

000
001
002
003 ENERGY SHIELDS FOR FAIRNESS
004
005
006
007008 **Anonymous authors**
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

Paper under double-blind review

ABSTRACT

Runtime fairness is not a one-time constraint but a dynamic property evaluated over a sequence of decisions. To ensure fairness at runtime it is necessary to account for past decisions, information neglected by conventional, static classifiers. Traditional fairness shields enforce runtime fairness abruptly, by intervening *deterministically* whenever a sequence of decisions violates the target for a running fairness measure. This motivates our *main conceptual contribution*: **energy shields**. An energy shield is a novel, lightweight, adaptive controller that monitors a sequence of decisions and intervenes *probabilistically* to ensure runtime fairness smoothly, by utilizing physics-inspired energy functions to nudge the sequence towards fairness: the more unfair the decisions, the stronger the nudging force becomes. This makes energy shields the *first* fairness shields to provide both *short-term safety and long-term liveness guarantees*. Safety ensures that the running fairness measure stays within a running target interval with high probability, and liveness ensures that the limit of the fairness measure lies within the limit target interval. Intuitively, the short-term specifies the tolerated fairness values and the long-term specifies the desired fairness values. We also provide a synthesis procedure for constructing the least intrusive energy shield for a given target specification, and demonstrate its efficiency experimentally. As a sanity check for the theoretical contributions, we evaluate our energy shields against existing fairness shields through the lens of short- and long-term fairness.

1 INTRODUCTION

Algorithmic decision making is ubiquitous in modern life, from hiring and lending to online advertising. In these settings, binary decisions, e.g., approving or denying a loan or displaying one ad over another, are often made sequentially Liu et al. (2018); D’Amour et al. (2020). Much research has focused on designing *statically fair* algorithms, which ensure fairness in expectation over a fixed distribution Caton & Haas (2020). This guarantees that in the long term the sequence of decisions will be fair (as long as the distribution does not shift). However, this perspective on fairness fails to address unfair behavior that arises in the short term from the cumulative history of decisions. In particular, a statically fair classifier may produce arbitrarily biased sequences in the short term Cano et al. (2025a). For illustration, we use a simplified binary online advertisement setting.

Example 1. Two companies, A and B , bid c_A and c_B for ad space. At each time $t \in \mathbb{N}$, a user sees an ad from A ($X_t = 1$) with decision probability p , or from B ($X_t = 0$) otherwise. In the static setting, the decision maker is fair if p matches the bid ratio, i.e., $p \in c_A/(c_A + c_B) \pm \varepsilon$ for $\varepsilon > 0$. In the sequential setting, a sequence x_1, \dots, x_t of ad placements is fair at time t if the empirical running average of ads matches the bid ratio, i.e., $\mu(x_1, \dots, x_t) = \frac{1}{t} \sum_{i=1}^t x_i \in c_A/(c_A + c_B) \pm \varepsilon$ for $\varepsilon > 0$. We want $\mu(x_1, \dots, x_t)$ to be always in the interval. Furthermore, the sequence of ad placements is fair in the limit if $\lim_{t \rightarrow \infty} \mu(x_1, \dots, x_t) = p$.

Running vs. limit fairness: safety and liveness. Ex. 1 illustrates two *fairness* properties defined w.r.t. a running and a limit target interval. The short-term property requires the running target interval to be met by the running average at a finite time, and the long-term property requires the limit target to be met by the limit average of the decision sequence. Intuitively, the running target expresses what fairness values we tolerate in the short-term, and the limit target expresses what fairness values we desire in the long-term. This distinction matches safety-liveness classification of properties studied in formal verification Baier & Katoen (2008). The violation of a safety property can be determined at

054 a finite time, thus matching our short-term requirement for fairness, which we call *running fairness*.
 055 The satisfaction of a liveness property can be determined in the limit only, thus matching our long-
 056 term requirement for fairness, which we call *limit fairness*. Our objective is to ensure that a decision
 057 sequence meets given target intervals for both running and limit fairness.
 058

059 **Traditional vs. energy shielding.** To enforce dynamic fairness requirements, *runtime shielding*
 060 has emerged as a promising approach Cano et al. (2025b). Originating in formal verification, a
 061 shield is a lightweight controller that monitors a decision sequence and intervenes minimally to
 062 correct decisions in order to enforce desired fairness requirements Alshiekh et al. (2018). Unfortu-
 063 nately, existing fairness shields act deterministically, alternating abruptly between full intervention
 064 and no intervention, to enforce highly restrictive short-term fairness requirements, and they fail to
 065 provide long-term fairness guarantees altogether Cano et al. (2025b). Here, we introduce gentle,
 066 probabilistic shielding mechanisms that have both short- and long-term fairness guarantees.
 067

068 **Contributions.** Our main *conceptual contribution* is the introduction of *energy shields*, a proba-
 069 bilistic shielding framework inspired by physics-based energy functions to ensure both running and
 070 limit fairness requirements. An energy shield smoothly nudges a sequence of decisions toward fair-
 071 ness by assigning higher “energy” to unfair sequences and intervening, proportional to the energy,
 072 to enter a lower-energy state. As a consequence, our *energy shields are the first fairness shields*
 073 *capable of providing short- and long-term fairness guarantees* To demonstrate that energy shields
 074 provide **short-term fairness** guarantees, we provide exponentially decaying tail bounds on the prob-
 075 ability and expected value of running fairness violations with respect to a given target interval. To
 076 demonstrate that energy shields provide **long-term fairness** guarantees, we characterize the limit
 077 behavior of the shielded process, deriving conditions on the energy function to ensure convergence
 078 of the fairness measure to a given target. Additionally, we quantify the long-run expected cost of in-
 079 tervention and prove that steeper energy functions yield fewer violations, an important monotonicity
 080 property. We exploit the monotonicity property of energy functions to propose a **synthesis pro-
 081 cedure** that combines binary search and dynamic programming with tail bounds to find the least
 082 intrusive energy shield satisfying the desired running and limit fairness properties. We validate the
 083 effectiveness and efficiency of the synthesis procedure experimentally. Moreover, we benchmark
 084 our energy shields against the seminal energy shields by Cano et al. (2025b), empirically supporting
 085 the claim that energy shields are the first shields to provide both short- and long-term guarantees.
 086

087 As the bulk of our contribution is theoretical, we focus in the main text on presenting results and
 088 their intuitions. The complete proofs of all statements are provided in the appendix.
 089

090 2 RELATED WORK

091 Most algorithmic fairness research focuses on fairness in a static setting. This includes: measures for
 092 the fairness of a decision maker at the level of groups Feldman et al. (2015); Hardt et al. (2016) and
 093 individuals Dwork et al. (2012); pre-, in-, and post-processing techniques to synthesize fair decision
 094 maker Hardt et al. (2016); Gordaliza et al. (2019); Zafar et al. (2019); Agarwal et al. (2018); Wen
 095 et al. (2021); verification techniques to check whether static decision makers are fair Albarghouthi
 096 et al. (2017); Bastani et al. (2019); Sun et al. (2021); Ghosh et al. (2021); Meyer et al. (2021); Li et al.
 097 (2023). Among the existing techniques for static fairness, our shields could be classified as a post-
 098 processing technique. The key difference is that those methods modify decision-makers *once* before
 099 deployment, whereas our work addresses fairness *during* deployment via runtime intervention.
 100

101 We are not the first to be concerned with algorithmic fairness over a sequence of decisions Alamdari
 102 et al. (2024); Cano et al. (2025a). A large body of work focuses on detecting unfair behavior at
 103 runtime, both for individual Gupta et al. (2025) and for group fairness Albarghouthi & Vinitksy
 104 (2019); Henzinger et al. (2023a); Baumeister et al. (2025). Beyond detection, Cano et al. (2025b)
 105 is the only work enforcing fairness at runtime. Their shields adopt the sequential fairness definition
 106 from Parand et al. Alamdari et al. (2024), ensuring that a sequence of decisions will be fair with
 107 probability 1 at predefined periodic intervals. Our shields soften this condition, providing high-
 108 probability short-term guarantees and are the first to provide limit guarantees.
 109

110 Although the monitoring and enforcement of fairness has only recently emerged as a topic of inter-
 111 est, classical runtime monitoring and enforcement have long been studied in the runtime verification
 112

108 community, with monitors Stoller et al. (2011); Faymonville et al. (2017); Maler & Nickovic (2004);
 109 Donzé & Maler (2010); Bartocci et al. (2018); Baier et al. (2003) and shields Carr et al. (2023); Al-
 110 shiekh et al. (2018); Córdoba et al. (2023) developed for Linear Temporal Logic specifications.
 111 Shielding has also been explored in probabilistic settings Jansen et al. (2020); Yang et al. (2023),
 112 drawing on techniques from probabilistic model checking Katoen (2016). We build on results from
 113 stochastic approximation Borkar (2008); Karandikar & Vidyasagar (2024) to design gentler proba-
 114 bilistic shields that preserve almost sure convergence guarantees.

3 SETTING

In this section, we formally introduce the setting, the fairness properties, and the notion of a shield.

Decision process. We model the setting in Ex. 1 using a single coin with decision probability $p \in [0, 1]$. At each point in time $t \in \mathbb{N}$ the coin is tossed, resulting in a decision $x_t \in \mathbb{B}$, where $\mathbb{B} = \{0, 1\}$, which is the realization of the random variable $X_t \sim \text{Bernoulli}(p)$. Combined they generate the decision process $X = (X_t)_{t \in \mathbb{N}}$ consisting of i.i.d. Bernoulli p random variables. A realization $x = (x_t)_{t \in \mathbb{N}} \in \{0, 1\}^\omega$ of X is an infinite sequence of binary values.

Fairness. We are interested in measuring the fairness of the process defined above. The measure we use is called average outcome fairness Cano et al. (2025a). Formally, given an infinite sequence of binary decisions $x \in \{0, 1\}^\omega$, e.g., a realization of the decision process, we measure the fairness of a finite prefix $x_{1:t} = (x_1, \dots, x_t)$ as its average decision, denoted by $\mu(x_{1:t})$, or μ_t if clear from the context. We use $\mu(x)$ to denote the fairness measure of the realization in the limit, if it exists:

$$\mu(x_{1:t}) = \frac{1}{t} \sum_{i=1}^t x_i, \quad \text{and} \quad \mu(x) = \lim_{t \rightarrow \infty} \mu(x_{1:t}). \quad (1)$$

A *fairness target* is a tuple $\varphi = (\tau, \mathcal{S}, \mathcal{L})$ consisting of a burn-in time $\tau \in \mathbb{N}$, a running target $\mathcal{S} \subseteq [0, 1]$, and a limit target $\mathcal{L} \subseteq [0, 1]$. The fairness target specifies the acceptable fairness measure values at every finite time greater than the burn-in τ and at the limit. For the target to be satisfiable we require that the running and limit target intersect $\mathcal{S} \cap \mathcal{L}$ and that the burn-in is sufficiently large to account for initial high variance. Given a fairness target, a infinite sequence $x \in \mathbb{B}^\omega$ satisfies:

- *point fairness* at time $t \geq \tau$, if the fairness measure is in the running target at t , i.e., $\mu(x_{1:t}) \in \mathcal{S}$;
- *running fairness*, if point fairness is always satisfied after the burn-in τ , i.e., $\forall t \geq \tau: \mu(x_{1:t}) \in \mathcal{S}$;
- *limit fairness*, if the fairness measure is in the limit target in the limit, i.e., $\lim_{t \rightarrow \infty} \mu(x_{1:t}) \in \mathcal{L}$;
- *fairness*, if both running fairness and limit fairness are satisfied.

Example 2 (Ex. 1 cont.). Assume the bid ratio is 0.5. In the long-term we require a fairness measure of 0.5, and in the short-term we accept a tolerance of 0.1 after some burn-in τ . The corresponding fairness target is $\varphi = (\tau, [0.4, 0.6] \{ 0.5 \})$. Note, the decision X_t follows a Bernoulli distribution, thus the fairness measure $\mu(X_{1:t})$ follows a binomial distribution scaled by $1/t$. Since $(1/t) \text{Bin}(t, p) \xrightarrow{t \rightarrow \infty} p$, limit fairness requires $p = 0.5$. The probability of satisfying point fairness at t is $\mathbb{P}[\mu(X_{1:t}) \in \mathcal{S}] = \mathbb{P}[\text{Bin}(t, 0.5) \in t(0.5 \pm 0.1)] \approx \sum_{i=\lfloor 0.4t \rfloor}^{\lceil 0.6t \rceil} \binom{t}{i} p^i (1-p)^{t-i}$.

Shielding. As illustrated in Ex. 2, without control of p , the only possible intervention on the process is to overwrite individual decisions. This is called shielding. A deterministic shield for a decision process is a program with the power to flip the decisions made at runtime. Formally, a deterministic shield $\mathfrak{S}: (\mathbb{B} \times \mathbb{B})^* \times \mathbb{B} \rightarrow \mathbb{B}$ uses the history of decisions $x_1, \dots, x_t \in \mathbb{B}^t$ at time $t \in \mathbb{N}$ and the history of intervention $y_1, \dots, y_{t-1} \in \mathbb{B}^{t-1}$ to compute the next intervention y_t . The intervention indicates whether the decision x_t is flipped, i.e., if $y_t = 1$ then the shielded decision is $z_t = 1 - x_t$, otherwise the shielded decision is $z_t = x_t$. The objective of the shield is to aid the satisfaction of the fairness target, evaluated over the sequence shielded decisions $z = (z_t)_{t \in \mathbb{N}}$ with as little interference as possible. To measure this, we define the average interference cost ν_t over a sequence of interventions y_1, \dots, y_t as the average number of interventions $\nu_t = (1/t) \sum_{i=1}^t y_i$.

Example 3 (Ex. 2 cont.). A trivial shield guaranteeing running and limit fairness, ensures that the company ads alternate, i.e., the shielded decision sequence is $(01)^\omega$, thus $\mu(z_{1:2t}) = 0.5$ for all t .

162 A less restrictive shield that satisfies running fairness, but not the limit fairness, is one that only
 163 interferes if the point fairness is about to be violated. We call this shield the naive shield. For
 164 example, assume at $t = 100$ we have $\mu(z_{1:t}) = 40/100$, if $x_{t+1} = 0$, then the fairness measure
 165 would be $40/101 < 0.4$, so the shield enforces, i.e., $z_{t+1} = 1$.
 166

167 **Probabilistic shields.** The deterministic shields act aggressively and abruptly once the fairness
 168 measure is at risk of leaving the target and remains idle most of the time. This leads to two regimes,
 169 one where the shield has full control over the decision and one where it has none. We introduce
 170 probabilistic shields where the intervention is done probabilistically, allowing for a much more
 171 gentle approach to shielding. Formally, a probabilistic shield is a function $\mathfrak{S}: (\mathbb{B} \times \mathbb{B})^* \times \mathbb{B} \rightarrow [0, 1]$
 172 mapping a history of decisions $x_1, \dots, x_t \in \mathbb{B}^t$ at time $t \in \mathbb{N}$ and the history of interventions
 173 (y_1, \dots, y_{t-1}) into an nudging probability $q_t = \mathfrak{S}(x_1, y_1, \dots, y_{t-1}, x_t)$, defining the distribution
 174 from which the next intervention is sampled, i.e., $Y_t \sim \text{Bernoulli}(q_t)$.
 175

176 **Problem sketch.** We consider four processes: the decision process $X = (X_t)_{t \in \mathbb{N}}$, the intervention
 177 process $Y = (Y_t)_{t \in \mathbb{N}}$, the shielded decision process $Z = (Z_t)_{t \in \mathbb{N}}$, and the shielded fairness process
 178 $M = (M_t)_{t \in \mathbb{N}}$. The processes are defined at every time $t \in \mathbb{N}$ as follows:
 179

$$X_t \sim \text{Bernoulli}(p), \quad Y_t \sim \text{Bernoulli}(\mathfrak{S}(X_1, Y_1, \dots, Y_{t-1}, X_t)), \\ Z_t = (1 - X_t) \cdot Y_t + X_t \cdot (1 - Y_t), \quad \text{and} \quad M_t = \mu(Z_1, \dots, Z_t).$$

181 Intuitively, the shielded fairness process, i.e., the fairness measure evaluated over the shielded decision
 182 process, should satisfy a given fairness target $\varphi = (\tau, \mathcal{S}, \mathcal{L})$ with a high-probability, i.e.,
 183

$$\mathbb{P}(\forall t \geq \tau: M_t \in \mathcal{S}) \geq 1 - \delta, \quad \text{and} \quad \mathbb{P}(\lim_{t \rightarrow \infty} M_t \in \mathcal{L}) \geq 1 - \delta \quad \text{for } \delta \in (0, 1).$$

4 ENERGY-BASED SHIELDS

188 We introduce energy shields, a family of physics inspired probabilistic shields. The shields acts by
 189 computing an energy state for the process and nudging the system toward lower-energy configura-
 190 tions. The energy state is given by an energy function and the current fairness measure.
 191

192 **Energy function.** An energy function is bowl-shaped with minimum at its pivot point.
 193

194 **Definition 1.** An *energy function* with *pivoting point* $\kappa \in [0, 1]$ is a function $\zeta: [0, 1] \rightarrow [0, 1]$
 195 satisfying: (i) ζ is continuously differentiable, i.e., ζ is differentiable, and ζ' is continuous; (ii)
 $\zeta'(c) \leq 0$ for $c < \kappa$, and $\zeta'(c) \geq 0$ for $c > \kappa$; (iii) $\zeta(\kappa) = \zeta'(\kappa) = 0$, and $\zeta(c) > 0$ for $c \in \{0, 1\}$.
 196

197 **Example 4.** Two families of energy functions are even-polynomials and exponential energy func-
 198 tions (see Fig. 1a) defined, as $\zeta_{\kappa, \alpha, \beta}^{\text{Pol}}(x) = \alpha|x - \kappa|^\beta$ and $\zeta_{\kappa, \rho, \sigma}^{\text{Exp}}(x) = \rho(1 - e^{-\sigma(x - \kappa)^2})$, where $\kappa \in$
 199 $(0, 1)$, $\beta \in (1, \infty)$, $\alpha \in (0, 1/\max(\kappa, 1 - \kappa)^\beta)$, $\sigma \in (0, \infty)$, and $\rho \in (0, 1/(1 - e^{-\sigma(\min\{\kappa, 1 - \kappa\})^2}))$.
 200 The parameter ranges ensure that the energy function does not exceed 1 without clipping.
 201

202 **Energy shield.** An energy shield is defined w.r.t. an energy function $\zeta: [0, 1] \rightarrow [0, 1]$ with pivot-
 203 ing point $\kappa \in [0, 1]$. The pivot point determines the favored decision, while the energy function de-
 204 termines the nudging probability. Formally, assume we have observed the decisions x_1, \dots, x_t and
 205 accumulated the interventions y_1, \dots, y_{t-1} , which determined the shielded decisions z_1, \dots, z_{t-1} at
 206 time $t \in \mathbb{N}$. Then if $\mu_{t-1} \leq \kappa$, the shield accepts $x_t = 1$ and flips $x_t = 0$ with probability $\zeta(\mu_{t-1})$,
 207 and if $\mu_{t-1} > \kappa$, the shield accepts $x_t = 0$, and flips $x_t = 1$ with probability. This determines the
 208 distribution over the next shielded decision Z_t and fairness value M_t .
 209

210 **Claim 1** (Shielded decision process). A decision process Z shielded by an energy shield forms a
 211 sequence of Bernoulli random variables with evolving bias, i.e., $Z_t \sim \text{Bernoulli}(p_t)$. The biases are
 212 defined recursively as $p_1 = 1$ and $p_{t+1} = f(\mu_t)$ for a given history z_1, \dots, z_t , where
 213

$$f(\mu) = \begin{cases} p + (1 - p)\zeta(\mu) & \text{if } \mu \leq \kappa, \\ p \cdot (1 - \zeta(\mu)) & \text{if } \mu > \kappa. \end{cases} \quad (2)$$

214 Moreover, the resulting shielded fairness process update can be written as
 215

$$M_t = \mu_{t-1} + \frac{1}{t}(Z_t - \mu_{t-1}) \quad (\text{with } \mu_0 = 0). \quad (3)$$

216 Intuitively, Eq. 2 shows that the decision maker pulls the fairness value μ toward p , while the shield
 217 exerts an opposing pull toward κ . Stronger energy values amplify this effect, shifting the process
 218 further to κ . Eq. 3 makes this dynamic explicit: the sequence of μ_t ’s evolves as a stochastic approx-
 219 imation process drifting in the direction of $\mathbb{E}[Z_{t+1}|\mu_t] - \mu_t = f(\mu_t) - \mu_t$. At a convergence point,
 220 this drift should vanish, which occurs at a fixed point of f , i.e., a value of μ satisfying $f(\mu) = \mu$.
 221

222 5 SHORT-TERM: SAFETY GUARANTEES

224 For the short-term running fairness property we require the shielded fairness process to stay within
 225 the running target $\mathcal{S} \subseteq [0, 1]$ at all finite times after the burn-in τ (see Ex. 2). To prove that our
 226 shield satisfies running fairness with high-probability we develop upper bounds on the probability
 227 and the expected number of point fairness violations over an interval.

228 **Definition 2.** Let M be the shielded fairness process generated by the decision probability p and
 229 an energy shield with energy function ζ . For a running target interval $\mathcal{S} \subseteq [0, 1]$, we define two
 230 violation measures, the probability of violating point fairness $\mathcal{P}_{\mathcal{S}}$ and the expected number of point
 231 fairness violations $\mathcal{E}_{\mathcal{S}}$ over a time interval $[T, T'] \subseteq \mathbb{N} \cup \{\infty\}$, respectively, as
 232

$$233 \mathcal{P}_{\mathcal{S}}(M_{T:T'}) = \mathbb{P}[\exists t \in [T, T'] : M_t \notin \mathcal{S}] \quad \text{and} \quad \mathcal{E}_{\mathcal{S}}(M_{T:T'}) = \sum_{t=T}^{T'} \mathbb{E}[\mathbf{1}[M_t \notin \mathcal{S}]].$$

236 5.1 UPPER-BOUNDS

238 Our main result is a bound for the probability of the shielded fairness process violating point fairness
 239 at time T . This result follows from constructing a martingale that upper-bounds the distance of the
 240 process to its converging point μ^* , and using Azuma-Hoeffding’s inequality on said martingale.

241 **Theorem 1.** Let $I = [L, U]$ such that $\kappa, p \in I$. Let $\tau = 4/\min(|L - \mu^*|, |U - \mu^*|)$, $K = (1/32) \cdot 4^{\beta}$,
 242 and $\beta = \sup_{r \in [0, 1]} f'(r)$, then for every $t \geq \tau$ we have

$$244 \mathbb{P}(M_t \notin I) \leq \exp(-Kt|L - \mu^*|^2) + \exp(-Kt|U - \mu^*|^2). \quad (4)$$

245 Note that, by Eq. 2, f is always non-increasing (see Fig. 1a), implying that $\beta \leq 0$ and $K \leq 1/32$.

247 **Property bounds.** Utilizing Theorem 1, the expected number of violations and the probability of
 248 a single violation in an interval, follows from Boole’s inequality.

249 **Corollary 1.** Let $r_B = e^{-K|B - \mu^*|^2}$ for $B \in \{L, U\}$. For every $T, T' \in \mathbb{N}$ s.t. $\tau \leq T < T'$, then

$$252 \mathcal{E}_{\mathcal{S}}(M_{T:T'}) \leq \sum_{t=T}^{T'} (r_L^t + r_U^t), \quad \text{and} \quad \mathcal{E}_{\mathcal{S}}(M_{T:\infty}) \leq \frac{r_L^T}{1 - r_L} + \frac{r_U^T}{1 - r_U}.$$

254 Hence, it follows that $\mathcal{P}_{\mathcal{S}}(M_{T:T'}) \leq \mathcal{E}_{\mathcal{S}}(M_{T:T'})$ for $T' \in \mathbb{N} \cup \{\infty\}$ s.t. $T' \geq T$.

256 5.2 MONOTONICITY WITH RESPECT TO THE ENERGY FUNCTION

258 In this section, we define a partial order among energy functions given by their steepness, and show
 259 how our violation properties are monotone with respect to steepness.

261 **Definition 3.** Let ζ_1, ζ_2 be two energy functions. We say that ζ_1 is *steeper* than ζ_2 , and denote it by
 262 $\zeta_1 \succeq \zeta_2$ if $\zeta_1(y) \geq \zeta_2(y)$ for all $y \in [0, 1]$.

263 Intuitively, a steeper energy function (see Fig. 1a) constitutes a more “aggressive” shield, leading to
 264 fewer safety violations, both in probability and in expectation.

265 **Theorem 2.** Let $\zeta_1 \succeq \zeta_2$ be two energy functions, with a common minimum at κ . Let $\mathcal{S} = [L, U]$.
 266 Let M^{ζ_1} and M^{ζ_2} be the shielded fairness process generated by enforcing the decision process of
 267 $p \in [0, 1]$ with ζ_1 and ζ_2 , respectively. Let $\tau = \lceil \max\{1/|\kappa - L|, 1/|\kappa - U|\} \rceil$, for all $T \in \mathbb{N}$ and
 268 $T' \in \mathbb{N} \cup \{\infty\}$ such that $T < T'$ we have

$$269 \mathcal{E}_{\mathcal{S}}(M_{T:T'}^{\zeta_1}) \leq \mathcal{E}_{\mathcal{S}}(M_{T:T'}^{\zeta_2}), \quad \text{and} \quad \mathcal{P}_{\mathcal{S}}(M_{T:T'}^{\zeta_1}) \leq \mathcal{P}_{\mathcal{S}}(M_{T:T'}^{\zeta_2}).$$

270 **6 LONG-TERM: LIVENESS GUARANTEES**
271272 For the long-term limit fairness property, we require the shielded fairness process to converge to a
273 point in the limit target $\mathcal{L} \subseteq [0, 1]$. To prove that our shield satisfies limit fairness, we identify the
274 conditions on the energy function to ensure that the shielded fairness process converges to a target
275 value μ^* under the decision probability p with an expected average interference cost of $|p - \mu^*|$.
276277 **Main result.** It is remarkable that both the convergence of M_t and the expected cost depend *only*
278 on p and the value of ζ at the target point μ^* and not on the pivot point κ .279 **Theorem 3.** *Let $\mu^* \in [0, 1]$. Given the shielded decision process from Eq. 3, with bias p and energy
280 function ζ , the shielded fairness process $(M_t)_{t \in \mathbb{N}}$ converges almost surely (a.s.) to μ^* if and only if*
281

282
$$\zeta(\mu^*) = \begin{cases} (\mu^* - p)/(1 - p) & \text{if } p < \mu^*, \\ (p - \mu^*)/p & \text{otherwise} \end{cases} . \quad (5)$$

283

284 Furthermore, the expected interference cost $(\mathbb{E}[\nu_t])_{t \in \mathbb{N}}$ converges almost surely to $|\mu^* - p|$.
285286 **Proof intuition.** We establish that the shielded fairness process $M = (M_t)_{t \in \mathbb{N}}$ converges at the
287 fixpoint of f . First, we show f has indeed a unique fixed point located between p and the pivot κ .
288289 **Lemma 1.** *The function $f: [0, 1] \rightarrow [0, 1]$ defined as in Eq. 2 is continuously differentiable, and has
290 a unique point $\mu^* \in [0, 1]$ such that $f(\mu^*) = \mu^*$. Furthermore, μ^* sits between p and κ .*291 Once we know f has a unique fixed point, we use stochastic approximation theory to prove that the
292 shielded fairness process M converges almost surely to the fixed point.
293294 **Lemma 2.** *The fairness process, defined in Eq. 3, converges a.s. to the unique fixpoint μ^* of f , as
295 defined in Eq. 2. The error $(M_t - \mu^*)^2$ converges a.s. at the rate of $o(1/t^\lambda)$ for all $\lambda \in (0, 1)$.*
296297 *Proof sketch.* Let $g(x) = f(x) - x$, then we rewrite the update rule in Eq. 3 as
298

299
$$M_{t+1} = M_t + \gamma_t(g(M_t) + \xi_{t+1}), \quad (6)$$

300

301 where $\xi_t = Z_t - f(M_{t-1})$, and $\gamma_t = 1/t$. This is a classical Robbins-Monro form for stochastic
302 approximation Borkar (2008). From stochastic approximation theory we know that, under certain
303 regularity conditions, M_t from Eq. 6 approximates the zero value of g , which is fixed point of f . We
304 use Karandikar & Vidyasagar (2024) to bound the convergence rate. \square
305306 Equation 5 in Theorem 3 follows from Lemma 2 by noticing that the fixed point μ^* has to satisfy
307 Eq. 7. Since the fixed point lies between p and κ , the branch in Eq. 7 is chosen based on $\mu^* \leq p$.
308

309
$$\mu^* = \begin{cases} p + (1 - p) \cdot \zeta(\mu^*) & \text{if } \mu \leq \kappa, \\ p \cdot (1 - \zeta(\mu^*)) & \text{if } \mu > \kappa \end{cases} . \quad (7)$$

310

311 We show that the expected cost of intervention at step $t + 1$ is the probability of not seeing the
312 favorable decision (1 when the current average is below κ and 0 otherwise) and having an energy
313 high enough to intervene. The expected intervention cost $h(\mu_t)$ converges $h(\mu_t) \rightarrow h(\mu^*)$, because
314 the fairness value converges $\mu_t \rightarrow \mu^*$, where h is defined as
315

316
$$h(\mu) = \begin{cases} (1 - p) \cdot \zeta(\mu) & \text{if } \mu \leq \kappa \\ p \cdot \zeta(\mu) & \text{otherwise} \end{cases} . \quad (8)$$

317

318 **Lemma 3.** *For the process described in Eq. 3, the corresponding sequence of average interference
319 (ν_t) $_{t \in \mathbb{N}}$ converges to $h(\mu^*)$, where μ^* is the fixpoint of f (Eq. 2) and h is as defined in Eq. 8.*
320321 Finally, Eq. 5 and Lemma 3 imply that the expected intervention cost converges to $|p - \mu^*|$.
322323 **7 SHIELD SYNTHESIS**324 In this section, we state the energy shield synthesis problem and propose Alg. 1 as a solution.
325

Problem statement. A problem instance $(\mathcal{V}, \Xi, p, \varphi, \delta, \varepsilon)$ consists of: a violation measure $\mathcal{V} \in \{\mathcal{P}, \mathcal{E}\}$, a totally ordered set of energy functions $\Xi = \{\zeta_k\}_{k \in R}$ indexed by some interval $R \subset \mathbb{R}$ where for all $i \leq j \in R$ we have $\zeta_i \preceq \zeta_j$, a decision probability $p \in [0, 1]$; a fairness target $\varphi = (\tau, \mathcal{S}, \mathcal{L})$, a violation threshold $\delta > 0$, and an approximation tolerance $\varepsilon > 0$. Given a problem instance, find the least invasive energy shield that satisfies the fairness violation constraint.

Problem 1. Given $(\mathcal{V}, \Xi, p, \varphi, \delta, \varepsilon)$ find an energy function $\zeta \in \Xi$ such that:

(i) the shielded fairness process M^ζ with parameter p satisfies

$$\mathcal{V}_S \zeta(M_{\tau:\infty}^\zeta) \leq \delta \quad \text{and} \quad \lim_{t \rightarrow \infty} M_t^\zeta \in \mathcal{L} \quad \text{almost surely, and} \quad (9)$$

(ii) ζ is ε -minimal, i.e., if $\zeta^* \in \Xi$ is the smallest valid element, then $|\mathcal{V}_S(M_{\tau:\infty}^{\zeta^*}) - \mathcal{V}_S(M_{\tau:\infty}^\zeta)| \leq \varepsilon$.

We remark that if the violation measure is \mathcal{P} , then running fairness should be satisfied with probability greater than $1 - \delta$, and if it is \mathcal{E} , then the total number of point fairness violations after τ should be bounded by δ . In Alg. 1 we show a synthesis procedure based on having a family of energy functions $\Xi = (\zeta_r)_{r \in R}$ that is indexed by some interval $R \subset \mathbb{R}$ and is monotonic with respect to the index. In Appendix B we describe such a family, denoted $(\zeta_{r;p,\mathcal{S},\mathcal{L}}^{\text{Mon}})_{r \in (0,1)}$, which is defined w.r.t. a decision probability p and a specification φ , as a piecewise exponential and polynomial function.

Algorithm 1 Shield synthesis

Require: problem instance $(\mathcal{V}, \Xi, p, \varphi, \delta, \varepsilon)$.

- 1: $l \leftarrow \min R; u \leftarrow \max R$ \triangleright Set lower and upper bound for energy function w.r.t. \prec .
- 2: $T_{\text{DP}} \leftarrow \min\{t \in \mathbb{N} \mid \frac{r_-^t}{1-r_-} + \frac{r_+^t}{1-r_+} \leq \varepsilon\}$ \triangleright smallest t s.t. bound (Cor. 1) satisfies tolerance
- 3: **if** $\text{CONDITION}(\zeta_u, T_{\text{DP}}) > \delta$ **then return** FAIL \triangleright The most strict energy function is not enough
- 4: **while** $l \neq u$ **do**
- 5: $m \leftarrow (l + u)/2$; $d \leftarrow \text{CONDITION}(\zeta_m, T_{\text{DP}})$
- 6: **if** $|d - \delta| < \varepsilon$ **then return** ζ_m
- 7: **if** $d \leq \delta$ **then** $l \leftarrow m$ **else** $u \leftarrow m$

Algorithm. The algorithm exploits the monotonicity of the energy function family to find the least steep energy function that satisfies the violation condition, which determines the least invasive shield. Inside of the binary-search it is necessary to approximate the violation measure \mathcal{V} . The approximation algorithm, in CONDITION , uses dynamic programming and the tail-bounds from Section 5. Concretely, we divide the interval $[\tau, \infty)$ into two intervals: a prefix interval $[\tau, T_{\text{DP}}]$ and tail interval (T_{DP}, ∞) . For the prefix, we compute the exact violation measure with standard dynamic programming techniques. For the tail term, we use the bound in Sec. 5.1. Since the bounds for a violation in the tail can be made arbitrarily small for large enough T_{DP} , we use them to approximate the exact violation value with as much precision as required—a time-precision trade-off.

8 UNKNOWN AND NON-STATIONARY INPUT DISTRIBUTIONS

Until now we have assumed that the bias of the system p is fixed and unknown. In this section, we explore what results are possible when relaxing said conditions.

Setting 1: p is fixed but unknown. This corresponds to shielding a sequence sampled from a fixed distribution, with unknown bias p . In this case, we can use the natural estimator $\hat{p}_t = (1/t) \sum_i x_i$ as a replacement for p and update our energy function at each step to reflect the current estimation of. Since \hat{p}_t converges to p a.s., the energy function we use also converges, and since it is designed to make the shielded process converge to a certain target, we retain the long-term a.s. guarantees.

Theorem 4. Let $\mu^* \in [0, 1]$ be a fairness target. Let $(\zeta_q)_{q \in [0,1]}$ be a family of energy functions satisfying

$$\zeta_q(\mu^*) = \begin{cases} (\mu^* - q)/(1 - q) & \text{if } q < \mu^*, \\ (q - \mu^*)/q & \text{otherwise.} \end{cases} \quad (10)$$

Then the shielded process that takes at each step the energy function $\zeta_{\hat{p}_t}$ converges a.s. to μ^* .

378 **Setting 2: $(p_t)_{t \in \mathbb{N}}$ is varying and unknown.** If $(p_t)_{t \in \mathbb{N}}$ is allowed to evolve arbitrarily, a non-
 379 trivial shield cannot guarantee convergence for the shielded process to a target point. The best effort
 380 solution we can give is that a fixed energy function, if it is steep enough, can guarantee that the
 381 process stays within a certain target interval almost surely.

382 **Theorem 5.** Let $\zeta: [0, 1] \rightarrow [0, 1]$ be an energy function, and L, R the unique point satisfying
 383

$$384 \quad L < \kappa \text{ and } \zeta(L) = L, \quad R > \kappa \text{ and } \zeta(R) = 1 - R.$$

385 Then, for the shielded process $(M_t)_{t \in \mathbb{N}}$ it holds almost surely that

$$387 \quad \liminf_{t \rightarrow \infty} M_t \geq L, \quad \text{and} \quad \limsup_{t \rightarrow \infty} M_t \leq R.$$

390 9 ENERGY SHIELDS FOR GROUP FAIRNESS

392 Group fairness formalizes equal treatment of demographic groups in binary decision settings
 393 Mehrabi et al. (2021). Group fairness metrics typically compare decision probabilities across
 394 groups. For example, demographic parity ensures that the ratio of positive decisions is independent
 395 of the group. With a tolerance of $\varepsilon > 0$, this means

$$397 \quad \mathbb{P}(X = 1 \mid \text{Group} = A) - \mathbb{P}(X = 1 \mid \text{Group} = B) \in [-\varepsilon, +\varepsilon]. \quad (11)$$

398 Probabilities are taken w.r.t. the population distribution and a decision maker.

400 In situations where the inputs come in pairs and we are forced to give a positive decision to one
 401 group and a negative decision to the other group, the fairness-relevant part of the sequence can be
 402 modeled with the stochastic processes described in Sec. 3. This is illustrated in Ex. 1.

404 **Two-group setting.** A more common sequential input setting is that at each time an instance of
 405 either group A or group B is presented, and the decision can be either positive or negative for
 406 that instance. To model group fairness in such problems, we need to extend our input space X to
 407 be $\mathbb{G} \times \mathbb{B}$, where $\mathbb{G} = \{A, B\}$. At each point in time, we obtain an input (g_t, x_t) , where $g_t \in$
 408 $\{A, B\}$ indicates whether the input belongs to the demographic group A or B , and $x_t \in \{0, 1\}$
 409 indicates whether the initial decision is positive or not. The value of g_t is sampled from a distribution
 410 Bernoulli(π), where π indicates the probability of seeing an instance of group A . The probability
 411 of a positive decision x_t is group dependent, i.e., x_t is sampled from a distribution Bernoulli(p_{g_t}),
 412 where p_A and p_B indicate the positive decision probability for each group, respectively. The values
 413 of $\pi \in (0, 1)$, and $p_A, p_B \in [0, 1]$ are the parameters of the setting.

414 The shield can modify the decision, but not the demographic group. Therefore a (probabilistic)
 415 shield is now a function $\mathfrak{S}: (\mathbb{G} \times \mathbb{B} \times \mathbb{B})^* \times (\mathbb{G} \times \mathbb{B}) \rightarrow [0, 1]$ mapping a history $(g_i, x_i, y_i)_{i=1}^{t-1}$ and
 416 an input $(g_t, x_t) \in \mathbb{G} \times \mathbb{B}$ to a nudging probability q_t , defining the distribution from which the next
 417 intervention is samples, i.e., $Y_t = \text{Bernoulli}(q_t)$.

418 **Updated problem sketch.** We consider the group process $G = (G_t)_{t \in \mathbb{N}}$, the decision process
 419 $X = (X_t)_{t \in \mathbb{N}}$ the intervention process $Y = (Y_t)_{t \in \mathbb{N}}$, the shielded input process $Z = (Z_t)_{t \in \mathbb{N}}$, and
 420 the shielded fairness process $M = (M_t)_{t \in \mathbb{N}}$, defined as follows:

$$422 \quad G_t \sim \text{Bernoulli}(\pi), \quad X_t \sim \text{Bernoulli}(p_{G_t}), \quad Y_t = \text{Bernoulli}(\mathfrak{S}(G_1, X_1, Y_1, \dots, Y_{t-1}, G_t, X_t))$$

$$423 \quad Z_t = (1 - X_t) \cdot Y_t + X_t \cdot (1 - Y_t), \quad M_t = \frac{\sum_{i=1}^t Z_i \cdot \mathbf{1}[G_i = A]}{\sum_{i=1}^t \mathbf{1}[G_i = A]} - \frac{\sum_{i=1}^t Z_i \cdot \mathbf{1}[G_i = B]}{\sum_{i=1}^t \mathbf{1}[G_i = B]}$$

427 Given a fairness target $\varphi = (\tau, \mathcal{S}, \mathcal{L})$ and a probability $\delta \in [0, 1]$ the goal of the shield is to guarantee

$$429 \quad \mathbb{P}(\forall t \geq \tau: M_t \in \mathcal{S}) \geq 1 - \delta, \quad \text{and} \quad \mathbb{P}(\lim_{t \rightarrow \infty} M_t \in \mathcal{L}) \geq 1 - \delta.$$

431 Note that now the fairness measure M_t can take values in the interval $[-1, +1]$, so we can expect
 432 the fairness targets to reflect that, i.e., in general $\mathcal{S}, \mathcal{L} \subseteq [-1, +1]$.

432 **Implementation of the energy shield.** In this setting, an energy function follows the same def-
 433inition as Def. 1, with the only modification that the domain changes from $[0, 1]$ to $[-1, +1]$, so
 434 $\zeta: [-1, +1] \rightarrow [0, 1]$.
 435

436 The shield monitors the evolution of the shielded fairness process M . At each time t , the current
 437 fairness value is μ_t . If $\mu_t > \kappa$, the shield favours the decision that tends to decrease μ_t , and if
 438 $\mu_t < \kappa$, the shield favours the decision that tends to increase μ_t . When the input is $(g_t = A, x_t, y_t)$,
 439 forcing $z_t = 1$ increases the value of μ_t , and forcing $z_t = 0$ decreases it. On the other hand, when
 440 $g_t = B$, forcing $z_t = 1$ decreases the value of μ_t , and forcing $z_t = 0$ increases it.
 441

442 The shield is implemented with the same rationale as in the setting with a single process: if the
 443 proposed decision agrees with the direction favoured by the shield, the shield accepts it. Otherwise,
 444 the shield may flip the decision with a probability given by $\zeta(\mu_t)$.
 445

446 In this setting, we obtain similar long-term guarantees as with Thm. 3.
 447

448 **Theorem 6.** *Let $\mu^* \in [-1, +1]$. The two-group shielded fairness process $(M_t)_{t \in \mathbb{N}}$ converges a.s.
 449 to μ^* if and only if*

$$\zeta(\mu^*) = \begin{cases} (\mu^* - (p_A - p_B))/(1 - (p_A - p_B)) & \text{if } p_A - p_B < \mu^*, \\ (p_A - p_B - \mu^*)/(p_A - p_B) & \text{otherwise.} \end{cases} \quad (12)$$

450 The expected interference cost $(\mathbb{E}[\nu_t])_{t \in \mathbb{N}}$ converges a.s. to a value smaller than $|\mu^* - (p_A - p_B)|$.
 451

452 10 EXPERIMENTS

453 We evaluate the impact of energy function on the fairness measure in the setting of Ex. 2, explore
 454 the time-precision trade-off in Alg. 1, and compare against existing fairness shields.
 455

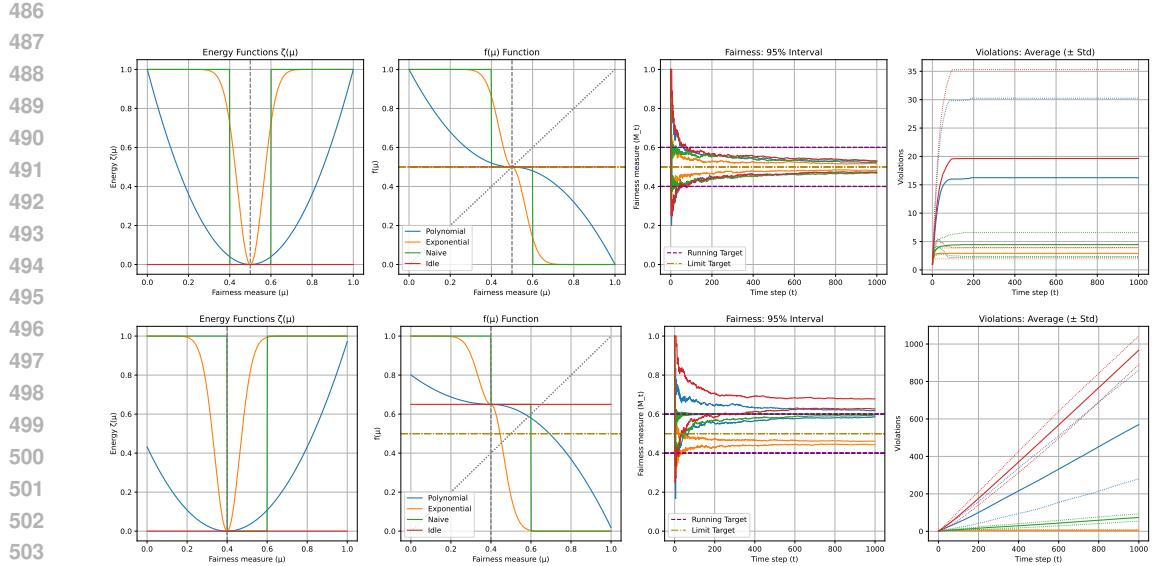
456 **Energy functions.** We consider a fair ($p = 0.5$) and biased ($p = 0.65$) decision maker with
 457 fairness target $\varphi = (\tau, [0.4, 0.6], \{0.5\})$. In addition to considering polynomial ζ^{Pol} and exponential
 458 ζ^{Exp} energy functions (see Ex. 4), we consider ζ^{Idle} , which is always 0, and ζ^{Naive} , which is 0 in
 459 the interior of the running target and 1 elsewhere. In Fig. 1a we can observe that: if the fixpoint
 460 of the functions—indicated by the intersections of 45°-line with the value graph of f —is contained
 461 within the running target, the number of violations decays quickly; if the decision maker is biased,
 462 a pivot within a given target interval does not guarantee the fairness value convergence to a point in
 463 the target. if the decision maker is fair, our shields remain useful by further reducing violations.
 464

465 **Approximation precision.** The synthesis Algorithm 1 performs a violation probability approxi-
 466 mation (VPA) by running dynamic programming (DP) up to a horizon T_{DP} and then utilizing the
 467 infinite horizon tail-bound from Sec. 5. Fig. 1b investigates the precision-time trade-off. We observe
 468 that an extension of the DP horizon T_{DP} leads to an exponential gain in VPA precision, at the cost
 469 of a quadratic increase in computation time. This impacts the runtime of Alg. 1, as higher precision
 470 requirements demands longer DP horizons. The culprit is the looseness of the tail-bounds over the
 471 short horizon, as demonstrated by the gain in precision, if larger burn-ins are considered.
 472

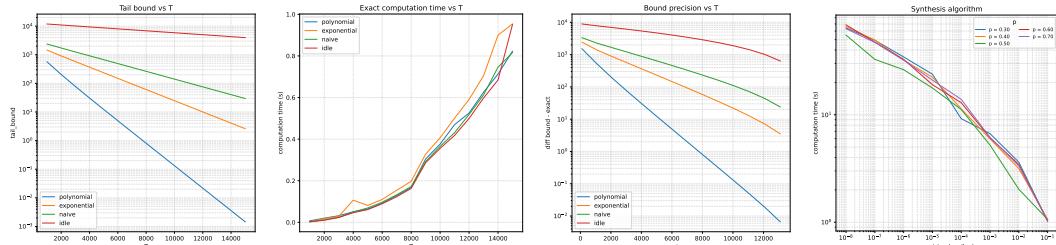
473 **Comparison.** We benchmark our energy shields w.r.t. $\zeta_{0,1;p,\mathcal{S},\mathcal{L}}^{\text{Mon}}$ against two baselines: the naive
 474 shield (Ex. 3), which enforces running fairness strictly, and periodic shields Cano et al. (2025b),
 475 which requires heavy-computation at runtime to enforce point fairness periodically with optimal
 476 expected cost. For a fairness target $\varphi = (100, [0.4, 0.6], [0.49, 0.51])$ and decision probability $p =$
 477 0.3, our shield (tuned to $\mu^* = 0.5$) achieves both running and limit fairness, unlike baselines, which
 478 do not support long-term guarantees. When configured for running fairness, baselines cluster at
 479 target boundaries without converging; when tuned for limit fairness, periodic shields over-intervene
 480 and naive shield are almost maximally invasive. This is reflected in the intervention costs: baselines
 481 prioritizing running fairness intervene less, while those targeting limit fairness intervene more.
 482

483 11 DISCUSSION

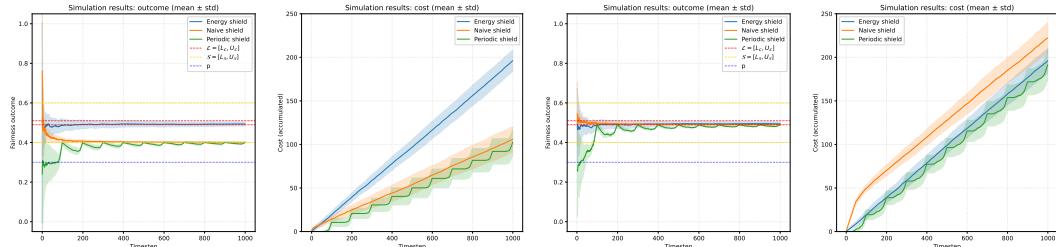
484 **Safety and liveness.** Formal verification classifies properties over infinite traces as safety or live-
 485 ness Lampert (1977); Henzinger et al. (2023b). A safety property asserts that “bad” things never



(a) Simulated impact of the energy functions ζ for fairness target $\varphi = (0, [0.4, 0.6], \{0.5\})$, decision probability p . **Rows (R):** (R1) $p = 0.5$ and $\zeta \in \{\zeta^{\text{Naive}}, \zeta^{\text{Idle}}, \zeta^{\text{Pol}}, \zeta^{\text{Exp}}_{0.5,4,2}, \zeta^{\text{Exp}}_{0.5,1,128}\}$; (R2) considers $p = 0.65$ and $\zeta \in \{\zeta^{\text{Naive}}, \zeta^{\text{Idle}}, \zeta^{\text{Pol}}, \zeta^{\text{Exp}}_{0.4,2,7,2}, \zeta^{\text{Exp}}_{0.4,1,128}\}$. **Columns (C):** (C1) the behavior of the energy function for each fairness value. (C2) illustration of the characteristic functions f and their respective fixpoints $f(\mu) = \mu$. (C3) 95% confidence interval of fairness values at each time. (C4) Average point fairness violations with standard deviation at each time. The simulation results are averaged over 1000 simulations for each p and ζ .



(b) Time-precision trade-off in the violation probability approximation (VPA) of Alg. 1 with fairness target $\varphi = (100, [0.3, 0.7], \{0.45, 0.55\})$, decision probability $p = 0.65$, and energy function $\zeta \in \{\zeta^{\text{Naive}}, \zeta^{\text{Idle}}, \zeta^{\text{Pol}}, \zeta^{\text{Exp}}_{0.4,2,7,2}, \zeta^{\text{Exp}}_{0.4,1,128}\}$. **Columns (C):** (C1) VPA with increasing until dynamic programming (DP) threshold T_{DP} . (C2) Computation time as T_{DP} increases. (C3) VPA precision, i.e., VPA with $T_{\text{DP}} = 0$ compared to $T_{\text{DP}} = 15000$, for increasing burn-ins τ . (C4) Runtime time of Alg. 1 as precision ε increases, for different decision probabilities p , with fixed $\delta = 0.1$.



(c) Energy shields compared against naive and periodic shield, for fairness target $\varphi = (\tau = 100, \mathcal{S} = [0.4, 0.6], \mathcal{L} = [0.49, 0.51])$ and decision probability $p = 0.3$. As energy function, we use $\zeta_{0.1;p,\mathcal{S},\mathcal{L}}^{\text{Mon}}$. **Columns (C):** The naive and periodic shield are tuned for the: (C1&2) running target $[0.4, 0.6]$; (C3&4) limit target $[0.49, 0.51]$. (C1&C3) depicts the fairness measure and (C2&C4) the accumulated number of interventions.

Figure 1: Experimental evaluation of energy shields: Fig. 1a shows the impact of different energy functions; Fig. 1b shows time-precision trade-off in Alg. 1; Fig. 1c shows the comparison with existing fairness shields.

540 occur, e.g., the car never crashes, thus its violation can be observed on a finite prefix. A liveness
 541 property asserts that “good” things will eventually happen, e.g., the car reaches its goal, thus its
 542 satisfaction can only be determined in the limit, and every finite trace can be extended to an infinite
 543 satisfiable trace. For stochastic processes one evaluates the satisfaction probability of the safety or
 544 liveness property w.r.t. the law of the process Baier & Katoen (2008). We deliberately connected
 545 short- and long-term fairness to safety and liveness respectively. The short-term running fairness
 546 property requires the fairness measure to remain within the running target at all times, which im-
 547 plies that once violated it remains violated along the infinite sequence, i.e., safety. The long-term
 548 limit fairness property requires the limit of the fairness measure to lie in the limit target, which im-
 549 plies it can only be determined in the limit, i.e., liveness. The convergence in the limit holds because
 550 the fairness measure is an average. This normalization allows the distance to μ^* to shrink and the
 551 process to converge almost surely. We emphasize that this implies that almost surely the fairness
 552 measure eventually remains within the limit target. It does not imply that the fairness measure even-
 553 tually remains within the limit target almost surely. Formally, $\mathbb{P}(\exists \tau \in \mathbb{N} \forall t \geq \tau : M_t \in \mathcal{L}) = 1$ is
 554 satisfied, and $\exists \tau \in \mathbb{N} : \mathbb{P}(\forall t \geq \tau : M_t \in \mathcal{L}) = 1$ is not.

555 12 CONCLUSION

556 We introduced *energy shields*, a physics-inspired framework for enforcing fairness at runtime. En-
 557 ergy shields are a lightweight, probabilistic mechanisms that provide rigorous short- and long-term
 558 fairness guarantees. Utilizing a bowl-shaped energy function, enforcing fairness is reduced to an
 559 energy minimization problem, which enables gentle, adaptive interventions. We provide a synthesis
 560 algorithm based on binary search and dynamic programming to find the least-invasive shield for
 561 a given specification. The experimental validation demonstrates the practicality of our synthesis
 562 procedure and supports, in the comparison with Cano et al. (2025b), the claim that energy shields
 563 are the first shields to satisfy both short- and long-term guarantees. While our results establish a
 564 foundational theory of energy-based fairness shielding, extending these guarantees to multi-group
 565 fairness, properties beyond fairness and dynamic environments remains an exciting direction.

566 DISCLAIMERS

567 ETHICS STATEMENT

568 This work is strongly motivated by the problem of ensuring algorithmic fairness in machine learning
 569 applications. Although this is a sensitive topic, the ethical considerations for this paper are minimal.
 570 Our contributions are largely theoretical, and we do not employ any datasets containing sensitive
 571 information.

572 REPRODUCIBILITY STATEMENT

573 The proofs for all theoretical results are stated in the paper. For most results, the main body of the
 574 text contains only a sketch of the argument and the full proof can be found in the appendix. The
 575 code required to reproduce the experiments, including the synthesis procedure, are provided in the
 576 supplementary materials. Our empirical evaluation requires only models computational resources.
 577 All computation time experiments reported were performed with a Macbook with the M2 chip.

585 USE OF LLMs

586 LLM tools were used to polish the text, to brainstorm arguments for proving the theoretical results,
 587 and to generate parts of the code used for the experimental evaluation.

590 REFERENCES

591 Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna Wallach. A reduc-
 592 tions approach to fair classification. In *International Conference on Machine Learning (ICML)*,
 593 pp. 60–69. PMLR, 2018.

594 Parand A. Alamdari, Toryn Q. Klassen, Elliot Creager, and Sheila A. Mcilraith. Remembering to be
 595 fair: Non-Markovian fairness in sequential decision making. In *Proceedings of the International*
 596 *Conference on Machine Learning (ICML)*, volume 235, pp. 906–920, 2024.

597

598 Aws Albarghouthi and Samuel Vinitsky. Fairness-aware programming. In *Proceedings of the Con-*
 599 *ference on Fairness, Accountability, and Transparency*, pp. 211–219, 2019.

600 Aws Albarghouthi, Loris D’Antoni, Samuel Drews, and Aditya V Nori. Fairsquare: probabilis-
 601 tic verification of program fairness. *Proceedings of the ACM on Programming Languages*, 1
 602 (OOPSLA):1–30, 2017.

603

604 Mohammed Alshiekh, Roderick Bloem, Rüdiger Ehlers, Bettina Könighofer, Scott Niekum, and
 605 Ufuk Topcu. Safe reinforcement learning via shielding. In *Proceedings of the AAAI conference*
 606 *on artificial intelligence*, volume 32, 2018.

607

608 C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen. Model-checking algorithms for continuous-
 609 time markov chains. *IEEE Transactions on Software Engineering*, 29(6):524–541, 2003.

610

611 Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT press, 2008.

612

613 Ezio Bartocci, Jyotirmoy Deshmukh, Alexandre Donzé, Georgios Fainekos, Oded Maler, Dejan
 614 Ničković, and Sriram Sankaranarayanan. Specification-based monitoring of cyber-physical sys-
 615 tems: a survey on theory, tools and applications. In *Lectures on Runtime Verification*, pp. 135–175.
 Springer, 2018.

616

617 Osbert Bastani, Xin Zhang, and Armando Solar-Lezama. Probabilistic verification of fairness prop-
 618 erties via concentration. *Proceedings of the ACM on Programming Languages*, 3(OOPSLA):
 1–27, 2019.

619

620 Jan Baumeister, Bernd Finkbeiner, Frederik Scheerer, Julian Siber, and Tobias Wagenpfeil. Stream-
 621 based monitoring of algorithmic fairness. In *International Conference on Tools and Algorithms*
 622 *for the Construction and Analysis of Systems*, pp. 60–81. Springer, 2025.

623

624 Vivek S. Borkar. *Stochastic approximation: a dynamical systems viewpoint*. Springer, 2008.

625

626 Filip Cano, Thomas A. Henzinger, and Konstantin Kueffner. Algorithmic fairness: A runtime per-
 627 spective. In *Proceedings of the International Conference on Runtime Verification (RV)*, pp. 1–21,
 2025a.

628

629 Filip Cano, Thomas A. Henzinger, Bettina Könighofer, Konstantin Kueffner, and Kaushik Mallik.
 630 Fairness shields: Safeguarding against biased decision makers. *Proceedings of the AAAI Confer-
 631 ence on Artificial Intelligence*, 39(15):15659–15668, 2025b.

632

633 Steven Carr, Nils Jansen, Sebastian Junges, and Ufuk Topcu. Safe reinforcement learning via shield-
 634 ing under partial observability. In *Proceedings of the AAAI conference on artificial intelligence*,
 volume 37, pp. 14748–14756, 2023.

635

636 Simon Caton and Christian Haas. Fairness in machine learning: A survey. *ACM Computing Surveys*,
 2020.

637

638 Filip Cano Córdoba, Alexander Palmisano, Martin Fränzle, Roderick Bloem, and Bettina
 639 Könighofer. Safety shielding under delayed observation. In *Proceedings of the International*
 640 *Conference on Automated Planning and Scheduling*, volume 33, pp. 80–85, 2023.

641

642 Alexander D’Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, David Sculley, and Yoni
 643 Halpern. Fairness is not static: deeper understanding of long term fairness via simulation studies.
 644 In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT)*, pp.
 525–534, 2020.

645

646 Alexandre Donzé and Oded Maler. Robust satisfaction of temporal logic over real-valued signals.
 647 In *International Conference on Formal Modeling and Analysis of Timed Systems*, pp. 92–106.
 Springer, 2010.

648 Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness
 649 through awareness. In *Proceedings of the 3rd innovations in theoretical computer science confer-*
 650 *ence*, pp. 214–226, 2012.

651

652 Peter Faymonville, Bernd Finkbeiner, Maximilian Schwenger, and Hazem Torfah. Real-time stream-
 653 based monitoring. 2017.

654

655 Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubra-
 656 manian. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD*
 657 *International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 259–268, 2015.

658

659 Bishwamitra Ghosh, Debabrota Basu, and Kuldeep S Meel. Justicia: A stochastic sat approach to
 660 formally verify fairness. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*,
 661 volume 35, pp. 7554–7563, 2021.

662

663 Paula Gordaliza, Eustasio Del Barrio, Gamboa Fabrice, and Jean-Michel Loubes. Obtaining fairness
 664 using optimal transport theory. In *International Conference on Machine Learning (ICML)*, pp.
 665 2357–2365. PMLR, 2019.

666

667 Ashutosh Gupta, Thomas A. Henzinger, Konstantin Kueffner, Kaushik Mallik, and David Pape.
 668 Monitoring robustness and individual fairness. In *KDD 2025*, 2025.

669

670 Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In
 671 *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 3315–3323, 2016.

672

673 Thomas A. Henzinger, Mahyar Karimi, Konstantin Kueffner, and Kaushik Mallik. Runtime moni-
 674 toring of dynamic fairness properties. In *Proceedings of the ACM Conference on Fairness, Ac-
 675 countability, and Transparency (FAccT)*, pp. 604–614. ACM, 2023a.

676

677 Thomas A Henzinger, Nicolas Mazzocchi, and N Ege Saraç. Quantitative safety and liveness. In
 678 *FoSSaCS*, pp. 349–370, 2023b.

679

680 Nils Jansen, Bettina Könighofer, Sebastian Junges, Alex Serban, and Roderick Bloem. Safe re-
 681 enforcement learning using probabilistic shields (invited paper). In *CONCUR*, volume 171 of
 682 *LIPics*, pp. 3:1–3:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

683

684 Rajeeva Laxman Karandikar and Mathukumalli Vidyasagar. Convergence rates for stochastic ap-
 685 proximation: Biased noise with unbounded variance, and applications. *Journal of Optimization*
 686 *Theory and Applications*, 203(3):2412–2450, 2024.

687

688 Joost-Pieter Katoen. The probabilistic model checking landscape. In *Proceedings of the 31st Annual*
 689 *ACM/IEEE Symposium on Logic in Computer Science*, pp. 31–45, 2016.

690

691 Leslie Lamport. Proving the correctness of multiprocess programs. *IEEE transactions on software*
 692 *engineering*, (2):125–143, 1977.

693

694 Yannan Li, Jingbo Wang, and Chao Wang. Certifying the fairness of knn in the presence of dataset
 695 bias. In *International Conference on Computer Aided Verification (CAV)*. Springer, 2023.

696

697 Lydia T Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt. Delayed impact of fair
 698 machine learning. In *International Conference on Machine Learning*, pp. 3150–3158. PMLR,
 699 2018.

700

701 Oded Maler and Dejan Nickovic. Monitoring temporal properties of continuous signals. In *Inter-
 702 national Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems*, pp. 152–166.
 703 Springer, 2004.

704

705 Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey
 706 on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35, 2021.

707

708 Anna Meyer, Aws Albarghouthi, and Loris D’Antoni. Certifying robustness to programmable data
 709 bias in decision trees. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:26276–
 710 26288, 2021.

702 Herbert Robbins and David Siegmund. A convergence theorem for non negative almost super-
 703 martingales and some applications. In *Optimizing methods in statistics*, pp. 233–257. Elsevier,
 704 1971.

705 Scott D Stoller, Ezio Bartocci, Justin Seyster, Radu Grosu, Klaus Havelund, Scott A Smolka, and
 706 Erez Zadok. Runtime verification with state estimation. In *International conference on runtime*
 707 *verification*, pp. 193–207. Springer, 2011.

708 Bing Sun, Jun Sun, Ting Dai, and Lijun Zhang. Probabilistic verification of neural networks against
 709 group fairness. In *International Symposium on Formal Methods*, pp. 83–102. Springer, 2021.

710 Min Wen, Osbert Bastani, and Ufuk Topcu. Algorithms for fairness in sequential decision making.
 711 In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 1144–1152.
 712 PMLR, 2021.

713 Wen-Chi Yang, Giuseppe Marra, Gavin Rens, and Luc De Raedt. Safe reinforcement learning via
 714 probabilistic logic shields. In *IJCAI*, pp. 5739–5749. ijcai.org, 2023.

715 Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P Gummadi. Fair-
 716 ness constraints: A flexible approach for fair classification. *The Journal of Machine Learning*
 717 *Research*, 20(1):2737–2778, 2019.

721 A DETAILED PROOFS

722 **Claim 1** (Shielded decision process). *The shielded decision process generated by the energy shield
 723 can be written as a sequence of Bernoulli random variables with evolving bias, $Z_t \sim \text{Bernoulli}(p_t)$.
 724 The biases are defined recursively as $p_1 = 1$ and $p_{t+1} = f(\mu_t)$ for a given history z_1, \dots, z_t , where*

$$725 \quad f(\mu) = \begin{cases} p + (1-p)\zeta(\mu) & \text{if } \mu \leq \kappa, \\ p \cdot (1 - \zeta(\mu)) & \text{if } \mu > \kappa. \end{cases} \quad (13)$$

726 Moreover, the resulting shielded fairness process can be written as

$$727 \quad M_t = M_{t-1} + \frac{1}{t}(Z_t - M_{t-1}) \quad (\text{with } M_0 = 0). \quad (14)$$

728 *Proof.* Both equations are just simple computations.

729 **Equation 13.** Suppose $\mu_t \leq \kappa$, i.e., the shield favors 1's. Then a 1 can be obtained either by
 730 $X_{t+1} = 1$ (which happens with probability p), or by flipping the decision $X_{t+1} = 0$ (which happens
 731 with probability $(1-p)\zeta(\mu_t)$). Therefore, when $\mu_t \leq \kappa$, Z_{t+1} behaves like a Bernoulli of bias
 732 $p + (1-p)\zeta(\mu_t)$. Analogously, when $\mu_t > \kappa$, the shield favors 0's, so to obtain a 1 we need to toss
 733 $X_{t+1} = 0$ and for the shield to fail to flip the decision, which happens with probability $(1 - \zeta(\mu_t))$.

734 Equation 14.

$$735 \quad M_t = \frac{\sum_{i=1}^t Z_i}{t} = \frac{Z_t}{t} + \frac{t-1}{t} \cdot \frac{\sum_{i=1}^{t-1} Z_i}{t-1} = \frac{Z_t}{t} + \frac{t-1}{t} M_{t-1} = M_{t-1} + \frac{1}{t}(Z_t - M_{t-1}).$$

736 \square

745 A.1 LONG TERM GUARANTEES

746 **Lemma 1.** *The function $f: [0, 1] \rightarrow [0, 1]$ defined as in Eq. 2 is continuously differentiable, and
 747 has a unique point $\mu^* \in [0, 1]$ such that $f(\mu^*) = \mu^*$. Furthermore, μ^* sits between p and κ .*

748 *Proof.* We need to prove smoothness, existence of the fixpoint, and that it sits between p and κ .

749 *Smoothness.* The function f inherits continuous differentiability from ζ clearly at all points except
 750 maybe $\mu = \kappa$. For $\mu = \kappa$, the assumption of the energy function being flat at the pivoting point κ
 751 guarantees continuous differentiability (Def. 1, item 3). We first write the expression for f' :

$$752 \quad f'(\mu) = \begin{cases} -p \cdot \zeta'(\mu) & \text{if } \mu > \kappa \\ (1-p)\zeta'(\mu) & \text{otherwise.} \end{cases} \quad (15)$$

756 We can check that f is continuous at $\mu = \kappa$, as
 757

$$758 \quad p \cdot (1 - \zeta(\kappa)) = p, \text{ and } p + (1 - p) \cdot \zeta(\kappa) = p.$$

759 Similarly, we can check that f' is continuous at $f = \kappa$, as
 760

$$761 \quad -p \cdot \zeta'(\kappa) = 0, \text{ and } (1 - p) \cdot \zeta'(\kappa) = 0.$$

762 Note that both are necessary conditions. For f to be continuous at $\mu = \kappa$, we need
 763

$$764 \quad p(1 - \zeta(\kappa)) = p + (1 - p)\zeta(\kappa) \iff -p\zeta(\kappa) = \zeta(\kappa) - p\zeta(\kappa) \iff \zeta(\kappa) = 0.$$

765 Similarly, for f' to be continuous as $\mu = \kappa$, we need
 766

$$767 \quad -p\zeta'(\kappa) = (1 - p)\zeta'(\kappa) \iff \zeta'(\kappa) = 0.$$

768 *Existence and uniqueness of a fixpoint.* Consider the function $g(\mu) = f(\mu) - \mu$. We have $g(0) = p(1 - \zeta(0)) - 0 > 0$ and $g(1) = p + (1 - p)\zeta(1) - 1 < 1 - 1 = 0$. Since g is continuous, there must be
 769 a point $\mu^* \in [0, 1]$ such that $g(\mu^*) = 0$. Suppose this point is not unique, i.e., there exist two point
 770 $x < \mu$ such that $g(x) = g(\mu) = 0$. As we can see in Eq. equation 15, $f'(\mu) \leq 0$ for all $\mu \in [0, 1]$,
 771 therefore g is non-increasing, so if $g(x) = g(\mu) = 0$, then $g(z) = 0$ for all $z \in [x, \mu]$. The interval
 772 $[x, \mu]$ contains at least two numbers that are either larger or lower than κ . Let $z_1, z_2 \in [x, \mu] \cap [0, \kappa)$
 773 be such numbers, with $z_1 < z_2$. Then we have
 774

$$775 \quad p + (1 - p)\zeta(z_1) - z_1 = p + (1 - p)\zeta(z_2) - z_2,$$

776 which implies that
 777

$$778 \quad z_2 - z_1 = (1 - p)(\zeta(z_2) - \zeta(z_1)).$$

779 On the left-hand side we have a positive number (because $z_1 < z_2$). On the right-hand side we have
 780 a non-positive number (because ζ is decreasing in the range $[0, \kappa]$), which is a contradiction. Since
 781 assuming $z_1, z_2 \in [x, \mu] \cap [0, \kappa)$ leads to a contradiction, it must be that $z_1, z_2 \in [x, \mu] \cap (\kappa, 1]$.
 782 However, with a similar argument this implies that
 783

$$784 \quad p(1 - \zeta(z_1)) - z_1 = p(1 - \zeta(z_2)) - z_2 \iff z_2 - z_1 = p(\zeta(z_1) - \zeta(z_2)).$$

785 Since ζ is in the decreasing range, in the right-hand side we have again a non-positive term, which
 786 is a contradiction.

787 *Placement of the fixpoint.* We write the case where $p \leq \kappa$, the case where $p > \kappa$ is analogous. If
 788 $p \leq \kappa$, we have

$$789 \quad g(p) = (1 - p)\zeta(p) \geq 0, \quad \text{and} \quad g(\kappa) = p - \kappa \leq 0.$$

790 By continuity, there exists a point $\mu^* \in [p, \kappa]$ such that $g(\mu^*) = \mu^*$. Note that this includes the
 791 special case where $p = \kappa$, in which the unique fixpoint is $\mu^* = \kappa = p$. \square
 792

793 **Lemma 2.** *The process described in Eq. 3 converges almost surely to the unique fixpoint μ^* of f ,
 794 as defined in Eq. 2. Furthermore, the convergence rate of the error satisfies $(M_t - \mu^*)^2 = o(1/t^\lambda)$
 795 for all $\lambda \in (0, 1)$ almost surely.*
 796

797 *Proof. Error construction.* For convenience, we will prove that the sequence converges to the unique
 798 root of $g(\mu) = f(\mu) - \mu$. Let $(V_t)_{t \in \mathbb{N}}$ be the sequence of squared distances to the target, i.e.,
 799 $V_t = (\mu_t - \mu^*)^2$. We will prove that $V_t \rightarrow 0$ a.s. (almost surely).

800 First we need to find a recurrence formula for V_t . Recall from Eq. 6 the recurrence for μ_t is
 801

$$802 \quad \mu_t = \mu_{t-1} + \gamma_t(g(\mu_{t-1}) + \xi_t),$$

803 where $\gamma_t = 1/t$ and ξ_t has null expectation. For V_t we have
 804

$$805 \quad V_t = (\mu_t - \mu^*)^2 = (\mu_{t-1} - \mu^* + \gamma_t(g(\mu_{t-1}) + \xi_t))^2 \quad (16)$$

$$806 \quad = (\mu_{t-1} - \mu^*)^2 + 2\gamma_t(\mu_{t-1} - \mu^*)(g(\mu_{t-1}) + \xi_t) + \gamma_t^2(g(\mu_{t-1}) + \xi_t)^2. \quad (17)$$

807 Taking conditional expectations we have
 808

$$809 \quad \mathbb{E}[V_t | \mu_{t-1}] = V_{t-1} + 2\gamma_t(\mu_{t-1} - \mu^*) + \gamma_t^2 \mathbb{E}[(g(\mu_{t-1}) + \xi_t)^2]. \quad (18)$$

810 Since $g(\mu_{t-1})$ is deterministic given μ_{t-1} and $\mathbb{E}[\xi_t | \mu_{t-1}] = 0$, we have
 811

$$812 \mathbb{E}[(g(\mu_{t-1}) + \xi_t)^2] = g(\mu_{t-1})^2 + \sigma_t^2,$$

813 where we define $\sigma_t^2 = \mathbb{E}[\xi_t^2 | \mu_{t-1}]$. Recall that $\xi_t = Z_t - f(\mu_{t-1})$, which is always in the interval
 814 $[-1, 1]$, therefore $\sigma_t^2 \leq 1$. Plugging everything into Eq. equation 17, we have
 815

$$816 \mathbb{E}[V_t | \mu_{t-1}] \leq V_{t-1} + 2\gamma_t(\mu_{t-1} - \mu^*)g(\mu_{t-1}) + \gamma_t^2(g(\mu_{t-1})^2 + \sigma_t^2). \quad (19)$$

817 Let $\alpha_t = -2\gamma_t \frac{g(\mu_{t-1})}{\mu_{t-1} - \mu^*}$ and $\beta_t = \gamma_t^2(g(\mu_{t-1})^2 + \sigma_t^2)$. Note that α_t is well defined, even when
 818 $\mu_{t-1} = \mu^*$ as, using L'Hôpital's rule:
 819

$$820 \lim_{\mu \rightarrow \mu^*} \frac{g(\mu)}{\mu - \mu^*} = \lim_{\mu \rightarrow \mu^*} \frac{f(\mu) - \mu}{\mu - \mu^*} = \lim_{\mu \rightarrow \mu^*} \frac{g'(\mu)}{1} = g'(\mu^*).$$

821 Also, g is decreasing around μ^* , so $g(\mu_{t-1})$ and $\mu_{t-1} - \mu^*$ have opposite signs. Therfore $\alpha_t \geq 0$.
 822 We can then rewrite Eq. equation 19 as
 823

$$824 \mathbb{E}[V_t | \mu_{t-1}] \leq (1 - \alpha_t)V_{t-1} + \beta_t. \quad (20)$$

825 *Proof of convergence.* This is a standard form for the condition in the classical Robbins-Siegmund
 826 theorem Robbins & Siegmund (1971), which guarantees convergence of V_t almost surely whenever
 827 the sequences (α_t) and (β_t) are non-negative and satisfy the limiting properties that $\sum_{t=1}^{\infty} \alpha_t = \infty$
 828 and $\sum_{t=1}^{\infty} \beta_t < \infty$ almost surely. We use the recent results in Karandikar & Vidyasagar (2024) to
 829 guarantee convergence. In particular, we use Thm. 5.1 to guarantee $V_t \rightarrow \infty$ almost surely, and
 830 Thm. 5.2 to guarantee rates of convergence. To apply both theorems, we need to guarantee that
 831

- 832 1. $\sum_{t=1}^{\infty} \alpha_t = \infty$ almost surely, and
- 833 2. $\sum_{t=1}^{\infty} \beta_t < \infty$ almost surely.

834 Note that both results are true in the strict sense (not only ‘‘almost surely’’). In the case of α_t , note
 835 that there exists $C > 0$ such that $g(\mu_t)/(\mu^* - \mu_t) \geq C$. Otherwise, it would imply that g has a zero
 836 other than μ^* or that $g'(\mu^*) = 0$, both which we know are not true. Therefore
 837

$$838 \sum_{t=1}^{\infty} \alpha_t \geq 2C \sum_{t=1}^{\infty} \gamma_t = \infty.$$

839 Since $\gamma_t = 1/t$, the sum $\sum_t \gamma_t$ diverges. For the case of β_t , recall by definition that $g(\mu_t)^2 \leq 1$. We
 840 have also previously established that $\sigma_t^2 \leq 1$:

$$841 \sum_{t=1}^{\infty} \beta_t \leq 2 \sum_{t=1}^{\infty} \gamma_t^2 < \infty.$$

842 Theorem 5.1 in Karandikar & Vidyasagar (2024) is stated in terms of a sequence h_t . We can use
 843 $h_t = V_t$ (since the identity is what Karandikar & Vidyasagar (2024) defines as a class \mathcal{B} function),
 844 to guarantee $V_t \rightarrow 0$ almost surely.

845 *Convergence rates.* From Theorem 5.2 in Karandikar & Vidyasagar (2024), we have that $V_t =$
 846 $o(1/t^{\lambda})$ for any λ such that $\alpha_t \geq \lambda/t$ for sufficiently large t . In our case, we can take $\lambda =$
 847 $2|g'(\mu^*)| - \epsilon$ for any $\epsilon > 0$. Therefore, $V_t = o(1/t^{\lambda})$ for any $\lambda \in (0, 1) \cap (0, 2|g'(\mu^*)|) =$
 848 $(0, \min\{1, 2|g'(\mu^*)|\})$.

849 However, we can be more fine grained and show that actually $\min\{1, 2|g(\mu^*)|\} = 1$. Recall from
 850 the definition of g that $g'(\mu) = f'(\mu) - 1$, and $f'(\mu) \leq 0$ for all $\mu \in [0, 1]$. Therefore $|g'(\mu)| =$
 851 $1 + |f'(\mu)| \geq 1$. Altogether, we obtain the expected result of $V_t = o(1/t^{\lambda})$ for all $\lambda \in (0, 1)$. \square

852 **Lemma 3.** *For the process described in Eq. 3 , the corresponding sequence of average interference
 853 $(\nu_t)_{t \in \mathbb{N}}$ converges to $h(\mu^*)$, where μ^* is the fixpoint of f (Eq. 2) and h is as defined in Eq. 8.*

864 *Proof.* By definition $N_t = \frac{1}{t} \sum_{i=1}^t Y_i$.
 865

866 Since $\mathbb{E}[Y_t | \mu_{t-1}] = h(\mu_{t-1})$, and $\mu_t \xrightarrow{t \rightarrow \infty} \mu^*$ almost surely. Since h is a continuous function,
 867

$$868 \lim_{t \rightarrow \infty} N_t = h\left(\lim_{t \rightarrow \infty} M_t\right) = h(\mu^*). \\ 869$$

□

870
 871
 872 **Theorem 3.** Let $\mu^* \in [0, 1]$. Given the shielded decision process as described in Eq. 3, with bias
 873 parameter p and energy function ζ , the shielded fairness process $(M_t)_{t \in \mathbb{N}}$ converges almost surely
 874 to μ^* if and only if

$$875 \zeta(\mu^*) = \begin{cases} |\mu^* - p|/(1-p) & \text{if } p < \mu^*, \\ 876 |\mu^* - p|/p & \text{otherwise.} \end{cases} \quad (21)$$

877 Furthermore, the expected interference cost $(\mathbb{E}[\nu_t])_{t \in \mathbb{N}}$ converges almost surely to $|\mu^* - p|$.
 878

879 *Proof.* We know from Lemma 2 that the fairness outcome converges to a fixpoint that is always
 880 between p and κ . If we want the process to converge to a desired outcome μ^* , then we need to find
 881 κ and ζ such that $f(\mu^*) = \mu^*$.
 882

883 The function $f(\mu)$ is defined by parts depending on whether μ is larger or smaller than κ . If $p \leq \mu^*$,
 884 then we need $\kappa \geq \mu^*$, so we will be using the expression $f(\mu) = p + (1-p)\zeta(\mu)$ for $\mu = \mu^*$.
 885 Imposing μ^* to be a fixpoint, we have

$$886 \mu^* = p + (1-p)\zeta(\mu^*) \iff \zeta(\mu^*) = \frac{\mu^* - p}{1-p}. \\ 887$$

888 Analogously, if $p \geq \kappa$, then we need $\kappa \leq \mu^*$, so we will use the expression $f(\mu) = p(1 - \zeta(\mu))$. If
 889 we set μ^* to be a fixpoint, we have
 890

$$891 \mu^* = p(1 - \zeta(\mu^*)) \iff \zeta(\mu^*) = 1 - \frac{\mu^*}{p}. \\ 892$$

893 To prove the expectation of cost, just recall from Lemma 3 that the expected cost converges to $h(\mu^*)$.
 894 If $p \leq \mu^*$, then $\kappa \geq \mu^*$, so we have
 895

$$896 h(\mu^*) = (1-p)\zeta(\mu^*), \quad \text{and} \quad \zeta(\mu^*) = \frac{\mu^* - p}{1-p} \implies h(\mu^*) = \mu^* - p. \\ 897$$

898 Analogously, if $p \geq \mu^*$, then $\kappa \leq \mu^*$, so we have
 899

$$900 h(\mu^*) = p\zeta(\mu^*), \quad \text{and} \quad \zeta(\mu^*) = 1 - \frac{\mu^*}{p} \implies h(\mu^*) = p - \mu^*. \\ 901$$

902 This completes the proof, since in both cases $h(\mu^*) = |\mu^* - p|$. □

904 A.2 SHORT-TERM GUARANTEES: UPPER BOUNDS

905 **Lemma 4.** Let $\Delta_t = \mu_t - \mu^*$. Recall $\xi_t = Z_t - f(M_{t-1})$ from Equation 6. The following holds:
 906

$$907 \Delta_t = A_{1,t} \Delta_1 + \sum_{i=1}^{t-1} w_{i+1,t} \xi_{i+1}, \quad \text{where} \quad (22) \\ 908$$

$$909 A_{i,t} = \prod_{j=i}^{t-1} \left(1 - \frac{\alpha_j}{j+1}\right), \quad w_{i,t} = \frac{A_{i,t}}{i}, \quad \text{and} \quad \alpha_t := 1 - \frac{f(\mu_t) - \mu^*}{\mu_t - \mu^*}. \quad (23) \\ 910$$

911 Furthermore, the following inequality holds:
 912

$$913 \sum_{i=1}^{t-1} w_{i+1,t}^2 \leq \frac{4^{1-\beta}}{t+1} \quad \text{with} \quad \beta = \sup_{r \in [0,1]} f'(r). \quad (24) \\ 914$$

918 *Proof.* From the definition of Δ_t and the recursion for μ_t expressed in Equation 3, we have
919

$$920 \quad 921 \quad \Delta_{t+1} = \left(1 - \frac{\alpha_t}{t+1}\right) \Delta_t + \frac{\xi_{t+1}}{t+1}. \quad (25)$$

922 Equation 22 follows by induction on t from Eq. 25. The base case $t = 1$ is trivial, since product
923 defining $A_{i,t}$ is empty, and so is the sum of $w_{i+1,t} \xi_{i+1}$. For the induction step, can apply the
924 induction hypothesis to Equation 25 to obtain
925

$$926 \quad \Delta_{t+1} = \left(1 - \frac{\alpha_t}{t+1}\right) \Delta_t + \frac{\xi_{t+1}}{t+1} \\ 927 \\ 928 \quad = \left(1 - \frac{\alpha_t}{t+1}\right) \left(A_{1,t} \Delta_1 + \sum_{i=1}^{t-1} w_{i+1,t} \xi_{i+1} \right) + \frac{\xi_{t+1}}{t+1} \cdot A_{1,t} \left(1 - \frac{\alpha_t}{t+1}\right) \Delta_1 +$$

930 Note that, from the definition of $A_{i,t}$, we have
931

$$932 \quad 933 \quad \left(1 - \frac{\alpha_t}{t+1}\right) A_{i,t} = A_{i,t+1}, \quad \text{and} \quad \left(1 - \frac{\alpha_t}{t+1}\right) w_{i,t} = w_{i,t+1}.$$

934 Using the previous identities, plus the case that $w_{t+1,t+1} = \frac{1}{t+1}$ we get
935

$$936 \quad \Delta_{t+1} = A_{1,t+1} \Delta_1 + \sum_{i=1}^{t-1} w_{i+1,t+1} \xi_{i+1} + \frac{\xi_{t+1}}{t+1} \\ 937 \\ 938 \quad = A_{1,t+1} \Delta_1 + \sum_{i=1}^t w_{i+1,t+1} \xi_{i+1},$$

939 which finishes the induction step.
940

941 To prove the lower bound expressed in Equation 24, consider $\beta = \sup_{r \in [0,1]} f'(r)$. Since $f(\mu^*) = \mu^*$, and f is differentiable on the compact interval $[0, 1]$, by the mean-value theorem
942

$$943 \quad \frac{f(\mu) - \mu^*}{\mu - \mu^*} = \frac{f(\mu) - f(\mu^*)}{\mu - \mu^*} \in \{f'(z) : z \text{ between } \mu \text{ and } \mu^*\} \subseteq (-\infty, \beta].$$

944 Therefore, we have the following lower bound for α_t that holds for every $t \in \mathbb{N}$:
945

$$946 \quad \alpha_t = 1 - \frac{f(\mu) - \mu^*}{\mu - \mu^*} \geq 1 - \beta. \quad (26)$$

947 For $u \in (0, 1)$, we can use the bound $\log(1 - u) \leq -u$. We can apply this to the definition of $A_{i,t}$
948

$$949 \quad \log A_{i,t} = \sum_{j=i}^{t-1} \log \left(1 - \frac{\alpha_j}{j+1}\right) \leq - \sum_{j=i}^{t-1} \frac{\alpha_j}{j+1} \leq (\beta - 1) \log \frac{t+1}{i+1} = \log \left(\frac{i+1}{t+1}\right)^{1-\beta}. \quad (27)$$

950 The last inequality stems from using twice that $\sum_{i=1}^t 1/i \leq \log(t+1)$. Using $(i+1)/i \leq 2$, we
951 have
952

$$953 \quad w_{i,t}^2 = \left(\frac{A_{i,t}}{i}\right)^2 \leq \frac{1}{i^2} \left(\frac{i+1}{t+1}\right)^{2-2\beta} \leq \left(\frac{2}{t+1}\right)^{2-2\beta} i^{-2\beta}.$$

954 Using this bound, we have
955

$$956 \quad \sum_{i=1}^{t-1} w_{i+1,t}^2 \leq \frac{4^{1-\beta}}{(t+1)^{2-2\beta}} \sum_{i=1}^{t-1} (i+1)^{-2\beta} \stackrel{(*)}{\leq} \frac{4^{1-\beta}}{(t+1)^{2-2\beta}} (t-1)t^{-2\beta} \leq \frac{4^{1-\beta}}{t+1},$$

957 where in the $(*)$ inequality we are bounding each element of the sum by the largest one, which is
958 $t^{-2\beta}$, and that there are $t-1$ summands. \square
959

960 **Lemma 5.** Let $\delta > 0$. Let $\tau = \frac{2^{1+\frac{1}{1-\beta}}}{\delta^{1-\beta}} - 1$, and $K = (1/32) \cdot 4^\beta$. For all $t \geq \tau$, we have
961

$$962 \quad \mathbb{P}[\Delta_t > \delta] \leq \exp(-K\delta^2 t), \quad \mathbb{P}[\Delta_t < -\delta] \leq \exp(-K\delta^2 t). \quad (28)$$

972 *Proof.* The idea is to use the expression in Equation 22 to bound Δ_t . Of the two summands, $A_{1,t}\Delta_1$
 973 will be bounded by bounding $A_{1,t}$ for large enough t . For the second summand, we will show it
 974 is a martingale and use a concentration inequality to bound its value. We show the case for the
 975 first inequality. The second inequality follows the same argument, with a symmetric use of the
 976 Azuma-Hoeffding's inequality.

977 Fix $t \in \mathbb{N}$ and consider $E_{i,t} = \sum_{j=1}^i w_{i+1,t}\xi_{i+1}$, for $i < t$. When conditioning over the decisions
 978 up to time i , we have $\mathbb{E}[\xi_{i+1} \mid \mu_i] = 0$, and therefore
 979

$$980 \mathbb{E}[E_{i,t} \mid \mu_i] = E_{i-1,t} + w_{i+1,t}\mathbb{E}[\xi_{i+1} \mid \mu_i] = E_{i-1,t}. \\ 981$$

982 So $(E_{i,t})_{i=1}^{t-1}$ is a martingale. Its increments are bounded by $|E_{i,t} - E_{i-1,t}| = |w_{i+1,t}\xi_{i+1}| \leq$
 983 $|w_{i+1,t}|$. Applying Azuma-Hoeffding's inequality to $E_{t-1,t}$ we have, for any $\delta > 0$

$$984 \mathbb{P}[E_{t-1,t} > \delta] \leq \exp\left(-\frac{\delta^2}{2\sum_{i=1}^{t-1} w_{i+1,t}^2}\right) \leq \exp\left(-\frac{(t+1)\delta^2}{2 \cdot 4^{1-\beta}}\right) \quad (29)$$

985 From Equation 27, applied to $i = 1$, we have
 986

$$987 A_{1,t} \leq \left(\frac{2}{t+1}\right)^{1-\beta}. \\ 988$$

989 We can apply the triangle inequality to Equation 22 to get
 990

$$991 |\Delta_t| \leq |A_{1,t}\Delta_1| + |M_{t-1,t}|. \quad (30)$$

992 So we can apply the bound in Equation 29 with $\delta/2$ to bound $\mathbb{P}[\Delta_t > \delta]$ as long as t is large enough
 993 to guarantee $|A_{1,t}\Delta_1| \leq \delta/2$. In such case

$$994 \mathbb{P}[\Delta_t > \delta] \leq \mathbb{P}[E_{t-1,t} > \delta/2] \leq \exp\left(-\frac{(t+1)(\delta/2)^2}{2 \cdot 4^{1-\beta}}\right) = \exp(-K(t+1)\delta^2), \quad (31)$$

995 with $K = (8 \cdot 4^{1-\beta})^{-1} = (1/32) \cdot 4^\beta$. For the previous bound to hold, we need t large enough so
 996 that $|A_{1,t}\Delta_1| \leq \delta/2$. Using $|\Delta_1| \leq 1$ and Equation 30, we have
 997

$$1000 |A_{1,t}\Delta_1| \leq \left(\frac{2}{t+1}\right)^{1-\beta}, \quad (32)$$

1001 which is smaller than $\delta/2$ for $t \geq \frac{2^{1+\frac{1}{1-\beta}}}{\delta^{1-\beta}} - 1$. □
 1002

1003 **Theorem 1.** Let $I = [L, U]$ such that $\kappa, p \in I$. Let $\tau = 4/\min(|L - \mu^*|, |U - \mu^*|)$. Then for every
 1004 $t \geq \tau$ we have

$$1005 \mathbb{P}(M_t \notin I) \leq \exp(-Kt|L - \mu^*|^2) + \exp(-Kt|U - \mu^*|^2), \quad (33)$$

1006 where K is a positive constant defined as $K = (1/32) \cdot 4^\beta$ and $\beta = \sup_{r \in [0,1]} f'(r)$.
 1007

1008 *Proof.* This is just a matter of unpacking the results from Lemma 5 into the case of having a concrete
 1009 interval. First note that
 1010

$$1011 \mathbb{P}[M_t \notin I] = \mathbb{P}[M_t < L] + \mathbb{P}[M_t > U] = \mathbb{P}[\Delta_t < -|L - \mu^*|] + \mathbb{P}[\Delta_t > |U - \mu^*|].$$

1012 We bound each of the summands using Lemma 5, which is guaranteed to hold as long as $t \geq$
 1013 $\frac{2^{1+\frac{1}{1-\beta}}}{\delta^{1-\beta}} - 1$. Since $\beta \leq 0$, we have that
 1014

$$1015 \frac{2^{1+\frac{1}{1-\beta}}}{\delta^{1-\beta}} - 1 \leq \frac{4}{\delta} - 1 \leq 4/\delta,$$

1016 so if $t \geq 4/\delta$, Lemma 5 holds. This is exactly the condition that $t \geq \tau$, for $\tau = 4/\min(|L -$
 $\mu^*|, |U - \mu^*|)$. □
 1017

1026 **Corollary 1.** Let $r_- = \exp(-K|L - \mu^*|^2)$, $r_+ = \exp(-K|U - \mu^*|^2)$. For every $T, T' \in \mathbb{N}$ such
 1027 that $\tau \leq T < T'$ we have:

$$1029 \quad \mathcal{E}_{\mathcal{S}}(M_{T:T'}) \leq \sum_{t=T}^{T'} (r_-^t + r_+^t) \quad \text{and} \quad \mathcal{E}_{\mathcal{S}}(M_{T:\infty}) \leq \frac{r_-^T}{1-r_-} + \frac{r_+^T}{1-r_+}.$$

1032 Moreover, this provides the upper bound $\mathcal{P}_{\mathcal{S}}(M_{T:T'}) \leq \mathcal{E}_{\mathcal{S}}(M_{T:T'})$ for $T' \in \mathbb{N} \cup \{\infty\}$ s.t. $T' \geq T$.

1034 *Proof.* This corollary is just an application of Boole's inequality (also known as union bound) on
 1035 the results from Theorem 1. \square

1037 A.3 SHORT-TERM GUARANTEES: MONOTONICITY

1039 To prove Theorem 2, we will use an inductive argument and the tower property of the expectation.

1040 **Definition 4** (Tower operator). Let $\mathcal{U} = \{u: [0, 1] \rightarrow \mathbb{R}\}$ be the space of all measurable functions
 1041 from $[0, 1]$ to \mathbb{R} . Let $t \in \mathbb{N}$, and ζ be an energy function. The *towering operator* is the function
 1042 $T_{\zeta,t}: \mathcal{U} \rightarrow \mathcal{U}$, that takes a function u and returns $T_{\zeta,t}u$ defined as

$$1043 \quad (T_{\zeta,t}u)(y) := f_{\zeta}(y) u(y_t^+(y)) + (1 - f_{\zeta}(y)) u(y_t^-(y)),$$

1045 where $y_t^+(y) := y + \frac{1-y}{t+1}$ and $y_t^-(y) := y - \frac{y}{t+1}$.

1047 The *iterated tower operator* is $T_{\zeta}^{(t)} = T_{\zeta,1} \circ T_{\zeta,2} \circ \dots \circ T_{\zeta,t}$

1048 **Lemma 6.** Let $u, v: [0, 1] \rightarrow \mathbb{R}$, $y \in [0, 1]$, $t \in \mathbb{N}$, ζ an energy function and (y_k) be a stochastic
 1049 process generated following Eq. equation 3. The following properties hold:

- 1051 • Expectation: $\mathbb{E}_{\zeta}[u(y_t)] = \mathbb{E}_{\zeta}[T_{\zeta,t-1}u(y_{t-1})]$
- 1053 • Iterated expectation: $\mathbb{E}_{\zeta}[u(y_t)] = \mathbb{E}[T_{\zeta}^{(t-1)}u(y_1)]$.
- 1055 • Monotonicity: If $u(y) \leq v(y)$, then $(T_{\zeta,t}u)(y) \leq (T_{\zeta,t}v)(y)$.

1057 *Proof.* The first property follows from a simple computation of the expectation, using that $\mathbb{P}[y_{t+1} =$
 1058 $y_t^+ \mid y_t] = f_{\zeta}(y_t)$. In fact, the tower operator is defined with the expectation property in mind.
 1059 Applying the expectation property consecutively leads to the next property

$$1060 \quad \mathbb{E}_{\zeta}[u(y_t)] = \mathbb{E}_{\zeta}[T_{\zeta}^{(t-1)}u(y_1)].$$

1062 Monotonicity follows directly from the definition, noting that both $f_{\zeta}(y^+)$ and $1 - f_{\zeta}(y^-)$ are always
 1063 non-negative. \square

1065 The iterated expectation property is useful because it lets us write the expectation of $u(y_m)$, which
 1066 depends on the distribution of y_m after m steps of the process, in terms of the expectation of a
 1067 function depending only on the much simpler distribution of y_1 . In particular, the distribution over
 1068 y_1 does not depend on ζ . If we want to prove a result of the form $\mathbb{E}_{\zeta_1}[u(y_t)] \leq \mathbb{E}_{\zeta_2}[u(y_t)]$ for some
 1069 u, ζ_1 , and ζ_2 , it is equivalent to prove that $\mathbb{E}[T_{\zeta_1}^{(t-1)}u(y_1)] \leq \mathbb{E}[T_{\zeta_2}^{(t-1)}u(y_1)]$, so it suffices to prove
 1070 that

$$1071 \quad T_{\zeta_1}^{(t-1)}u(y) \leq T_{\zeta_2}^{(t-1)}u(y), \quad \text{for } y \in \{0, 1\}. \quad (34)$$

1072 This observation lets us go from a local comparison to a global result. We study two properties on
 1073 the short term: the probability of violating point fairness, and the expected number of point fairness
 1074 violations. For the first one, we can take $u(y) = \mathbf{1}\{y \notin (L, 1 - U)\}$. For the second one, we can
 1075 take $u(y) = (L - y_k)_+ + (y_k - (1 - U))_+$.

1076 **Lemma 7.** Consider the functions y_t^+, y_t^- as defined in Def. 4. Then

- 1078 1. If $y \leq \kappa$ (start below):

1079 (a) If $y_t^+(y) \leq \kappa$ (no overshoot), then $|y_t^+(y) - \kappa| < |y_t^-(y) - \kappa|$.

1080 (b) If $y_t^+(y) > \kappa$ (overshoot), then
 1081

$$1082 |y_t^+(y) - \kappa| \leq \frac{1 - \kappa}{t + 1}, \quad |y_t^-(y) - \kappa| \leq \frac{1}{t + 1}. \quad (35)$$

1084 2. If $y > \kappa$ (start above):
 1085

1086 (a) If $y_t^-(y) \geq \kappa$ (no overshoot), then $|y_t^+(y) - \kappa| > |y_t^-(y) - \kappa|$,
 1087 (b) If $y_t^-(y) < \kappa$ (overshoot), then

$$1089 |y_t^-(y) - \kappa| \leq \frac{\kappa}{t + 1}, \quad |y_t^+(y) - \kappa| \leq \frac{1}{t + 1}. \quad (36)$$

1091 *Proof.* We follow the proof case by case.
 1092

- 1093 • *Case 1(a).* $|y_t^-(y) - \kappa| - |y_t^+(y) - \kappa| = y_t^+(y) - y_t^-(y) = 1/(t + 1) > 0$.
- 1094 • *Case 1(b)(left).*

$$1096 |y_t^+(y) - \kappa| = \frac{1 - y}{t + 1} - (\kappa - y) = \frac{1 - \kappa t - \kappa + y t}{t + 1}.$$

1099 Using $y \leq \kappa$ in the previous equation, we can cancel $-\kappa t$ with κt , and get the result.

- 1100 • *Case 1(b)(right).* Using $y_t^-(y) \geq 0$, we have $|y_t^-(y) - \kappa| = \kappa - y_t^-(y) < \kappa$.
- 1102 • *Case 2(a).* $|y_t^+(y) - \kappa| - |y_t^-(y) - \kappa| = y_t^+(y) - y_t^-(y) = 1/(t + 1) > 0$.
- 1104 • *Case 2(b)(left).*

$$1105 |y_t^-(y) - \kappa| = \frac{y}{t + 1} - (y - \kappa) = \frac{y - y t - y + \kappa t + \kappa}{t + 1}.$$

1107 We can cancel y and $-y$, and using $\kappa \leq y$, we can cancel $y t$ with $-y t$, and get the result.

- 1109 • *Case 2(b)(right).* Simply note that $|y_t^+(y) - \kappa| \leq |y_t^+(y) - y_t^-(y)| = 1/(t + 1)$.

1110 \square

1112 **Lemma 8.** Let $\zeta_1 \succeq \zeta_2$. Let $T \in \mathbb{N}$. Let u be a unimodal function, minimized at κ , and with
 1113 $u(y) = 0$ for $y \in [a, b]$, where

$$1115 a \leq \kappa - \frac{1}{T + 1}, \quad b \geq \kappa + \frac{1}{T + 1}.$$

1117 Then for all $t \geq T$, we have for $y \in \{0, 1\}$ that

$$1119 T_{\zeta_1}^{(t)} u(y) \leq T_{\zeta_2}^{(t)} u(y).$$

1121 *Proof. Step 1: a pointwise one-step inequality.* We first prove that for any $t \in \mathbb{N}$, $y \in [0, 1]$, and any
 1122 measurable function v , symmetric around κ , minimized at κ and null on $[a, b]$, we have

$$1123 (T_{\zeta_1, t} v)(y) \leq (T_{\zeta_2, t} v)(y). \quad (37)$$

1125 Using the definition of the tower operator, we have

$$\begin{aligned} 1127 (T_{\zeta_1, t} v)(y) - (T_{\zeta_2, t} v)(y) &= \left(f_{\zeta_1}(y) v(y_t^+) + (1 - f_{\zeta_1}(y)) v(y_t^-) \right) - \left(f_{\zeta_2}(y) v(y_t^+) + (1 - f_{\zeta_2}(y)) v(y_t^-) \right) \\ 1128 &= f_{\zeta_1}(y) \left(v(y_t^+) - v(y_t^-) \right) - f_{\zeta_2}(y) \left(v(y_t^+) - v(y_t^-) \right) \\ 1129 &= \left(f_{\zeta_1}(y) - f_{\zeta_2}(y) \right) \cdot \left(v(y_t^+) - v(y_t^-) \right). \end{aligned}$$

1133 Therefore, to check $(T_{\zeta_1, t} v)(y) - (T_{\zeta_2, t} v)(y) \leq 0$ we just need to prove that both factors of the
 1134 previous equation are of different sign, or one of them is zero.

1134 We check both sides of κ separately.
 1135
 1136 Case $y \leq \kappa$. Since $\zeta_1 \succeq \zeta_2$, we have $f_{\zeta_1}(y) \geq f_{\zeta_2}(y)$. If we are in the no-overshoot regime
 1137 ($y_t^+(y) \leq \kappa$), we have v in its descending mode in the whole interval $[y_t^-(y), y_t^+(y)]$, so $v(y_t^+(y)) \leq$
 1138 $v(y_t^-(y))$. Therefore $f_{\zeta_1}(y) - f_{\zeta_2}(y)$ and $v(y_t^+) - v(y_t^-)$ have opposed signs.

1139 If we are in the overshoot regime ($y_t^+(y) \leq \kappa$), we want to make sure that $v(y_t^+) = v(y_t^-) = 0$.
 1140 Using Eq. equation 35, we can guarantee it with
 1141

$$1142 \quad a \leq \kappa - \frac{1}{t+1}, \quad \text{and} \quad b \geq \kappa + \frac{1-\kappa}{t+1},$$

1144 which is guaranteed for $t \geq T$ in the hypothesis of the lemma.
 1145

1146 Case $y > \kappa$. This case is analogous, following the same argument as in $y \leq \kappa$, taking into account
 1147 the switch in signs.

1148 *Step 2: iterate the one-step inequality.* For $r = 0, 1, \dots, t$ define

$$1149 \quad \Phi_r := B_1 B_2 \cdots B_r A_{r+1} A_{r+2} \cdots A_t u,$$

1150 with the conventions

$$1152 \quad A_s := T_{\zeta_1, s}, \quad B_s := T_{\zeta_2, s}, \quad \Phi_0 = A_1 A_2 \cdots A_t u = T_{\zeta_1}^{(t)} u, \quad \Phi_t = B_1 B_2 \cdots B_t u = T_{\zeta_2}^{(t)} u.$$

1154

1155 Our goal is to show the chain of inequalities

$$1156 \quad \Phi_0(y) \leq \Phi_1(y) \leq \cdots \leq \Phi_t(y), \quad \text{for } y \in \{0, 1\}, \quad (38)$$

1158 which implies $T_{\zeta_1}^{(t)} u(y) \leq T_{\zeta_2}^{(t)} u(y)$ at the endpoints.

1159 *Telescoping principle.* For each $r = 1, \dots, t$, set

$$1161 \quad v^{(r)} := A_{r+1} A_{r+2} \cdots A_t u.$$

1162 Then

$$1163 \quad \Phi_{r-1} = B_1 \cdots B_{r-1} (A_r v^{(r)}), \quad \Phi_r = B_1 \cdots B_{r-1} (B_r v^{(r)}).$$

1165 Thus, if we can show

$$1166 \quad A_r v^{(r)}(y) \leq B_r v^{(r)}(y), \quad y \in \{0, 1\}, \quad (39)$$

1167 then applying the prefix operator $B_1 \cdots B_{r-1}$ (which is order-preserving because of the monotonicity property in Lemmas 6) yields

$$1169 \quad \Phi_{r-1}(y) \leq \Phi_r(y) \quad \text{for } y \in \{0, 1\}.$$

1171 Chaining over $r = 1, \dots, t$ establishes the desired inequality.

1172 *Verification of equation 39.* Fix r and write $y^+ := y_r^+(y)$, $y^- := y_r^-(y)$ for brevity. By the algebraic
 1173 identity from Step 1,

$$1175 \quad (A_r v^{(r)} - B_r v^{(r)})(y) = (f_{\zeta_1}(y) - f_{\zeta_2}(y)) (v^{(r)}(y^+) - v^{(r)}(y^-)).$$

1177 At the endpoints $y \in \{0, 1\}$, the sign of $f_{\zeta_1} - f_{\zeta_2}$ is fixed by the assumption $\zeta_1 \succeq \zeta_2$, while the sign
 1178 of $v^{(r)}(y^+) - v^{(r)}(y^-)$ is determined by the geometry of the successors and the plateau assumption:

1179 - For $y = 0$: we have $f_{\zeta_1}(0) \geq f_{\zeta_2}(0)$. If $y^+ \leq \kappa$ (no overshoot), then $|y^+ - \kappa| < |y^- - \kappa|$, hence
 1180 $v^{(r)}(y^+) \leq v^{(r)}(y^-)$. If $y^+ > \kappa$ (overshoot), then by the plateau condition both $y^\pm \in [a, b]$, so
 1181 $v^{(r)}(y^\pm) = 0$. In both cases, $v^{(r)}(y^+) - v^{(r)}(y^-) \leq 0$, so the product is ≤ 0 .

1183 - For $y = 1$: we have $f_{\zeta_1}(1) \leq f_{\zeta_2}(1)$. If $y^- \geq \kappa$ (no overshoot), then $|y^- - \kappa| < |y^+ - \kappa|$, hence
 1184 $v^{(r)}(y^-) \leq v^{(r)}(y^+)$, i.e. $v^{(r)}(y^+) - v^{(r)}(y^-) \geq 0$. If $y^- < \kappa$ (overshoot), then $y^\pm \in [a, b]$, so
 1185 $v^{(r)}(y^\pm) = 0$. In both cases, $v^{(r)}(y^+) - v^{(r)}(y^-) \geq 0$, so the product is ≤ 0 .

1187 Thus $(A_r v^{(r)} - B_r v^{(r)})(y) \leq 0$ for $y \in \{0, 1\}$, proving equation 39, which in turn proves Eq. equation 38, as we wanted. \square

1188 **Theorem 2.** Let $\zeta_1 \succeq \zeta_2$ be two energy functions, with a common minimum at κ . Let $\mathcal{S} = [L, U]$.
 1189 Let M^{ζ_1} and M^{ζ_2} be the shielded fairness process generated by enforcing the decision process of
 1190 $p \in [0, 1]$ with ζ_1 and ζ_2 , respectively. Let $\tau = \lceil \max\{1/|\kappa - L|, 1/|\kappa - U|\} \rceil$, for all $T \in \mathbb{N}$ and
 1191 $T' \in \mathbb{N} \cup \{\infty\}$ such that $T < T'$ we have

1192
$$\mathcal{E}_{\mathcal{S}}(M_{T:T'}^{\zeta_1}) \leq \mathcal{E}_{\mathcal{S}}(M_{T:T'}^{\zeta_2}), \quad \text{and} \quad \mathcal{P}_{\mathcal{S}}(M_{T:T'}^{\zeta_1}) \leq \mathcal{P}_{\mathcal{S}}(M_{T:T'}^{\zeta_2}).$$

1193
 1194 *Proof.* This is a direct consequence of Lemma 8. The condition on τ follows from it, and the
 1195 condition on the values of either probability or expectation of point fairness violation come from
 1196 the characterization of both of the safety measures as the expectation with certain functions u . In
 1197 particular, for the first one, we can take $u(y) = \mathbf{1}\{y \notin (L, U)\}$, and for the second one, we can take
 1198 $u(y) = (L - y_k)_+ + (y_k - U)_+$. The expectation of the outcome at a certain timestep corresponds
 1199 to the expectation of the iterated tower operator at the first timestep because of Lemma 6. \square
 1200

1201 A.4 SHIELD SYNTHESIS
 1202

1203 **Markov chain construction.** We construct the acyclic Markov chain $\mathcal{M} = (S, P, s_0)$ consisting
 1204 of a set of states S , a Markov transition kernel $P: S \times S \rightarrow [0, 1]$, and an initial state $s_0 = (0, 0)$.
 1205 Intuitively, the set of states is stratified into time steps $S = \bigcup_{t \in [0; T]} S_t$ where $S_t = \{(t, m) \mid m \in [t]\}$ for all $t \in [0; T]$. Intuitively, a state $s = (t, c) \in S$ is labeled with a point in time t , the positive
 1206 decision counter c clearly upper bound by the point in time t . The Markov transition kernel P is
 1207 defined as follows: for every state $s = (t, c)$ with $t \in [T - 1]$ we transition to state $s = (t + 1, c + 1)$ with probability $f(c/t)$, and to state $s' = (t + 1, c)$ with probability $1 - f(c/t)$.
 1208

1209 **Dynamic programming.** We compute the value of both the probabilistic and expected violation
 1210 measure for the interval $[1; T]$ using dynamic programming on the Markov chain (S, P, s_0) . Let
 1211 $\gamma(s) = \mathbf{1}[c/t \notin \mathcal{S}]$, then for every $s = (t, c) \in S$ we define its value $V_{\mathcal{E}}$ for the expected violation
 1212 measure as

1213
$$\begin{aligned} V_{\mathcal{E}}(s) &= \gamma(s) + (f(c/t)V_{\mathcal{E}}(t + 1, c + 1) + (1 - f(c/t))V_{\mathcal{E}}(t + 1, c)) & \text{if } t \in [T - 1] \\ V_{\mathcal{E}}(s) &= \gamma(s) & \text{if } t = T \end{aligned}$$

1214 Moreover, we define its value $V_{\mathcal{P}}$ for the expected violation measure as

1215
$$\begin{aligned} V_{\mathcal{P}}(s) &= \max \left(\gamma(s), f(c/t)V_{\mathcal{P}}(t + 1, c + 1) + (1 - f(c/t))V_{\mathcal{P}}(t + 1, c) \right) & \text{if } t \in [T - 1] \\ V_{\mathcal{P}}(s) &= \gamma(s) & \text{if } t = T \end{aligned}$$

1216 Notice that if $\gamma(c, t) = 1$ no further computation recursion is required.

1217 **Intuition.** The value function $V_{\mathcal{E}}(s)$ represents the expected number of violations incurred from
 1218 state $s = (t, c)$ onward, including a possible violation at time t . Formally, our process M takes on
 1219 the value $M_t = c/t$ at time t , then

1220
$$V_{\mathcal{E}}(t, c) = \mathbf{1}[M_t \notin \mathcal{S}] + \mathbb{E} \left[\sum_{i=t+1}^T \mathbf{1}[M_i \notin \mathcal{S}] \mid M_t \right].$$

1221 Hence, the value at the initial state $s_0 = (0, 0)$ equals the total expected number of violations up to
 1222 time T :

1223
$$V_{\mathcal{E}}(s_0) = \mathbb{E} \left[\sum_{i=1}^T \mathbf{1}[M_i \notin \mathcal{S}] \right].$$

1224 The value function $V_{\mathcal{P}}(s)$ instead captures the probability of encountering at least one violation from
 1225 state s onward. This probability equals 1 if $M_t \notin \mathcal{S}$, and otherwise it coincides with the probability
 1226 of observing a violation at some later time:

1227
$$V_{\mathcal{P}}(t, c) = \mathbb{P} \left[\max_{i \in [t; T]} \mathbf{1}[M_i \notin \mathcal{S}] = 1 \mid M_t \right].$$

1228 In particular, $V_{\mathcal{P}}(s_0)$ gives the probability of observing any violation within the time horizon $[1, T]$.

1242 **B FAMILIES OF ENERGY FUNCTIONS OF INTEREST**
 1243

1244 In this section we propose two families of general-purpose energy shields, and one family of energy
 1245 shields specifically designed to fit a given specification. Note that some of these functions have
 1246 isolated non-smooth points, where the function is continuous but not differentiable. These do not
 1247 pose a theoretical issue, since we can always “glue” the non-smooth endpoints of the two smooth
 1248 pieces using infinitely differentiable bump functions¹.
 1249

1250 **B.1 POLYNOMIAL**
 1251

1252
$$\zeta_{\kappa, \alpha, \beta}^{Pol}(x) = \alpha|x - \kappa|^\beta,$$

 1253

1254 where $\beta \in (1, \infty)$, $\kappa \in (0, 1)$ and $\alpha \in \left(0, \frac{1}{\max\{\kappa, 1-\kappa\}^\beta}\right)$. In this family, κ marks the pivoting
 1255 point, and α and β control the shape, with larger α and β producing steeper functions.
 1256

1257 **B.2 EXPONENTIAL**
 1258

1259
$$\zeta_{\kappa, \alpha, \beta}^{Exp}(x) = \alpha \left(1 - e^{-\beta(x-\kappa)^2}\right),$$

 1260

1262 where $\beta \in (0, \infty)$, $\kappa \in (0, 1)$, and $\alpha \in \left(0, \frac{1}{1-e^{-\beta(\min\{\kappa, 1-\kappa\})^2}}\right)$.
 1263

1264 **B.3 CONSTRUCTION OF A MONOTONIC FAMILY OF ENERGY FUNCTIONS**
 1265

1266 Let $p \in (0, 1)$, $\varphi = (\tau, \mathcal{S}, \mathcal{L})$ be a specification, with $\mathcal{S} = [L_{\mathcal{S}}, U_{\mathcal{S}}]$ and $\mathcal{L} = [L_{\mathcal{L}}, U_{\mathcal{L}}]$, $\mathcal{L} \subset \mathcal{S}$. We
 1267 give the construction for $\zeta_{r; p, \mathcal{S}, \mathcal{L}}^{\text{Mon}}$. For ease of notation, within this section we will call it just ζ_r . The
 1268 family of monotonic function is built differently depending on the relative position of p and \mathcal{L} .
 1269

1270 • **If $p < L$.** The family of energy functions is $(\zeta_r)_{r \in R}$, with $R = (0, 1)$, and is built as
 1271 follows. Let $\kappa = (U_{\mathcal{L}} + U_{\mathcal{S}})/2$, $a_r = (1 - r)L_{\mathcal{L}} + rU_{\mathcal{L}}$, $C_r = (a_r - p)/(1 - p)$,
 1272 $\alpha_r = (1 - r)/r$.

1273
$$\zeta_r(x) = \begin{cases} \zeta_r^1(x) & \text{if } x < a_r, \\ \zeta_r^2(x) & \text{if } a_r \leq x \leq \kappa, \\ \zeta_r^3(x) & \text{otherwise.} \end{cases} \quad (40)$$

 1274
 1275
 1276

1277 With the following definitions:
 1278

1279
$$\zeta_r^1(x) = C_r + (1 - C_r) \left(1 - e^{\frac{x-a_r}{\alpha_r}}\right),$$

 1280
 1281
$$\zeta_r^2(x) = C_r \left(1 - \frac{x-a_r}{\kappa-a_r}\right)^{\alpha_r},$$

 1282
 1283
$$\zeta_r^3(x) = 1 - \exp\left(-\left(\frac{x-\kappa}{\alpha_r}\right)^m\right).$$

 1284

1285 where m is a fixed integer $m \geq 2$, we choose $m = 2$.
 1286

1287 • **If $p > U$.** We follow a symmetric construction as the previous case. The family of energy
 1288 functions is $(\zeta_r)_{r \in R}$, with $R = (0, 1)$, and is built as follows. Let $\kappa = (L_{\mathcal{S}} + L_{\mathcal{L}})/2$,
 1289 $a_r = rL_{\mathcal{L}} + (1 - r)U_{\mathcal{L}}$, $C_r = (p - a_r)/p$, $\alpha_r = (1 - r)/r$.

1290
$$\zeta_r(x) = \begin{cases} \zeta_r^1(x) & \text{if } x < \kappa, \\ \zeta_r^2(x) & \text{if } \kappa \leq x \leq a_r, \\ \zeta_r^3(x) & \text{otherwise.} \end{cases} \quad (41)$$

 1291
 1292
 1293

1294 ¹This follows the same rationale as why most theoretical results in convergence of machine learning algo-
 1295 rithms work with the non-smooth ReLU activation function.

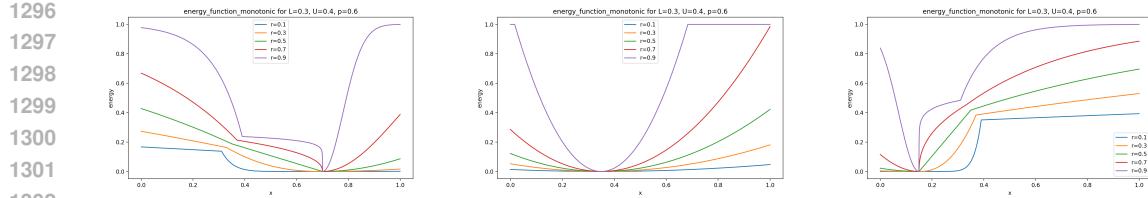


Figure 2: Monotonic families of functions

With the following definitions:

$$\begin{aligned}\zeta_r^1(x) &= 1 - e^{-\left(\frac{x-\kappa}{\alpha_r}\right)^m} \\ \zeta_r^2(x) &= C_r \left(1 - \frac{a_r - x}{a_r - \kappa}\right)^{\alpha_r}, \\ \zeta_r^3(x) &= C_r + (1 - C_r) \left(1 - e^{\frac{a_r - x}{\alpha_r}}\right).\end{aligned}$$

- **If $p \in [L_{\mathcal{L}}, U_{\mathcal{L}}]$.** In this case, the natural bias of the process already aligns with the short-term requirements, so we can choose $\kappa = p$ and use a family of either polynomial or exponential functions. We show here a family of modified polynomial functions that satisfy the monotonicity requirements. The family of energy functions is $(\zeta_r)_{r \in R}$, with $R = (0, 1)$, built as follows. Let $\alpha_r = r/(1-r)$, and $m \geq 2$ fixed, $l_r = \kappa - \frac{1}{\alpha_r^{1/m}}$, $u_r = \kappa + \frac{1}{\alpha_r^{1/m}}$. Then

$$\zeta_r(x) = \begin{cases} \alpha_r |x - \kappa|^m & \text{if } x \in [l_r, u_r] \\ 1 & \text{otherwise.} \end{cases}$$

We use $m = 2$.

Claim 2. *The family of functions $(\zeta_r)_{r \in R}$ satisfies the following properties:*

- *It is monotonic in r , i.e., for each $x \in [0, 1]$ and each pair $s, r \in (0, 1)$, if $s \leq r$, then $\zeta_s(x) \leq \zeta_r(x)$.*
- *The corresponding characteristic function f_r (as defined in Eq. 2) has a fixpoint at $\mu = a_r$.*
- *If $p < L_{\mathcal{L}}$, for all $x \neq \kappa$:*

$$\lim_{r \rightarrow 0} \zeta_r(x) = \begin{cases} (L_{\mathcal{L}} - p)/(1 - p) & \text{if } x \leq L_{\mathcal{L}} \\ 0 & \text{otherwise.} \end{cases}, \quad \lim_{r \rightarrow 1} \zeta_r(x) = \begin{cases} (L_{\mathcal{L}} - p)/(1 - p) & \text{if } x \in [U_{\mathcal{L}}, \kappa] \\ 1 & \text{otherwise.} \end{cases}$$

- *If $p > U_{\mathcal{L}}$, for all $x \neq \kappa$:*

$$\lim_{r \rightarrow 0} \zeta_r(x) = \begin{cases} (p - L_{\mathcal{L}})/p & \text{if } x \geq U_{\mathcal{L}} \\ 0 & \text{otherwise.} \end{cases}, \quad \lim_{r \rightarrow 1} \zeta_r(x) = \begin{cases} (p - L_{\mathcal{L}})/p & \text{if } x \in [\kappa, L_{\mathcal{L}}] \\ 1 & \text{otherwise.} \end{cases}$$

- *If $p \in \mathcal{L}$, for all $x \neq \kappa$ and all $l \in \{0, 1\}$, $\lim_{r \rightarrow l} \zeta_r(x) = l$.*

All the properties are satisfied by construction. Fig. 2 illustrates them.