

---

# Can a Bayesian Oracle Prevent Harm from an Agent?

---

Yoshua Bengio<sup>1,2,3,4,\*</sup>

Michael K. Cohen<sup>5,\*</sup>

Nikolay Malkin<sup>6,4,\*</sup>

Matt MacDermott<sup>1,2,7</sup>

Damiano Fornasiero<sup>1,2,8,†</sup>

Pietro Greiner<sup>1,2,9,†</sup>

Younesse Kaddar<sup>10,†</sup>

<sup>1</sup>Mila, Montréal, Canada

<sup>2</sup>LawZero, Montréal, Canada

<sup>3</sup>Université de Montréal, Montréal, Canada

<sup>4</sup>CIFAR, Learning in Machines and Brains, Canada

<sup>5</sup>UC Berkeley, Berkeley, USA

<sup>6</sup>University of Edinburgh, Edinburgh, UK

<sup>7</sup>Imperial College London, London, UK

<sup>8</sup>Universitat de Barcelona, Barcelona, Spain

<sup>9</sup>Università degli studi di Padova, Padua, Italy

<sup>10</sup>University of Oxford, Oxford, UK

\*,<sup>†</sup>First three authors and last three authors listed alphabetically

## Abstract

Is there a way to design powerful AI systems based on machine learning methods that would satisfy probabilistic safety guarantees? With the long-term goal of obtaining a probabilistic guarantee that would apply in every context, we consider estimating a context-dependent bound on the probability of violating a given safety specification. Such a risk evaluation would need to be performed at run-time to provide a guardrail against dangerous actions of an AI. Noting that different plausible hypotheses about the world could produce very different outcomes, and because we do not know which one is right, we derive bounds on the safety violation probability predicted under the true but unknown hypothesis. Such bounds could be used to reject potentially dangerous actions. Our main results involve searching for cautious but plausible hypotheses, obtained by a maximization that involves Bayesian posteriors over hypotheses. We consider two forms of this result, in the i.i.d. case and in the non-i.i.d. case, and conclude with open problems towards turning such theoretical results into practical AI guardrails.

## 1 INTRODUCTION

Ensuring that an AI system will not misbehave is a challenging open problem [Bengio et al., 2024], particularly in the current context of rapid growth in AI capabilities. Governance measures and evaluation-based strategies have been proposed to mitigate the risk of harm from highly capable AI systems, but do not provide any form of safety guarantee when no undesired behavior is detected. In contrast, the *safe-by-design* paradigm involves AI systems with quantitative safety guarantees, and therefore could represent a stronger form of protection [Dalrymple et al., 2024]. However, how to design such systems remains an open problem too.

Since testing an AI system for violations of a safety specification in every possible context, *e.g.*, every (query, output) pair, is impossible, we consider a rejection sampling approach that declines a candidate output or action if it has too high a probability of violating an unknown safety constraint; we refer to such violations as “harm”. The approach we outline does not require the developer to write down a safety constraint. The only requirement is the construction of a data generating process that produces information about harm, which we leave to future work. We also note that maintaining a Bayesian belief distribution about how to interpret a human-specified safety specification would protect the AI from committing to an incorrect interpretation. Here we instead focus on a question inspired by risk-management practice [McNeil et al., 2015]: even though the true probability of harm following from some proposed action is unknown,

because the true data-generating process is unknown, can we bound that risk using quantities that can be estimated by machine learning methods given the observed data?

To illustrate this question, consider a committee of “wise” humans whose theories about the world are all equally compatible with the available data, knowing that an unknown member of the committee has the correct theory. Each committee member can make a prediction about the probability of future harm that would result from following some action in some context. Marginalizing this harm probability over the committee members amounts to making them vote with equal weights. If the majority is aligned with the correct member’s prediction, then all is good, *i.e.*, if the correct theory predicts harm, then the committee will predict harm and can choose to avoid the harmful action. But what if the correct member is in the minority regarding their harm prediction? To get a *guarantee* that the true harm probability is below a given threshold, we could simply consider the committee member whose theory predicts the highest harm probability, and we would be sure that their harm probability prediction upper bounds the true harm probability. In practice, committee members are not equally “wise”, so we can correct this calculation based on how plausible the theory harbored by each committee member is. In a Bayesian framework, the plausibility of a theory corresponds to its posterior over all theories given the observed data, which is proportional to the data likelihood given the theory multiplied by the prior probability of that theory.

In this paper, we show how results about posterior consistency can provide probabilistic risk bounds. All the results have the form of inequalities, where the true probability of harm is upper bounded by a quantity that can in principle be estimated, given enough computational resources to approximate Bayesian posteriors over theories given the data. In addition, these are not hard bounds but only hold with some probability, and there is generally a trade-off between that probability and the tightness of the bound. We study two scenarios in the corresponding sections: the i.i.d. data setting in §3 and the non-i.i.d. data setting in §4, followed by an experiment in §5. In all cases, a key intermediate result is a bound relating the Bayesian posterior on the unknown true theory and the probability of other theories (with propositions labeled **True theory dominance**). The idea is that because the true theory generated the data, its posterior tends to increase as more data is acquired, and in the i.i.d. case it asymptotically dominates other theories. From such a relationship, the harm risk bound can be derived with very little algebra (yielding propositions labeled **Harm probability bound**). These bounds depend on the unknown true theory having nonzero prior weight.

We conclude the paper with a discussion of open problems that should be considered in turning such bounds into a safe-by-design AI system, taking into account the challenge of representing the notion of harm and the imperfect

estimation of the required conditional probabilities.

**Related Work.** The concept of blocking actions based on probabilistic criteria resembles probabilistic shielding in Markov Decision Processes (MDPs) [Jansen et al., 2020], but our bounds do not require knowledge of the true model, extending beyond Carr et al. [2023]’s work on partially observable MDPs. While Beckers et al. [2023] and Richens et al. [2022] propose specific frameworks for quantifying harm, our approach remains agnostic, by only requiring harmful outcomes to be representable as binary events  $H = 1$ , allowing various harm definitions, while providing conservative probability bounds for safety-critical contexts. Osband and Van Roy [2017] study translating concentration bounds from a pure predictive setting to an MDP setting with exploration, whereas we address an orthogonal question: providing safety guarantees without relying on potentially harmful exploration to gain information.

## 2 SAFE-BY-DESIGN AI?

Before an AI is built and deployed, it is important that the developers have high assurances that the AI will behave well. Dalrymple et al. [2024] propose an approach to “guaranteed safe AI” designs with built-in high-assurance quantitative safety guarantees, although these guarantees can sometimes be probabilistic and only asymptotic. It remains an open question whether and how that research program can be realized. The authors take existing examples of quantitative guarantees in safety-critical systems and motivate why such a framework should be adopted if we ever build AI systems that match or exceed human cognitive abilities and could potentially act in dangerous ways. Their program is motivated by current known limitations of state-of-the-art AI systems based on deep learning, including the challenge of engineering AI systems that robustly act as intended [Cohen et al., 2022b, Krakovna et al., 2020, Pan et al., 2021, Pang et al., 2023, Zhuang and Hadfield-Menell, 2020, Skalse et al., 2022, 2023, Karwowski et al., 2024, Skalse et al., 2024].

The approach proposed by Dalrymple et al. [2024] has the following components: a *world model* (which can be a distribution about hypotheses explaining the data), a *safety specification* (what are considered unacceptable states of the world), and a *verifier* (a computable procedure that checks whether a policy or action violates the safety specification).

Here, we study a system that infers a probabilistic world model, or *theory*,  $\tau$  and updates its estimate of  $\tau$  via machine learning, using the stream of observed data  $D$ . The observations  $D$  are assumed to come from a data-generating process given by a ground-truth world model  $\tau^*$ , which lies in the system’s space of possible theories. We do not assume that the agent observes any sort of Markovian state. The inference of the theory  $\tau$  is Bayesian, meaning that the system maintains an estimate  $q$  of the true posterior  $P(- | D)$  over

theories:  $q(\tau | D) \approx P(\tau | D)$ , where  $P(\tau | D)$  is proportional to the product of the prior probability  $P(\tau)$  with the likelihood of the observations under the theory,  $P(D | \tau)$ . In the simplest case,  $q$  is a point estimate, optimally placing its mass on the mode of the posterior. Assuming an observation  $x$  and a theory  $\tau$  are independent given  $D$ , inference of the latent theory  $\tau$  allows the system to approximate conditional probabilities  $P(y | x, D) \approx \mathbb{E}_{\tau \sim q(\tau | D)} [P_\tau(y | x, D)]$  over any random variables  $X, Y$  known to the world model.

The safety specification is given in the form of a binary random variable  $H$  (called ‘‘harm’’ below) whose probability given the other variables depends on the theory  $\tau$ . We are interested in predicting the probability of harm under the true theory  $\tau^*$ . Because  $\tau^*$  is unknown, we propose to estimate upper bounds on this probability using the estimated posteriors. These upper bounds can be used as thresholds for a *verifier* that checks whether the risk of harm falls below some acceptable level.

Following Dalrymple et al. [2024], we assume that the notion of harm has been specified, possibly in natural language, and that the ambiguities about its interpretation are encoded within the Bayesian posterior  $P(\tau | D)$ . This paper focuses on the verifier under different assumptions of i.i.d. or non-i.i.d. data.

**What do the observations and context represent?** We give a possible interpretation of the objects introduced in the preceding discussion in the simple case of an agent acting in a fully observed environment (Markov Decision Process, or MDP), where the theory is a transition model and the occurrence of harm at a state  $s$  is conditionally independent of all other variables given  $s$ .

- Observations  $Z$  are transitions  $z = (s, a, s', r)$ , where  $s$  is a state,  $a$  is an action,  $s'$  is the next state, and  $r$  is the reward received.
- Theories  $\tau$  encode the state visitation, transition probabilities<sup>1</sup>, as well as the behavior policy from which observations are collected.
- The dataset  $D$  is a sequence of observed transitions.
  - In the non-i.i.d. setting,  $D$  could consist, for example, of the observations from a finite rollout in the order in which they occurred.
  - In the i.i.d. setting,  $D$  would need to be a sequence of independent samples from a fixed state-action-reward visit distribution. This could be achieved, for example, by rolling out a behavior policy multiple times and randomly sampling transitions from the resulting trajectories.

In the common special case of a contextual bandit MDP under a fixed policy, the two coincide.

- The context  $X$  is a pair  $x = (s, a)$ , where  $s$  is a state and  $a$

<sup>1</sup>To be more precise, we can obtain transition probabilities from  $\tau$  by conditioning on  $(s, a)$  to get  $P_\tau(s', r | s, a)$ , but only for state-action pairs with non-zero probability under  $\tau$ .

is an action being considered at state  $s$ .

- The harm probability  $P(H = 1 | X = (s, a), \tau, D)$  can be any function of the theory  $\tau$ , the context  $x = (s, a)$ , and the data  $D$ . For example, this probability could be derived from a fixed specification of what it means for a state  $s'$  to be harmful,  $P_{\text{harm}}(H = 1 | s')$ . Then, the harm probability could be computed as  $P(H = 1 | X = (s, a), \tau, D) = \sum_{s', r} P_\tau(s', r | s, a) P_{\text{harm}}(H = 1 | s')$ .

We note that the interpretation of harm probability in the example above includes the case where the occurrence of harm is an observed variable  $s'_{\text{harm}}$  that is part of the state  $s'$ : in that case, we set  $P_{\text{harm}}(H = 1 | s') = 1$  if  $s'$  is harmful (i.e.,  $s'_{\text{harm}} = 1$ ), and  $P_{\text{harm}}(H = 1 | s') = 0$  otherwise. Then the harm probability is just the probability, under  $\tau$ , of reaching a harmful state, and observations of harm in  $D$  affect the Bayesian posterior over theories.

This interpretation also includes the case where the harm probability is a function of the state  $s'$ , but (non)occurrence of harm is not observed in  $D$ . For example, a language model encoding world knowledge and human preferences or constraints, or an iterative reasoning procedure that uses those constraints, could generate some specification of harm  $P_{\text{harm}}(H = 1 | s')$ , perhaps unreliably.

Finally, a setting that separates the predicted next state  $s'$  from the harm variable  $H$  in this way gives a framework for studying how an agent might tamper with harm guardrails. If the state  $s'$  decomposes as  $s' = (s'_{\text{harm}}, s'_{\text{rest}})$ , and  $P_{\text{harm}}$  is deterministic as a function of  $s'_{\text{harm}}$ , except for some difficult-to-reach values of  $s'_{\text{rest}}$ , then the agent can try to reach those values of  $s'_{\text{rest}}$ , so that harm is ‘recorded’ as not having occurred, even though it has. (See the end of §4.)

### 3 I.I.D. DATA

Following the notation introduced in the previous section, here, we consider the easier-to-analyze case where the observed examples  $D = (z_1, z_2, \dots, z_n)$  are sampled i.i.d. from the unknown distribution  $\tau^*$ . Assuming that the prior assigns nonzero mass to  $\tau^*$ , and all theories are distinct distributions, it can be shown that the posterior  $P(\tau | D)$  converges to a point mass at  $\tau^*$ . We show that for sufficiently large  $n$ , we can bound the probability under  $\tau^*$  of a harm event  $H = 1$  given conditions  $x$  (e.g., a context and an action) by considering the probability of  $H = 1$  given  $x$  and  $D$ , under a plausible but ‘cautious’ theory  $\tilde{\tau}$  that maximizes  $P(\tilde{\tau} | D) P(H = 1 | x, \tilde{\tau}, D)$ .

**Setting.** Fix a complete separable metric space  $\mathcal{Z}$ , called the *observation space*, let  $\mathcal{F}$  be its Borel  $\sigma$ -algebra, and fix a  $\sigma$ -finite measure  $\mu$  on  $\mathcal{F}$ . A *theory* is a probability distribution on the measurable space  $(\mathcal{Z}, \mathcal{F})$  that is absolutely continuous w.r.t.  $\mu$ . If  $\tau$  is a theory, we denote by  $P_\tau(\cdot)$  the Radon-Nikodym derivative  $\frac{d\tau}{d\mu} : \mathcal{Z} \rightarrow \mathbb{R}_{\geq 0}$ , which is

uniquely defined up to  $\mu$ -a.e. equality.

One can keep in mind two cases:

- (1)  $\mathcal{Z}$  is a finite or countable set and  $\mu$  is the counting measure. Theories  $\tau$  are equivalent to probability mass functions  $P_\tau : \mathcal{Z} \rightarrow \mathbb{R}_{\geq 0}$ .
- (2)  $\mathcal{Z} = \mathbb{R}^d$  and  $\mu$  is the Lebesgue measure. Theories are equivalent to their probability density functions  $P_\tau : \mathcal{Z} \rightarrow \mathbb{R}_{\geq 0}$  up to a.e. equality.

Consider a countable (possibly finite) set of theories  $\mathcal{M}$  containing a *ground truth* theory  $\tau^*$  and fix a choice of a (measurable) density function  $P_\tau$  for each  $\tau \in \mathcal{M}$ .

**Definition of posterior as a random variable.** If  $P$  is a prior distribution<sup>2</sup> on  $\mathcal{M}$  and  $z \in \mathcal{Z}$ , we define the posterior to be the distribution with mass function

$$P(\tau | z) = \frac{P(\tau) P_\tau(z)}{\sum_{\tau' \in \mathcal{M}} P(\tau') P_{\tau'}(z)} \propto P(\tau) P_\tau(z), \quad (1)$$

assuming the denominator converges and the sum is nonzero. Otherwise, the posterior is considered to be undefined. As written, the posterior depends on the choice of density functions  $P_\tau$ , but any two  $P_\tau$  that are  $\mu$ -a.e. equal yield the same posterior for  $\mu$ -a.e.  $z$ .

For  $z_1, z_2 \in \mathcal{Z}$ , we write  $P(\cdot | z_1, z_2)$  for the posterior given observation  $z_2$  and prior  $P(\cdot | z_1)$ , and similarly for a longer sequence of observations. It can be checked that  $P(\cdot | z_1, \dots, z_t)$  is invariant to the order of  $z_1, \dots, z_t$  and that it is defined in one order if and only if it is defined in all orders. This allows us to unambiguously write  $P(\cdot | D)$  where  $D$  is a finite multiset of observations, and we have

$$P(\tau | D) \propto P(\tau) \prod_{z \in D} P_\tau(z). \quad (2)$$

Let  $\tau^* \in \mathcal{M}$  be the ground truth theory and  $P(\cdot)$  a prior over  $\mathcal{M}$ . Consider a sequence of i.i.d.  $\mathcal{Z}$ -valued random variables  $Z_1, Z_2, \dots$  (whose realizations are the *observations*), where each  $Z_i$  follows the distribution  $\tau^*$ . For any  $t \in \mathbb{N}$ , the posterior  $P(\cdot | Z_{1:t})$  is then a random variable taking values in the space of probability mass functions on  $\mathcal{M}$ .<sup>3</sup>

**Bayesian posterior consistency.** We state, in our setting, a result about the concentration of the posterior at the ground truth theory  $\tau^*$  as the number of observations increases.

<sup>2</sup>To be precise,  $\mathcal{M}$  is endowed with the counting measure and we flexibly interchange distributions and mass functions on  $\mathcal{M}$ .

<sup>3</sup>To be precise, if the  $Z_i$ 's are measurable functions from a sample space  $\Omega$  to  $\mathcal{Z}$  and  $\langle Z_1, \dots, Z_t \rangle$  is their pairing, the random variable  $P(\cdot | Z_{1:t}) : \Omega \xrightarrow{\langle Z_1, \dots, Z_t \rangle} \mathcal{Z}^t \xrightarrow{P(\cdot | -)} \mathbb{P}(\mathcal{M})$  has codomain the space  $\mathbb{P}(\mathcal{M})$  of functions  $\mathcal{M} \rightarrow \mathbb{R}_{\geq 0}$  summing to 1. The function  $P(\cdot | -)$  mapping a sequence of observations to the posterior probability mass function is measurable, due to each  $P_\tau(z)$  being measurable in  $z$  and elementary facts.

**Proposition 3.1 (True theory dominance).** *Under the above conditions and supposing that  $P(\tau^*) > 0$ , the posterior  $P(\cdot | Z_{1:t})$  is almost surely defined for all  $n$ , and the following almost surely hold:*

- (a)  $P(\cdot | Z_{1:t}) \xrightarrow{t \rightarrow \infty} \delta_{\tau^*}$  as measures, where  $\delta_{\tau^*}$  is the Dirac measure, which assigns mass 1 to the theory  $\tau^*$  and 0 elsewhere; equivalently,  $\lim_{t \rightarrow \infty} P(\tau | Z_{1:t}) = \mathbb{1}[\tau = \tau^*]$ .
- (b) There exists  $N \in \mathbb{N}$  such that  $\arg \max_{\tau \in \mathcal{M}} P(\tau | Z_{1:t}) = \tau^*$  for all  $t \geq N$ .

(All proofs can be found in appendix A.) Note that this result assumes that all theories in  $\mathcal{M}$  are distinct as probability measures (so no two of the  $P_\tau$  are  $\mu$ -a.e. equal).

**On necessity of conditions.** The i.i.d. assumption in Prop. 3.1 is necessary; see Remark B.1 for an example where  $\limsup_{t \rightarrow \infty} P(\tau^* | Z_{1:t})$  does not almost surely approach 1.

However, in practice, Prop. 3.1 can be adapted to more general scenarios, by substituting the subset  $\mathcal{T} \subseteq \mathcal{M}$  of theories with minimum relative entropy to  $\tau^*$  for  $\tau^*$  (when  $\tau^*$  is not in  $\mathcal{M}$ ). Then, we can replace convergence to  $\delta_{\tau^*}$  with  $P(\mathcal{T} | Z_{1:t}) \rightarrow 1$  in (a), and replace  $\tau^*$  with  $\mathcal{T}$  in (b).

**On generalizations to uncountable sets of theories.** The proof of Prop. 3.1 critically uses that the set of theories  $\mathcal{M}$  is countable when passing from almost sure convergence under  $\tau^*$  sampled from the prior to almost sure convergence for any particular  $\tau^*$  with positive prior mass. This argument fails for uncountable  $\mathcal{M}$ ; indeed, characterization of the  $\tau^*$  for which the posterior converges to  $\delta_{\tau^*}$  is a delicate problem [Freedman, 1963, 1965, Diaconis and Freedman, 1986]. Concentration of the posterior in neighborhoods of  $\tau^*$  under some topology on  $\mathcal{M}$  has been studied by Schwartz [1965], Barron et al. [1999], Miller [2021], among others. For *parametric* families of theories with parameter  $\theta \in \mathbb{R}^d$ , under smoothness and nondegeneracy assumptions, the Bernstein-von Mises theorem guarantees convergence of the posterior  $P(\theta | Z_{1:t})$  to the true parameter  $\theta^*$  at a rate that is asymptotically Gaussian with inverse covariance  $I(\theta^*)t$ , where  $I(\cdot)$  denotes the Fisher information matrix.

**On convergence rates.** While we do not handle the *rate* of convergence in Prop. 3.1, guarantees can be obtained under specific assumptions on the prior and the set of theories. For example, for any  $\tau \in \mathcal{M}$ , the quantity  $D_\tau^t := \log \frac{P(\tau^* | Z_{1:t})}{P(\tau | Z_{1:t})}$  is a process with  $D_\tau^0 = \log \frac{P(\tau^*)}{P(\tau)}$  and i.i.d. increments, with

$$\begin{aligned} \mathbb{E}[D_\tau^{t+1} - D_\tau^t] &= D_{\text{KL}}(\tau^* \parallel \tau) \\ \mathbb{E}[(D_\tau^{t+1} - D_\tau^t)^2] &= \mathbb{E}_{Z \sim \tau^*} \left[ \left( \log \frac{P_{\tau^*}(Z)}{P_\tau(Z)} \right)^2 \right]. \end{aligned} \quad (3)$$

Under the assumption that the variances are finite and uniformly bounded in  $\tau$ , the central limit theorem would give posterior convergence rate guarantees.

However, note that Prop. 3.1 is a law-of-large-numbers-like result that holds even if the variances in (3) are not finite and uniformly bounded.

**Harm probability bounds.** So far, we have considered a collection  $\mathcal{M}$  of distributions over an observation space. Now, we show bounds when each theory computes probabilities over some additional variables. The following lemma extends Prop. 3.1 (b) to estimates of real-valued functions of the theories and observations.

**Lemma 3.2.** *Under the conditions of Prop. 3.1, let  $f : \mathcal{M} \times \bigcup_{t=0}^{\infty} \mathcal{Z}^t \rightarrow \mathbb{R}_{\geq 0}$  be a bounded measurable function. Then there exists  $N \in \mathbb{N}$  such that for all  $t \geq N$  and any  $\tilde{\tau} \in \arg \max_{\tau} [\mathbb{P}(\tau \mid Z_{1:t}) f(\tau, Z_{1:t})]$ , it holds that  $f(\tau^*, Z_{1:t}) \leq f(\tilde{\tau}, Z_{1:t})$ .*

A particular case of interest is when each theory is associated with estimates of probabilities of harm ( $H = 1$ ) given a context  $x$  and past observations  $Z_{1:t}$ . That is,  $\mathcal{M}$  gives rise to a collection of conditional probability mass functions over the possible harm outcomes, denoted  $\mathbb{P}(\cdot \mid x, \tau, Z_{1:t})$ , for every  $x$  lying in some space of possible contexts. In this setting, we have the following corollary:

**Proposition 3.3 (Harm probability bound).** *Under the same conditions as Prop. 3.1, there exists  $N \in \mathbb{N}$  such that for all  $t \geq N$  and  $\tilde{\tau} \in \arg \max_{\tau} \mathbb{P}(\tau \mid Z_{1:t}) \mathbb{P}(H = 1 \mid x, \tau, Z_{1:t})$ , it holds that*

$$\mathbb{P}(H = 1 \mid x, \tau^*, Z_{1:t}) \leq \mathbb{P}(H = 1 \mid x, \tilde{\tau}, Z_{1:t}). \quad (4)$$

Intuitively, this means that once we have seen enough i.i.d. data, we can act as though the ‘most cautious yet still plausible’ model upper-bounds reality: if that model judges an action safe enough, the true world will be at least as safe.

## 4 NON-I.I.D. DATA

In this section, we remove the assumption made in §3 that the  $Z_i$ ’s are i.i.d. given a theory  $\tau^*$ .

**Setting.** As before, let  $(\mathcal{Z}, \mathcal{F}, \mu)$  be a  $\sigma$ -finite Borel measure space of observations. For the results below to hold, we must also assume that  $(\mathcal{Z}, \mathcal{F})$  is a Radon space (e.g., any countable set or manifold), so as to satisfy the conditions of the disintegration theorem.

Let  $(\mathcal{Z}^{\infty}, \mathcal{F}^{\infty}, \mu^{\infty})$  be the space of infinite sequences of observations,  $\mathcal{Z}^{\infty} = \{(z_1, z_2, \dots) : z_i \in \mathcal{Z}\}$ , with the associated product  $\sigma$ -algebra and  $\sigma$ -finite measure. This object is the projective limit of the measure spaces  $(\mathcal{Z}^t, \mathcal{F}^{\otimes t}, \mu^{\otimes t})$ , where  $\mathcal{Z}^t = \{(z_1, \dots, z_t) : z_i \in \mathcal{Z}\}$  and the projection

$\mathcal{Z}^{t+1} \rightarrow \mathcal{Z}^t$  ‘forgets’ the observation  $z_{t+1}$ . A theory  $\tau$  is a probability distribution on  $(\mathcal{Z}^{\infty}, \mathcal{F}^{\infty})$  that is absolutely continuous w.r.t.  $\mu^{\infty}$ . For  $A \in \mathcal{F}^{\otimes t}$ , we write  $\tau_{1:t}(A)$  for the measure of the cylindrical set,  $\tau(A \times \mathcal{Z} \times \mathcal{Z} \times \dots)$ , so  $\tau_{1:t}$  is a measure on  $(\mathcal{Z}^t, \mathcal{F}^{\otimes t})$ . Because  $\mathcal{F}^{\infty}$  is generated by cylindrical sets, the absolute continuity condition on  $\tau$  is equivalent to absolute continuity of  $\tau_{1:t}$  w.r.t.  $\mu^{\otimes t}$  for all  $t$ .<sup>4</sup> This condition allows to define measurable probability density functions  $\mathbb{P}_{\tau} : \mathcal{Z}^t \rightarrow \mathbb{R}_{\geq 0}$  as Radon-Nikodym derivatives, so that

$$\forall A \in \mathcal{F}^{\otimes t}, \quad \tau_{1:t}(A) = \int_{z_{1:t} \in A} \mathbb{P}_{\tau}(z_{1:t}) d\mu^{\otimes t},$$

and measurable conditional probability densities  $\mathbb{P}_{\tau}(z_{t+1} \mid z_{1:t}) := \frac{\mathbb{P}_{\tau}(z_{1:t}, z_{t+1})}{\mathbb{P}_{\tau}(z_{1:t})}$  when  $\mathbb{P}_{\tau}(z_{1:t}) > 0$ . The disintegration theorem for product measures implies that these conditionals and marginals over finitely many observations can be manipulated algebraically using the usual rules of probability for  $\mu^{\infty}$ -a.e. collection of values, e.g., one has the autoregressive decomposition  $\mathbb{P}_{\tau}(z_{1:t}) = \prod_{i=1}^t \mathbb{P}_{\tau}(z_i \mid z_{1:i-1})$ , with the conditional  $\mathbb{P}_{\tau}(z_1 \mid z_{1:0})$  understood to be the marginal  $\mathbb{P}_{\tau}(z_1)$ .

A theory is canonically associated with a  $\mathcal{Z}^{\infty}$ -valued random variable  $Z_{1:\infty}$ . We denote its components by  $Z_1, Z_2, \dots$  and the collection of the first  $t$  observations by  $Z_{1:t}$ .

**Definition of posterior as a random variable.** Let  $\mathcal{M} = (\tau_i)_{i \in I}$  be a collection of theories indexed by a countable set  $I^2$  and let  $\mathbb{P}$  be a prior distribution on  $I$ . We define the posterior over indices to be

$$\mathbb{P}(i \mid z_{1:t}) := \frac{\mathbb{P}(i) \mathbb{P}_{\tau_i}(z_{1:t})}{\sum_{j \in I} \mathbb{P}(j) \mathbb{P}_{\tau_j}(z_{1:t})}, \quad (5)$$

assuming the denominator converges to a positive value.

Consider a ground truth index  $i^* \in I$  and abbreviate  $\tau^* := \tau_{i^*}$ . Let  $Z_{1:\infty}$  be the random variable taking values in  $\mathcal{Z}^{\infty}$  corresponding to  $\tau^*$ . Similarly to the i.i.d. case, the posterior  $\mathbb{P}(\cdot \mid Z_{1:t})$  is a random variable taking values in the space of probability mass functions on  $I$ .

For all results below, we assume that  $\mathbb{P}(i^*) > 0$ .

**Bayesian posterior convergence.** Previous work (e.g., [Cohen et al., 2022a]) has shown that if  $Z_{1:\infty} \sim \tau^*$ , then the limit inferior of  $\mathbb{P}(i^* \mid Z_{1:t})$  is almost surely positive. More generally, with probability at least  $1 - \delta$ , the posterior on the truth will not asymptotically go below  $\delta$  times the prior on the truth. We repeat that result here in our notation.

<sup>4</sup>This is, in turn, equivalent to absolute continuity of conditional distributions, i.e., for every measurable subset  $A \subseteq \mathcal{Z}^t$  such that  $\tau_{1:t}(A) > 0$ ,  $\frac{1}{\tau_{1:t}(A)} \tau_{1:t+1}|_{A \times \mathcal{Z}} \ll \mu^{t+1}|_{A \times \mathcal{Z}}$ , where  $A \times \mathcal{Z} \subseteq \mathcal{Z}^t \times \mathcal{Z} \cong \mathcal{Z}^{t+1}$ .

<sup>5</sup>Unlike in §3, we no longer require theories to be distinct.

**Lemma 4.1** (Martingale). *The process  $W_t := P(i^* | Z_{1:t})^{-1}$  is a supermartingale, i.e., it does not increase over time in expectation.*

**Proposition 4.2** (Posterior on truth). *For all  $\delta > 0$ , with probability at least  $1 - \delta$ ,  $\inf_t P(i^* | Z_{1:t}) \geq \delta P(i^*)$ ; that is,  $\tau^*(\{z_{1:\infty} : \inf_t P(i^* | z_{1:t}) < \delta P(i^*)\}) \leq \delta$ , or equivalently:*

$$\tau^* \left( \sup_{t \geq 0} W_t \geq (\delta P(i^*))^{-1} \right) \leq \delta \quad (6)$$

In the language of financial markets, if  $W_t$  was the price of a stock at time  $t$ , you could never make money in expectation by holding it. Suppose that you “bought shares” at time 0, paying  $W_0$ , and waited for their value to increase by a factor of  $\delta^{-1}$ . If (6) did not hold and the probability of such an increase occurring was greater than  $\delta$ , then you could make an expected profit by “ $\delta^{-1}$ -tupling” your money with probability greater than  $\delta$ .

The bound in Prop. 4.2 is “tight”; see Remark B.2.

**Harm probability bounds.** We now state analogues of Prop. 3.3 in the non-i.i.d. setting. As above, let  $H_t$  be a binary random variable that may depend on  $Z_{1:t}$ ,  $\tau$ , and a context variable  $x_t$ .

**Proposition 4.3** (Weak harm probability bound). *For any  $\delta > 0$ , with probability at least  $1 - \delta$ , the following holds for all  $t \in \mathbb{N}$  and all  $x_t$ :*

$$P(H_t = 1 | Z_{1:t}, \tau^*, x_t) \leq \sup_{i \in I} \frac{P(i | Z_{1:t}) P(H_t = 1 | Z_{1:t}, \tau_i, x_t)}{\delta P(i^*)}.$$

Next, we show how the bound in Prop. 4.3 can be strengthened by restricting to theories that have sufficiently high posterior mass relative to theories that are “better” than them.

Let  $i_{Z_{1:t}}^1, i_{Z_{1:t}}^2, i_{Z_{1:t}}^3, \dots$  be an enumeration of  $I$  in order of decreasing posterior weight  $P(i | Z_{1:t})$ , breaking ties arbitrarily, for example, following some fixed enumeration of  $I$  (i.e., we have  $P(i_{Z_{1:t}}^n | Z_{1:t}) \geq P(i_{Z_{1:t}}^{n+1} | Z_{1:t})$  for all  $n$ ). Each  $i_{Z_{1:t}}^n$  is an  $I$ -valued random variable (i.e., the index of a theory in  $\mathcal{M}$ ). For any  $0 < \alpha \leq 1$ , we also define the  $\mathcal{P}(I)$ -valued random variable

$$\mathcal{I}_{Z_{1:t}}^\alpha := \left\{ i_{Z_{1:t}}^n \in I : P(i_{Z_{1:t}}^n | Z_{1:t}) \geq \alpha \sum_{m \leq n} P(i_{Z_{1:t}}^m | Z_{1:t}) \right\}, \quad (7)$$

which is the set of indices that contain at least  $\alpha$  of the posterior mass of all indices that are more likely than it under the posterior. If  $\alpha = 1$ , this set is the singleton  $\{i_{Z_{1:t}}^1\}$ .

For any  $0 < \alpha < 1$ , it is nonempty, because it contains  $i_{Z_{1:t}}^1$ , and finite, since  $|\mathcal{I}_{Z_{1:t}}^\alpha| \geq N$  implies (easily by induction)

that  $\sum_{i \in \mathcal{I}_{Z_{1:t}}^\alpha} P(i | Z_{1:t}) \geq \left(\frac{1}{1-\alpha}\right)^{N-1} P(i_{Z_{1:t}}^1 | Z_{1:t})$ .

The following is a variant of Cohen et al. [2022a, Thm 2].

**Proposition 4.4** (True theory dominance). *If  $\alpha < \delta P(i^*)$ , then with probability at least  $1 - \delta$ , for all  $t \in \mathbb{N}$ ,  $i^* \in \mathcal{I}_{Z_{1:t}}^\alpha$ .*

**Proposition 4.5** (Harm probability bound). *If  $\alpha < \delta P(i^*)$ , then with probability at least  $1 - \delta$ , for all  $t \in \mathbb{N}$  and all  $x_t$ ,*

$$P(H_t = 1 | Z_{1:t}, \tau^*, x_t) \leq \max_{i \in \mathcal{I}_{Z_{1:t}}^\alpha} P(H_t = 1 | Z_{1:t}, \tau_i, x_t) \quad (8)$$

Intuitively, this means that even when the data distribution shifts over time, it is enough to check a small ‘credible but cautious’ set of high-posterior theories: if all of them say the action is safe below the threshold, the real environment is very unlikely to be worse. Because the conclusion of Prop. 4.5 is much stronger than that of Prop. 4.3, it would be much safer (or more useful, depending on the value of  $\alpha$ ) to use  $\arg \max_{i \in \mathcal{I}_{Z_{1:t}}^\alpha} P(H_t = 1 | Z_{1:t-1}, \tau_i, x_t)$  as a ‘paranoid’ theory rather than  $\arg \max_{i \in I} P(\tau_i | Z_{1:t-1}) P(H_t = 1 | Z_{1:t-1}, \tau_i, x_t)$ . The factor of  $(\delta P(i^*))^{-1}$  in Prop. 4.3 could render the upper bound on harm probability much larger than the trivial upper bound of 1. However, we note that approximating  $\mathcal{I}_{Z_{1:t}}^\alpha$  – such as by amortization or by Monte Carlo methods – is much more difficult than approximating the posterior alone.

**On the harm-recording mechanism.** Suppose that  $\tau_{i^*} = \tau^*$  is a data-generating process meeting the description “ $H_t = 1$  when harm has occurred”, while  $\tau_{i^\dagger} = \tau^\dagger$  is a data-generating process meeting the description “ $H_t = 1$  when harm is recorded as having occurred” and agreeing with  $\tau^*$  in its observational predictions otherwise. If, and only if, the recording process is functioning correctly,  $\tau^* = \tau^\dagger$ . For as long as the recording process is functioning correctly,  $P(i^* | Z_{1:t})/P(i^\dagger | Z_{1:t}) = P(i^*)/P(i^\dagger)$ . If the recording process ever fails at time  $t$ , then  $Z_t \sim P_{\tau^\dagger}$ , not  $P_{\tau^*}$ , since  $Z_t$  is the result of this recording process; therefore,  $P(i^* | Z_{1:t})/P(i^\dagger | Z_{1:t})$  would *decrease* in expectation, perhaps dramatically. We should not expect  $P(i^*)$  to naturally win out over  $P(i^\dagger)$ , even if there are no mistakes when recording how harmful certain situations are. However, the following holds with probability approaching 1 as  $\alpha \rightarrow 0$ : for all  $t$ , if the recording process has not failed by time  $t$ ,  $\mathcal{I}_{Z_{1:t}}^\alpha$  contains both  $i^*$  and  $i^\dagger$ . If  $\tau^*$  considers tampering with the recording process to be a ‘harmful’ outcome, then an AI system could attempt to avoid a first instance of tampering at time  $t$ , for all  $t$ .

## 5 EXPERIMENTS

**Exploding bandit setting.** We evaluate<sup>6</sup> the performance of safety guardrails based on Prop. 3.3 and Prop. 4.5 in a bandit MDP with 10 arms (actions). Each arm  $a \in \{1, \dots, 10\}$

<sup>6</sup>Code available at <https://github.com/saifh-github/conservative-bayesian-public>

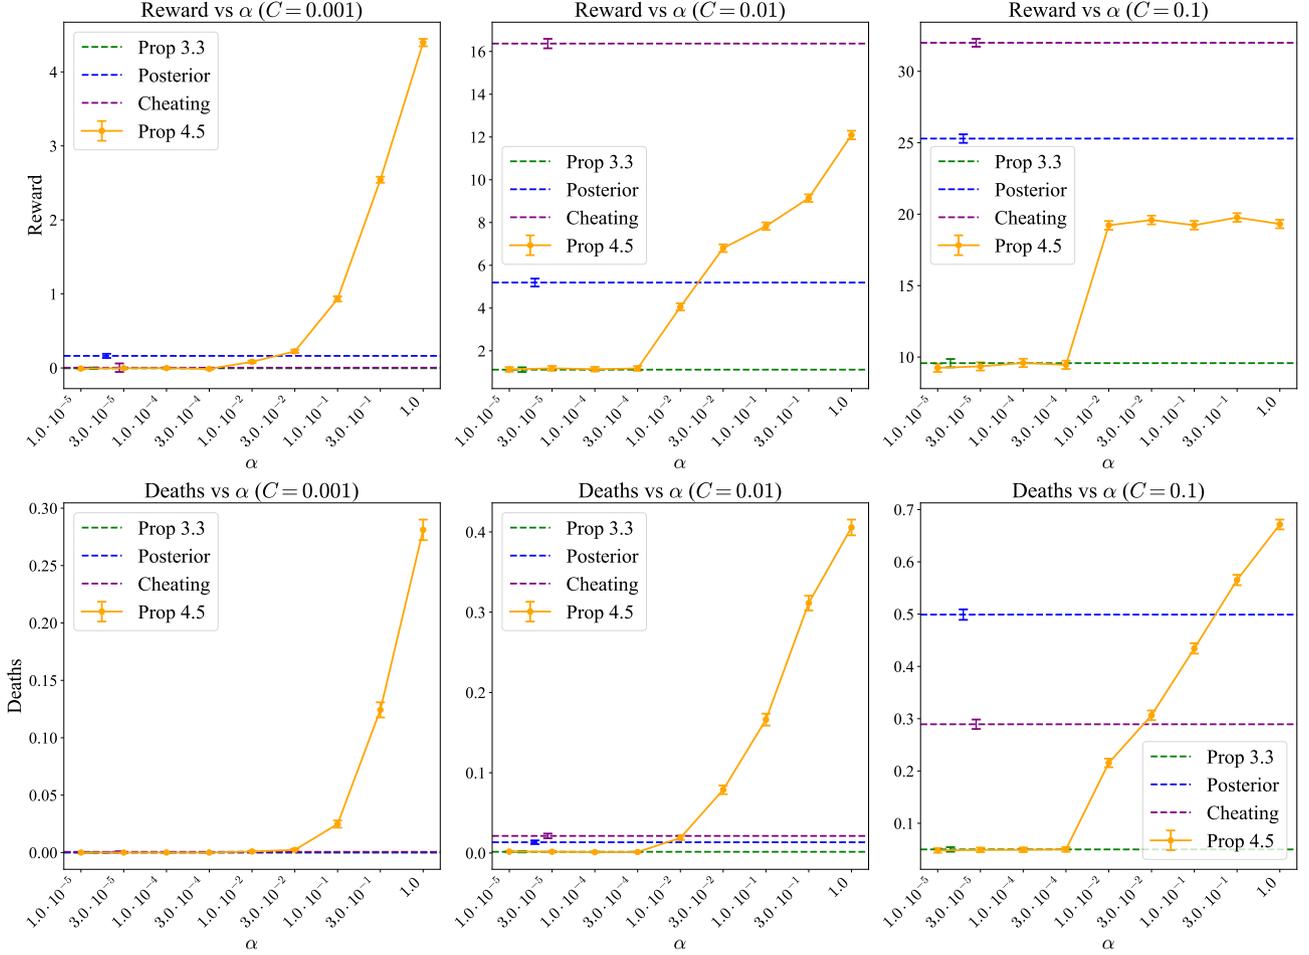


Figure 1: Mean episode deaths and reward for different guardrails in the exploding bandit setting.

is represented by a feature vector  $f_a \in \{0, 1\}^d$  (we take  $d = 10$ , but  $d$  is not necessarily equal to the number of arms), which is sampled uniformly at random at the start of each episode and known to the agent. The reward distribution of each arm is fixed for the duration of each episode and assumed to be of the following form: the reward received after taking action  $a$  follows a unit-variance normal distribution,  $r(a) \sim \mathcal{N}(f_a \cdot v^*, 1)$ , where  $v^* \in \{0, 1\}^d$  is some vector sampled uniformly at random at the start of each episode and unknown to the agent. Taking any action and observing the reward gives evidence about the identity of  $v^*$  and thus about the reward distributions of the other actions. The agent maintains a belief over the vector used to compute the reward, beginning with a uniform prior over  $\{0, 1\}^d$  and updating its posterior with each observation of an action-reward pair.

We assume that the agent samples its actions from a Boltzmann policy (with temperature 2) using the expected reward of each action under its posterior given the data seen so far, meaning that a reward vector  $v \in \{0, 1\}^d$  determines a distribution over sequences of action-reward pairs. Thus each  $v \in \{0, 1\}^d$  can be naturally associated

with a theory  $\tau_v$ <sup>7</sup>, and thus  $I := \{0, 1\}^d$  is an indexing set for a collection of theories  $\mathcal{M} = (\tau_v)_{v \in I}$ . Inference of  $v$  with evidence collected on-policy is equivalent to inference of  $\tau_v$  given data generated by a true theory  $\tau^* := \tau_{v^*}$ . Since the policy changes across timesteps, so does the distribution of action-reward pairs, so we are in the non-i.i.d. setting.

The bandit comes with a notion of harm: if the reward received at a given timestep exceeds some threshold  $E$ , the bandit explodes<sup>8</sup>, terminating the episode. In other words, we define harm as  $H_t := \mathbb{1}[R_t > E]$ , where  $R_t$  is the random variable representing the reward received when taking action  $a_t$ .  $E$  is set to a Monte Carlo approximation of the expected highest mean reward of any action (i.e.,  $\mathbb{E}[\max_a (f_a \cdot v^*)]$ ). The maximum episode length is 25 timesteps.

<sup>7</sup>The mapping  $v \mapsto \tau_v$  may not be injective. Different vectors may represent the same collection of reward distributions and therefore the same distribution over sequences of action-reward pairs.

<sup>8</sup>This emulates the important and problematic scenario where the user’s utility to be maximized, e.g., profits, conflicts with safety.

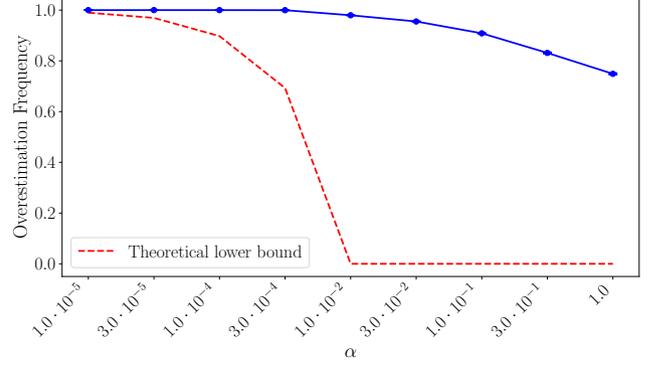
**Safety guardrails.** A *guardrail* is an algorithm that, given a possible action and context (e.g., current state and history), determines whether taking the action in the context is admissible. A guardrail can be used to mask the policy to forbid certain actions, such as those whose estimated harm exceeds some threshold  $C$ . We compare several guardrails (formally defined below): those constructed from Prop. 3.3 and Prop. 4.5, one that marginalizes across the posterior over  $\tau$  to get the posterior predictive harm probability, and one that ‘cheats’ by using the probability of harm under the true theory  $\tau^*$ . Recall that  $Z_{1:t}$  consists of the observations (i.e., actions taken and rewards received) at previous timesteps.

- **Prop. 3.3 guardrail:** rejects an action  $a_{t+1}$  if there exists  $\bar{v} \in \arg \max_v P(v | Z_{1:t}) P(H_{t+1} = 1 | \tau, Z_{1:t}, a_{t+1})$  with  $P(H_{t+1} = 1 | \tau_{\bar{v}}, Z_{1:t}, a_{t+1}) > C$  (note that the assumptions of i.i.d. observations and distinct theories do not hold).
- **Prop. 4.5 guardrail:** rejects an action  $a_{t+1}$  if  $\max_{v \in \mathcal{I}_{Z_{1:t}}^\alpha} P(H_{t+1} = 1 | Z_{1:t}, \tau_v, a_{t+1}) > C$ .
- **Posterior predictive guardrail:** rejects an action  $a_{t+1}$  if  $P(H_{t+1} = 1 | Z_{1:t}, a_{t+1}) > C$ .
- **Cheating guardrail:** rejects an action  $a_{t+1}$  if  $P(H_{t+1} = 1 | Z_{1:t}, \tau^*, a_{t+1}) > C$  (note that this guardrail assumes knowledge of the true theory  $\tau^*$ ).

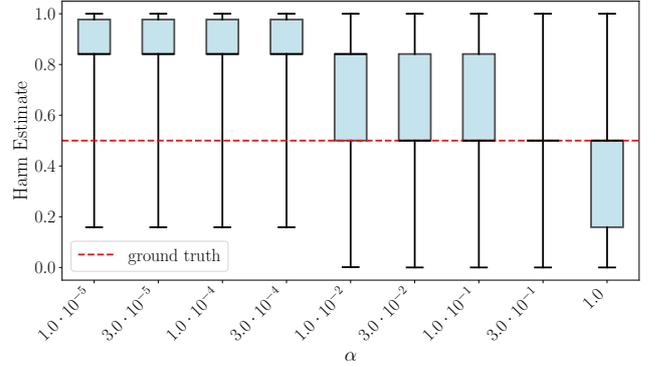
The guardrail is run at every sampling step, and actions that the guardrail rejects are forbidden to be sampled by the agent. If all actions are rejected, the episode terminates.

**Results.** Fig. 1 shows mean episode rewards and episode deaths under each guardrail across 10000 episodes, for different values of the rejection threshold  $C$ . The cheating guardrail achieves near zero deaths for sufficiently small  $C$ , but for  $C = 0.1$  its death probability is high.<sup>9</sup> The posterior predictive guardrail achieves zero deaths for small  $C$ , while for larger  $C$  it dies frequently, generally receiving lower reward compared to the cheating guardrail. The behavior of the Prop. 4.5 guardrail depends strongly on  $\alpha$ . When  $\alpha$  is close to 1, actions are rarely rejected, leading to frequent deaths. Up to a point, this riskier behavior allows the agent to get more reward, but for  $C = 0.1$  and high  $\alpha$  the trend starts to reverse, as early deaths become frequent enough to preclude the opportunity. At the other extreme, when  $\alpha$  is close to 0, the candidate set of theory indices  $\mathcal{I}_{Z_{1:t}}^\alpha$  is larger and the guardrail is extremely conservative. It rejects almost all actions, resulting in low deaths and low reward. This is the case even for larger  $C$ , since the estimated probability used to filter actions tends to overestimate an action’s harm probability under the true theory. For middling values of  $\alpha$ , Prop. 4.5 guardrail performs more similarly to the posterior predictive, sometimes with lower reward and higher deaths, and sometimes the opposite. The Prop. 3.3 guardrail, which makes the incorrect assumptions of i.i.d. data and distinct theories, is similarly conservative to the Prop. 4.5 guardrail with low  $\alpha$ .

<sup>9</sup>If every action taken had a harm probability of 0.1, the probability of death across an episode would be  $1 - ((1 - 0.1)^{25}) \approx 0.93$ .



(a) The frequency with which the inequality in Prop. 4.5 is satisfied.



(b) The distribution of the right-hand side of (8), for an action with a true harm probability of 0.5.

Figure 2: Overestimate frequency and harm estimate distribution for the Prop. 4.5 guardrail for varying  $\alpha$ .

**Tightness of bounds.** Fig. 2 shows how often and how tightly the inequality in Prop. 4.5 is satisfied. For an agent following a uniform policy across 10000 bandit episodes without action rejection or death, Fig. 2a shows the frequency with which  $\max_{v \in \mathcal{I}_{Z_{1:t}}^\alpha} P(H_{t+1} = 1 | Z_{1:t}, \tau_v, a_{t+1})$  overestimates the true harm probability. Prop. 4.5 gives a lower bound of  $1 - \frac{\alpha}{P(v^*)}$  (which may be negative) on the overestimation frequency, but the frequency significantly exceeds the bound for larger  $\alpha$ . Fig. 2b shows the distribution of harm estimates for actions with a ground truth harm probability of 0.5. For large  $\alpha$  the harm of these dangerous actions is usually *underestimated* – so the high overestimation rate (Fig. 2a) comes from less dangerous actions.

## 6 CONCLUSION AND OPEN PROBLEMS

The approach to safety verification proposed here is based on context-dependent run-time verification because the set of possible inputs for a machine learning system is generally astronomical, whereas the safety of the answer to a specific question is more likely to be tractable. It focuses on the risk of wrongly interpreting the data, including the safety specification itself (called ‘harm’ above) and exploits the fact that, as more evidence is gathered (necessary with i.i.d. data)

and when different theories predict different observations, the true interpretation rises towards the maximal value of the Bayesian posterior. The bound is tighter with the i.i.d. data, but the i.i.d. assumption is also unrealistic, and for safety-critical decisions, we would prefer to err on the side of prudence and fewer assumptions. However, it provides a template to think about variants of this idea in future work. Several challenges remain for turning such bounds into an operational run-time safeguard:

- (1) **Upper-bounding overcautiousness.** Can we ensure that we do not underestimate the probability of harm but do not massively overestimate it? Some simple theories consistent with the dataset (even an arbitrarily large one) might deem non-harmful actions harmful. Can we bound how much this harm-avoidance hampers the agent? A plausible approach would be to make use of a mentor for the agent that demonstrates non-harmful behavior [Cohen and Hutter, 2020].
- (2) **Tractability of posterior estimation.** How can we efficiently estimate the required Bayesian posteriors? For computational tractability, a plausible answer would rely on amortized inference, which turns the difficult estimation of these posteriors into the task of training a neural net probabilistic estimator which will be fast at run-time. Recent work on amortized Bayesian inference for symbolic models, such as causal structures [Deleu et al., 2022, 2023], and for intractable posteriors in language models [Guo et al., 2021, Hu et al., 2024, Venkatraman et al., 2024, Song et al., 2024, Yu et al., 2024] – which are useful when prior knowledge is encoded in a pretrained foundation model – suggests that this is feasible. Advances in efficient and adaptive Monte Carlo methods, *e.g.*, for language models [Phan et al., 2023, Zhao et al., 2024, Lew et al., 2023], can also be useful for this purpose, and MCMC approaches can complement and aid amortization [Hu et al., 2023, Kim et al., 2024, Sendera et al., 2024, Kim et al., 2025].
- (3) **Efficient search for a cautious theory.** How can we efficiently identify a cautious but plausible theory that upper-bounds the risk, since this requires an optimization at run-time? Again, a plausible answer is to rely on amortized probabilistic inference, *e.g.*, by sampling theories with a low or annealed temperature, a technique that has been used for decades in the nonconvex optimization literature and more recently combined with amortized sequential sampling methods [Zhang et al., 2023a,b, Zhou et al., 2024, *inter alia*].
- (4) **Partial specification of theories.** Another issue is that a full explanatory theory  $\tau$  for the evidence  $D$  will generally be too large an object to make the above bound calculations and optimization sufficiently efficient. An alternative would be to derive bounds that only require partial explanations, similarly to how humans tell cautious ‘stories’ about what could go wrong in uncertain

circumstances when thinking about taking a potentially dangerous decision [Bengio, 2024]. Work along these lines could allow us to avoid the requirement that we have a prior that puts positive weight on a single model that correctly models the entire model.

- (5) **Harm specification in natural language.** How can natural language safety specifications, as in Constitutional AI [Bai et al., 2022], be turned into a random variable over which conditional probabilities can be estimated? Manually labeling numerous examples matching the desired constraint is impractical. Instead, Bengio [2024] propose learning to interpret natural language statements as formal random variables over which probabilistic causal reasoning can occur. The ambiguity of natural language can then be captured by a distribution over formal interpretations.
- (6) **Effect of approximation error.** Finally, because we plan to use machine learning to estimate the bound, we have to consider the fact that a learned predictor will be imperfect. How do we conservatively correct for the potentially dangerous errors made by this predictor? This may be crucial if we use a neural network to estimate probabilities, since they are subject to lack of robustness and adversarial attacks [Goodfellow et al., 2015, Wei et al., 2023], *e.g.*, how do we ensure that the agent proposing the actions has not learned to find loopholes in the risk estimator?

## References

- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional AI: Harmlessness from AI feedback. *arXiv preprint arXiv:2212.08073*, 2022.
- Andrew Barron, Mark J. Schervish, and Larry Wasserman. The consistency of posterior distributions in nonparametric problems. *The Annals of Statistics*, 27(2):536–561, 1999.
- Sander Beckers, Hana Chockler, and Joseph Y. Halpern. Quantifying harm. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2023.
- Yoshua Bengio. Towards a cautious scientist AI with convergent safety bounds, February 2024. URL <https://yoshuabengio.org/2024/02/26/towards-a-cautious-scientist-ai-with-convergent-safety-bounds/>.
- Yoshua Bengio, Daniel Privitera, Tamay Besiroglu, Rishi Bommasani, Stephen Casper, Yejin Choi, Danielle Goldfarb, Hoda Heidari, Leila Khalatbari, Shayne Longpre,

- et al. International scientific report on the safety of advanced AI. Technical report, Department for Science, Innovation and Technology, May 2024.
- Steven Carr, Nils Jansen, Sebastian Junges, and Ufuk Topcu. Safe reinforcement learning via shielding under partial observability. *Association for the Advancement of Artificial Intelligence*, 2023.
- Michael K Cohen and Marcus Hutter. Pessimism about unknown unknowns inspires conservatism. *Conference on Learning Theory (COLT)*, 2020.
- Michael K Cohen, Marcus Hutter, and Neel Nanda. Fully general online imitation learning. *Journal of Machine Learning Research*, 23(1):15066–15095, 2022a.
- Michael K. Cohen, Marcus Hutter, and Michael A. Osborne. Advanced artificial agents intervene in the provision of reward. *AI magazine*, 43(3):282–293, 2022b.
- David Dalrymple, Joar Skalse, Yoshua Bengio, Stuart Russell, Max Tegmark, Sanjit Seshia, Steve Omohundro, Christian Szegedy, Ben Goldhaber, Nora Ammann, et al. Towards guaranteed safe AI: A framework for ensuring robust and reliable AI systems. *arXiv preprint arXiv:2405.06624*, 2024.
- Tristan Deleu, António Góis, Chris Emezue, Mansi Rankawat, Simon Lacoste-Julien, Stefan Bauer, and Yoshua Bengio. Bayesian structure learning with generative flow networks. *Uncertainty in Artificial Intelligence (UAI)*, 2022.
- Tristan Deleu, Mizu Nishikawa-Toomey, Jithendaraa Subramanian, Nikolay Malkin, Laurent Charlin, and Yoshua Bengio. Joint Bayesian inference of graphical structure and parameters with a single generative flow network. *Neural Information Processing Systems (NeurIPS)*, 2023.
- Persi Diaconis and David A. Freedman. On the consistency of Bayes estimates (with discussion). *The Annals of Statistics*, 14:1–26, 1986.
- J.L. Doob. Application of the theory of martingales. *Colloque International Centre Nat. Rech. Sci.*, pages 22–28, 1949.
- David A. Freedman. On the asymptotic behavior of Bayes’ estimates in the discrete case. *The Annals of Mathematical Statistics*, 34(4):1386–1403, 1963.
- David A. Freedman. On the asymptotic behavior of Bayes estimates in the discrete case II. *The Annals of Mathematical Statistics*, 36(2):454–456, 1965.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *International Conference on Learning Representations (ICLR)*, 2015.
- Han Guo, Bowen Tan, Zhengzhong Liu, Eric P. Xing, and Zhiting Hu. Efficient (soft) Q-learning for text generation with limited good data. *arXiv preprint arXiv:2106.07704*, 2021.
- Edward J. Hu, Nikolay Malkin, Moksh Jain, Katie Everett, Alexandros Graikos, and Yoshua Bengio. GFlowNet-EM for learning compositional latent variable models. *International Conference on Machine Learning (ICML)*, 2023.
- Edward J. Hu, Moksh Jain, Eric Elmoznino, Younesse Kaddar, Guillaume Lajoie, Yoshua Bengio, and Nikolay Malkin. Amortizing intractable inference in large language models. *International Conference on Learning Representations (ICLR)*, 2024.
- Nils Jansen, Bettina Könighofer, Sebastian Junges, Alex Serban, and Roderick Bloem. Safe Reinforcement Learning Using Probabilistic Shields. *International Conference on Concurrency Theory (CONCUR)*, 2020.
- Jacek Karwowski, Oliver Hayman, Xingjian Bai, Klaus Kiendlhofer, Charlie Griffin, and Joar Skalse. Goodhart’s law in reinforcement learning. *International Conference on Learning Representations (ICLR)*, 2024.
- Minsu Kim, Taeyoung Yun, Emmanuel Bengio, Dinghui Zhang, Yoshua Bengio, Sungsoo Ahn, and Jinkyoo Park. Local search GFlowNets. *International Conference on Learning Representations (ICLR)*, 2024.
- Minsu Kim, Sanghyeok Choi, Jiwoo Son, Hyeonah Kim, Jinkyoo Park, and Yoshua Bengio. Ant colony sampling with GFlowNets for combinatorial optimization. *Artificial Intelligence and Statistics (AISTATS)*, 2025.
- Victoria Krakovna, Jonathan Uesato, Vladimir Mikulik, Matthew Rahtz, Tom Everitt, Ramana Kumar, Zac Kenton, Jan Leike, and Shane Legg. Specification gaming: the flip side of AI ingenuity, 2020. URL [deepmind.com/blog/specification-gaming-the-flip-side-of-ai-ingenuity](https://deepmind.com/blog/specification-gaming-the-flip-side-of-ai-ingenuity).
- Alexander K Lew, Tan Zhi-Xuan, Gabriel Grand, and Vikash K Mansinghka. Sequential Monte Carlo steering of large language models using probabilistic programs. *arXiv preprint arXiv:2306.03081*, 2023.
- Alexander J McNeil, Rüdiger Frey, and Paul Embrechts. *Quantitative risk management: concepts, techniques and tools-revised edition*. Princeton university press, 2015.
- Jeffrey W. Miller. A detailed treatment of Doob’s theorem. *arXiv preprint arXiv:1801.03122*, 2018.
- Jeffrey W. Miller. Asymptotic normality, concentration, and coverage of generalized posteriors. *Journal of Machine Learning Research*, 22(168):1–53, 2021.

- Ian Osband and Benjamin Van Roy. Why is posterior sampling better than optimism for reinforcement learning? *International Conference on Machine Learning (ICML)*, 2017.
- Alexander Pan, Kush Bhatia, and Jacob Steinhardt. The effects of reward misspecification: Mapping and mitigating misaligned models. *International Conference on Learning Representations (ICLR)*, 2021.
- Richard Yuanzhe Pang, Vishakh Padmakumar, Thibault Selam, Ankur P. Parikh, and He He. Reward gaming in conditional text generation. *Association for Computational Linguistics (ACL)*, 2023.
- Du Phan, Matthew D. Hoffman, David Dohan, Sholto Douglas, Tuan Anh Le, Aaron Parisi, Pavel Sountsov, Charles Sutton, Sharad Vikram, and Rif A. Saurous. Training chain-of-thought via latent-variable inference. *Neural Information Processing Systems (NeurIPS)*, 2023.
- Jonathan Richens, Rory Beard, and Daniel H. Thompson. Counterfactual harm. *Neural Information Processing Systems (NeurIPS)*, 2022.
- Lorraine Schwartz. On Bayes procedures. *Probability Theory and Related Fields*, 4(1):10–26, 1965.
- Marcin Sendera, Minsu Kim, Sarthak Mittal, Pablo Lemos, Luca Scimeca, Jarrid Rector-Brooks, Alexandre Adam, Yoshua Bengio, and Nikolay Malkin. Improved off-policy training of diffusion samplers. *Neural Information Processing Systems (NeurIPS)*, 2024.
- Joar Skalse, Lucy Farnik, Sumeet Ramesh Motwani, Erik Jenner, Adam Gleave, and Alessandro Abate. STARC: A general framework for quantifying differences between reward functions. *International Conference on Learning Representations (ICLR)*, 2024.
- Joar Max Viktor Skalse, Nikolaus H. R. Howe, Dmitrii Krasheninnikov, and David Krueger. Defining and characterizing reward gaming. *Neural Information Processing Systems (NeurIPS)*, 2022.
- Joar Max Viktor Skalse, Matthew Farrugia-Roberts, Stuart Russell, Alessandro Abate, and Adam Gleave. Invariance in policy optimisation and partial identifiability in reward learning. *International Conference on Machine Learning (ICML)*, 2023.
- Zitao Song, Chao Yang, Chaojie Wang, Bo An, and Shuang Li. Latent logic tree extraction for event sequence explanation from LLMs. *International Conference on Machine Learning (ICML)*, 2024.
- Siddarth Venkatraman, Moksh Jain, Luca Scimeca, Minsu Kim, Marcin Sendera, Mohsin Hasan, Luke Rowe, Sarthak Mittal, Pablo Lemos, Emmanuel Bengio, Alexandre Adam, Jarrid Rector-Brooks, Yoshua Bengio, Glen Berseth, and Nikolay Malkin. Amortizing intractable inference in diffusion models for vision, language, and control. *Neural Information Processing Systems (NeurIPS)*, 2024.
- Jean Ville. *Étude critique de la notion de collectif*. 1939. URL <http://eudml.org/doc/192893>.
- Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. Jailbroken: How does LLM safety training fail? *Neural Information Processing Systems (NeurIPS)*, 2023.
- Fangxu Yu, Lai Jiang, Haoqiang Kang, Shibo Hao, and Lianhui Qin. Flow of reasoning: Efficient training of LLM policy with divergent thinking. *arXiv preprint arXiv:2406.05673*, 2024.
- David Zhang, Corrado Rainone, Markus Peschl, and Roberto Bondesan. Robust scheduling with GFlowNets. *International Conference on Learning Representations (ICLR)*, 2023a.
- Dinghui Zhang, Hanjun Dai, Nikolay Malkin, Aaron Courville, Yoshua Bengio, and Ling Pan. Let the flows tell: Solving graph combinatorial problems with GFlowNets. *Neural Information Processing Systems (NeurIPS)*, 2023b.
- Stephen Zhao, Rob Brekelmans, Alireza Makhzani, and Roger Baker Grosse. Probabilistic inference in language models via twisted sequential Monte Carlo. *International Conference on Machine Learning (ICML)*, 2024.
- Ming Yang Zhou, Zichao Yan, Elliot Layne, Nikolay Malkin, Dinghui Zhang, Moksh Jain, Mathieu Blanchette, and Yoshua Bengio. PhyloGFN: Phylogenetic inference with generative flow networks. *International Conference on Learning Representations (ICLR)*, 2024.
- Simon Zhuang and Dylan Hadfield-Menell. Consequences of misaligned AI. *Neural Information Processing Systems (NeurIPS)*, 2020.

## A PROOFS

**Proposition 3.1 (True theory dominance).** *Under the above conditions and supposing that  $P(\tau^*) > 0$ , the posterior  $P(\cdot | Z_{1:t})$  is almost surely defined for all  $n$ , and the following almost surely hold:*

- (a)  $P(\cdot | Z_{1:t}) \xrightarrow{t \rightarrow \infty} \delta_{\tau^*}$  as measures, where  $\delta_{\tau^*}$  is the Dirac measure, which assigns mass 1 to the theory  $\tau^*$  and 0 elsewhere; equivalently,  $\lim_{t \rightarrow \infty} P(\tau | Z_{1:t}) = \mathbb{1}[\tau = \tau^*]$ .
- (b) There exists  $N \in \mathbb{N}$  such that  $\arg \max_{\tau \in \mathcal{M}} P(\tau | Z_{1:t}) = \tau^*$  for all  $t \geq N$ .

*Proof of Proposition 3.1.* This is an application of Doob's posterior consistency theorem (Doob [1949]; see also Miller [2018] for a modern summary). This result, which follows from the theory of martingales, assumes that  $\tau^*$  is sampled from the prior distribution  $P(\tau)$  and the observations  $Z_i$  are defined as above. Doob's theorem states that if for every  $S \in \mathcal{F}$ , the map  $\tau \mapsto P_\tau(S)$  is measurable, then the posteriors  $P(\cdot | Z_{1:t})$  are almost surely defined and (a) holds P-almost surely with respect to the choice of  $\tau^*$ .

In our case, because  $\mathcal{M}$  is countable, the measurability condition is satisfied, showing that (a) holds for P-almost every  $\tau^* \in \mathcal{M}$ . In particular, if  $P(\tau^*) > 0$ , then (a) holds.

Finally, by (a), we have that for any  $\varepsilon > 0$ , there exists  $N$  such that for every  $t \geq N$ ,  $P(\tau^* | Z_{1:t}) > 1 - \varepsilon$ , or, equivalently,  $\sum_{\tau \neq \tau^*} P(\tau | Z_{1:t}) < \varepsilon$ , and therefore  $P(\tau | Z_{1:t}) < \varepsilon$  for all  $\tau \neq \tau^*$ . In particular, taking  $\varepsilon = 1/2$ , we get that for sufficiently large  $t$ ,  $P(\tau^* | Z_{1:t}) > P(\tau | Z_{1:t})$  for every  $\tau$ , which shows (b).  $\square$

**Lemma 3.2.** *Under the conditions of Prop. 3.1, let  $f : \mathcal{M} \times \bigcup_{t=0}^{\infty} \mathcal{Z}^t \rightarrow \mathbb{R}_{\geq 0}$  be a bounded measurable function. Then there exists  $N \in \mathbb{N}$  such that for all  $t \geq N$  and any  $\tilde{\tau} \in \arg \max_{\tau} [P(\tau | Z_{1:t}) f(\tau, Z_{1:t})]$ , it holds that  $f(\tau^*, Z_{1:t}) \leq f(\tilde{\tau}, Z_{1:t})$ .*

*Proof of Lemma 3.2.* First, note that the argmax exists by boundedness of  $f$  and  $P(\cdot | Z_{1:t})$ . By Prop. 3.1 (b), there exists  $N \in \mathbb{N}$  such that for all  $t \geq N$  and  $\tau \neq \tau^*$ ,  $P(\tau^* | Z_{1:t}) > P(\tau | Z_{1:t}) \geq 0$ . Let  $t \geq N$  and  $\tilde{\tau} \in \arg \max_{\tau} [P(\tau | Z_{1:t}) f(\tau, Z_{1:t})]$ . Then

$$P(\tau^* | Z_{1:t}) f(\tilde{\tau}, Z_{1:t}) \geq P(\tilde{\tau} | Z_{1:t}) f(\tilde{\tau}, Z_{1:t}) \geq P(\tau^* | Z_{1:t}) f(\tau^*, Z_{1:t}).$$

When  $\tilde{\tau} \neq \tau^*$ , the result follows since  $P(\tau^* | Z_{1:t}) > 0$ . The case  $\tilde{\tau} = \tau^*$  is trivial.  $\square$

**Proposition 3.3 (Harm probability bound).** *Under the same conditions as Prop. 3.1, there exists  $N \in \mathbb{N}$  such that for all  $t \geq N$  and  $\tilde{\tau} \in \arg \max_{\tau} P(\tau | Z_{1:t}) P(H = 1 | x, \tau, Z_{1:t})$ , it holds that*

$$P(H = 1 | x, \tau^*, Z_{1:t}) \leq P(H = 1 | x, \tilde{\tau}, Z_{1:t}). \quad (4)$$

*Proof of Proposition 3.3.* Apply Lemma 3.2 to the function  $f(\tau, Z_{1:t}) = P(H = 1 | x, \tau, Z_{1:t})$ .  $\square$

**Lemma 4.1 (Martingale).** *The process  $W_t := P(i^* | Z_{1:t})^{-1}$  is a supermartingale, i.e., it does not increase over time in expectation.*

*Proof of Lemma 4.1.* We have

$$\begin{aligned} \mathbb{E}_{\tau^*} [W_{t+1} | Z_{1:t} = z_{1:t}] &= \int_{\{z_{t+1} \in \mathcal{Z} : P_{\tau^*}(z_{t+1} | z_{1:t}) > 0\}} P(i^* | z_{1:t+1})^{-1} P_{\tau^*}(z_{t+1} | z_{1:t}) d\mu \\ &\stackrel{(a)}{=} \int_{\{z_{t+1} \in \mathcal{Z} : P_{\tau^*}(z_{t+1} | z_{1:t}) > 0\}} \frac{\sum_{j \in I} P(j | z_{1:t}) P_{\tau_j}(z_{t+1} | z_{1:t})}{P(i^* | z_{1:t}) P_{\tau^*}(z_{t+1} | z_{1:t})} P_{\tau^*}(z_{t+1} | z_{1:t}) d\mu \\ &\stackrel{(b)}{\leq} \int_{\mathcal{Z}} \frac{\sum_{j \in I} P(j | z_{1:t}) P_{\tau_j}(z_{t+1} | z_{1:t})}{P(i^* | z_{1:t})} d\mu \\ &= w_t \sum_{j \in I} P(j | z_{1:t}) \int_{\mathcal{Z}} P_{\tau_j}(z_{t+1} | z_{1:t}) d\mu \\ &\stackrel{(c)}{=} w_t \end{aligned} \quad (9)$$

where (a) is by the definition (5), (b) follows from cancellation and positivity of the integrand,  $w_t := P(i^* | Z_{1:t} = z_{1:t})^{-1}$  is the realization of  $W_t$ , and (c) follows because both the posterior and the conditional probability measure integrate to 1.

This holds for any  $z_{1:t}$ , so we can remove the conditional:

$$\begin{aligned}\mathbb{E}_{\tau^*}[W_{t+1} | w_t] &= \int_{z_1, \dots, z_t} \mathbb{E}_{\tau^*}[W_{t+1} | z_{1:t}, w_t] \mathbb{P}_{\tau^*}(z_1, \dots, z_t | W_t = w_t) d\mu^{\otimes t} \\ &\stackrel{(a)}{=} \int_{z_{1:t}} \mathbb{E}_{\tau^*}[W_{t+1} | z_{1:t}] \mathbb{P}_{\tau^*}(z_{1:t} | w_t) d\mu^{\otimes t} \\ &\stackrel{(b)}{\leq} \int_{z_{1:t}} w_t \mathbb{P}_{\tau^*}(z_{1:t} | w_t) d\mu^{\otimes t} \\ &= w_t\end{aligned}$$

where (a) follows because  $w_t$  is a function of  $z_{1:t}$  and (b) is Inequality 9. \(\square\)

**Proposition 4.2** (Posterior on truth). *For all  $\delta > 0$ , with probability at least  $1 - \delta$ ,  $\inf_t \mathbb{P}(i^* | Z_{1:t}) \geq \delta \mathbb{P}(i^*)$ ; that is,  $\tau^*(\{z_{1:\infty} : \inf_t \mathbb{P}(i^* | z_{1:t}) < \delta \mathbb{P}(i^*)\}) \leq \delta$ , or equivalently:*

$$\tau^*\left(\sup_{t \geq 0} W_t \geq (\delta \mathbb{P}(i^*))^{-1}\right) \leq \delta \quad (6)$$

*Proof of Proposition 4.2.* By Ville's inequality [Ville, 1939] for the supermartingale  $W_t := \mathbb{P}(i^* | Z_{1:t})^{-1}$ , for any  $\lambda > 0$ :

$$\tau^*\left(\sup_{t \geq 0} W_t \geq \lambda\right) \leq \frac{\mathbb{E}[W_0]}{\lambda} = \frac{1}{\lambda \mathbb{P}(i^*)}$$

Setting  $\lambda = (\delta \mathbb{P}(i^*))^{-1}$ , we get

$$\tau^*\left(\sup_{t \geq 0} W_t \geq (\delta \mathbb{P}(i^*))^{-1}\right) \leq \delta$$

and given that

$$\{z_{1:\infty} : \sup_{t \geq 0} w_t := \sup_{t \geq 0} \mathbb{P}(i^* | z_{1:t})^{-1} > (\delta \mathbb{P}(i^*))^{-1}\} = \{z_{1:\infty} : \inf_{t \geq 0} \mathbb{P}(i^* | z_{1:t}) < \delta \mathbb{P}(i^*)\},$$

the result follows. \(\square\)

**Proposition 4.3** (Weak harm probability bound). *For any  $\delta > 0$ , with probability at least  $1 - \delta$ , the following holds for all  $t \in \mathbb{N}$  and all  $x_t$ :*

$$\mathbb{P}(H_t = 1 | Z_{1:t}, \tau^*, x_t) \leq \sup_{i \in I} \frac{\mathbb{P}(i | Z_{1:t}) \mathbb{P}(H_t = 1 | Z_{1:t}, \tau_i, x_t)}{\delta \mathbb{P}(i^*)}.$$

*Proof of Proposition 4.3.* Substituting  $i$  for  $i^*$  on the r.h.s. can never increase the r.h.s., since  $i^* \in I$ . Then, after canceling and rearranging the terms, the proposition is readily implied by Prop. 4.2. \(\square\)

**Proposition 4.4** (True theory dominance). *If  $\alpha < \delta \mathbb{P}(i^*)$ , then with probability at least  $1 - \delta$ , for all  $t \in \mathbb{N}$ ,  $i^* \in I_{Z_{1:t}}^\alpha$ .*

*Proof of Proposition 4.4.* For any  $t \geq 1$ , by Prop. 4.2,

$$\begin{aligned}\delta &\geq \tau^*\left(\{z_{1:\infty} : \inf_{t'} \mathbb{P}(i^* | z_{1:t'}) < \delta \mathbb{P}(i^*)\}\right) \\ &\geq \tau^*(\{z_{1:\infty} : \mathbb{P}(i^* | z_{1:t}) < \delta \mathbb{P}(i^*)\}) \\ &\geq \tau^*(\{z_{1:\infty} : \mathbb{P}(i^* | z_{1:t}) < \alpha\}).\end{aligned}$$

So  $\tau^*(\{z_{1:\infty} : \mathbb{P}(i^* | z_{1:t}) \geq \alpha\}) \geq 1 - \delta$ , and the result follows by the fact that  $I_{Z_{1:t}}^\alpha \supseteq \{i \in I : \mathbb{P}(i | Z_{1:t}) \geq \alpha\}$ , since the sum in (7) never exceeds 1. \(\square\)

**Proposition 4.5** (Harm probability bound). *If  $\alpha < \delta \mathbb{P}(i^*)$ , then with probability at least  $1 - \delta$ , for all  $t \in \mathbb{N}$  and all  $x_t$ ,*

$$\mathbb{P}(H_t = 1 | Z_{1:t}, \tau^*, x_t) \leq \max_{i \in I_{Z_{1:t}}^\alpha} \mathbb{P}(H_t = 1 | Z_{1:t}, \tau_i, x_t) \quad (8)$$

*Proof of Proposition 4.5.* This follows directly from Prop. 4.4. \(\square\)

## B ACCOMPANYING EXAMPLES

**Remark B.1** (Necessity of conditions, i.i.d. case). *The assumption that the data-generating process  $\tau^*$  lies in  $\mathcal{M}$  and has positive prior mass is also necessary for convergence of the posterior. To illustrate this, we give a simple example in which the theories are Bernoulli distributions and the posterior does not converge to any distribution over  $\mathcal{M}$ .*

Take  $\mathcal{Z} = \{-1, 1\}$  and  $\mathcal{M} = \{\tau_p, \tau_{1/2}, \tau_{1-p}\}$  for some  $\frac{1}{2} < p < 1$ , where  $P_{\tau_c}(1) = c$ . Assume a prior with  $P(\tau_p) = P(\tau_{1-p}) = \frac{1}{2}$  and take the true data-generating process  $\tau^*$  to be  $\tau_{1/2}$ , which has prior mass 0. The log-ratio of posterior masses is then an unbiased random walk:

$$\log \frac{P(\tau_p | Z_{1:t})}{P(\tau_{1-p} | Z_{1:t})} = \log \frac{P_{\tau_p}(Z_{1:t})}{P_{\tau_{1-p}}(Z_{1:t})} = \left( \log \frac{p}{1-p} \right) \sum_{i=1}^t Z_i.$$

This quantity almost surely takes on arbitrarily large and small values infinitely many times. In fact, by the law of iterated logarithms, for any  $\epsilon > 0$  there are infinitely many  $t$  such that

$$\log \frac{P(\tau_p | Z_{1:t})}{P(\tau_{1-p} | Z_{1:t})} \geq (1 - \epsilon) \left( \log \frac{p}{1-p} \right) \sqrt{2t \log \log t}$$

and the same holds for  $\log \frac{P(\tau_{1-p} | Z_{1:t})}{P(\tau_p | Z_{1:t})}$ , by symmetry. In particular,  $\lim_{t \rightarrow \infty} P(\tau | Z_{1:t})$  almost surely does not exist for any  $\tau \neq \tau^*$ , and the  $\liminf$  and  $\limsup$  are almost surely 0 and 1, respectively.

**Remark B.2** (Tightness, non-i.i.d. case). *Prop. 4.2 is “tight” in the following sense: for all  $\delta, \epsilon > 0$ , there exist  $\mathcal{M}, P$ , and  $\tau_{i^*} \in \mathcal{M}$ , such that with probability at least  $\delta$ ,  $\limsup_t P(i^* | Z_{1:t}) \leq (\delta + \epsilon) P(i^*)$ .*

We construct such an example. Consider the following setting:  $\mathcal{M} = \{\tau^*, \tau'\}$  (indexed by  $I = \{i^*, i'\}$  as  $\tau_{i^*} = \tau^*, \tau_{i'} = \tau'$ ),  $\mathcal{Z} = \{0, 1\}$ , and the theories are defined by

$$\begin{aligned} P_{\tau^*}(1) &= \delta, & P_{\tau'}(1) &= 1, \\ P_{\tau_i}(1 | z_{1:t}) &= \frac{1}{2} \quad \forall i \in I, t \geq 1, z_{1:t} \in \mathcal{Z}^t. \end{aligned}$$

One has  $P(i^* | Z_1 = 1) = \frac{\delta P(i^*)}{\delta P(i^*) + P(i')}$ . Since  $\tau^*$  and  $\tau'$  give the same conditional probabilities of  $Z_t$  given  $Z_{1:t-1}$  for  $t > 1$ , one has  $P(i^* | Z_{1:t}) = P(i^* | Z_1)$ . So, for all  $t \geq 1$ ,  $P(i^* | Z_1 = 1, Z_{2:t} = z_{2:t}) < \delta(1 - P(i^*))^{-1} P(i^*)$ , and hence

$$\begin{aligned} \tau^* \left( \left\{ z_{1:\infty} : \limsup_t P(i^* | z_{1:t}) < \delta(1 - P(i^*))^{-1} P(i^*) \right\} \right) &\geq \\ \tau^* \left( \left\{ z_{1:\infty} : z_1 = 1 \right\} \right) &= P_{\tau^*}(1) = \delta. \quad (10) \end{aligned}$$

So by choosing  $P(i^*) < 1 - 1/(1 + \frac{\epsilon}{\delta})$ , so that  $\delta(1 - P(i^*))^{-1} < \delta + \epsilon$ , we get the desired property.