorch.org/docs/stable/notes/randomness.html>. Accessed 2025-11-22.
- PyTorch Developers. 2024. Reproducibility — controlling sources of non-determinism in PyTorch. <https://pytorch.org/docs/stable/notes/randomness.html>. Accessed: 2025-11-27.
- Sudha Rao and Joel Tetreault. 2018. [Dear sir or madam, may i introduce the GYAFC dataset: Corpus, benchmarks and metrics for formality style transfer](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 129–140, New Orleans, Louisiana. Association for Computational Linguistics.
- Shiori Sagawa, Jason Wei, Yiding Li, Yi Tay, Josip Djolonga, Tatsunori B. Hashimoto, Behnam Neyshabur, Ed H. Chi, Jeff Dean, William Fedus, Deep Ganguli, et al. 2023. [The many faces of emergent abilities in large language models](#). *arXiv preprint arXiv:2304.15004*.
- Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo. 2023. [Are emergent abilities of large language models a mirage?](#) In *Advances in Neural Information Processing Systems*, volume 36, pages 55565–55581.
- Mostafa Shahriari, Rudolf Ramler, and Lukas Fischer. 2022. [How do deep-learning framework versions affect the reproducibility of neural network models?](#) *Machine Learning and Knowledge Extraction*, 4(4):888–911.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Dipankar Srirag, Aditya Joshi, Jordan Painter, and Diptesh Kanojia. 2025. [BESSTIE: A benchmark for sentiment and sarcasm classification for varieties of english](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, Vienna, Austria. Association for Computational Linguistics. Datasets and code: <https://github.com/unswnlp/BESSTIE>.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*. ArXiv:1804.07461.

- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023a. [Self-consistency improves chain-of-thought reasoning in language models](#). *arXiv preprint*, arXiv:2203.11171.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, and Denny Zhou. 2023b. Self-consistency improves chain of thought reasoning in large language models. *International Conference on Learning Representations (ICLR)*. ArXiv:2203.11171.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022a. [Emergent abilities of large language models](#). *Transactions on Machine Learning Research*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022b. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems 35 (NeurIPS)*. ArXiv:2201.11903.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Graham Neubig, and Yuan Cao. 2024. Tree of thoughts: Deliberate problem solving with large language models. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*. ArXiv:2305.10601.
- Shunyu Yao, Dian Zhao, Yuan Yu, et al. 2023. Tree of thoughts: Deliberate problem solving with large language models. *arXiv preprint arXiv:2305.10601*.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems*, volume 28.
- Ziyang Zhang, Xinheng Ding, Jiayi Yuan, Rixin Liu, Huizi Mao, Jiarong Xing, and Zirui Liu. 2025. [Deterministic inference across tensor parallel sizes that eliminates training-inference mismatch](#). *arXiv preprint*, arXiv:2511.17826.
- Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Sun, Jeff Huang, Cody Hao Yu, Shiyi Cao, Christos Kozyrakis, Ion Stoica, Joseph E. Gonzalez, Clark Barrett, and Ying Sheng. 2024. [Sglang: Efficient execution of structured language model programs](#). arXiv preprint arXiv:2407.01239.
- Yang Zhuang, Zachary C. Lipton, Srinivasan Parthasarathy, and Weiwei Tu. 2021. [Randomness in neural network training: Characterizing the impact of tooling](#). In *Proceedings of the 3rd Conference on Machine Learning and Systems (MLSys)*.