Runtime Analysis of Coevolutionary Algorithms on a Class of Symmetric Zero-Sum Games

Benford, Alistair; Lehre, Per Kristian

Runtime Analysis of Coevolutionary Algorithms on a Class of Symmetric Zero-Sum Games

Alistair Benford
University of Birmingham
Birmingham, UK

Per Kristian Lehre
University of Birmingham
Birmingham, UK

ABSTRACT

A standard aim in game theory is to find a pure or mixed Nash equilibrium. For strategy spaces too large for a Nash equilibrium to be computed classically, this can instead be approached using a coevolutionary algorithm. How to design coevolutionary algorithms which avoid pathological behaviours (such as cycling or forgetting) on challenging games is then a crucial open problem.

We argue that runtime analysis can provide insight and inform the design of more powerful and reliable algorithms for this purpose. To this end, we consider a class of symmetric zero-sum games for which the role of population diversity is pivotal to an algorithm's success. We prove that a broad class of algorithms which do not utilise a population have superpolynomial runtime for this class. In the other direction we prove that, with high probability, a coevolutionary instance of the univariate marginal distribution algorithm finds the unique Nash equilibrium in time $O(n(\log n)^2)$.

Together, these results demonstrate the importance of generating diverse search points for evolving better strategies. The corresponding proofs develop several techniques that may benefit future analysis of estimation of distribution and coevolutionary algorithms.

CCS CONCEPTS

• **Theory of computation → Theory of randomized search heuristics**.

KEYWORDS

runtime analysis, coevolution, game theory

ACM Reference Format:
Alistair Benford and Per Kristian Lehre. 2024. Runtime Analysis of Coevolutionary Algorithms on a Class of Symmetric Zero-Sum Games. In *Genetic and Evolutionary Computation Conference (GECCO '24), July 14–18, 2024, Melbourne, VIC, Australia.* ACM, New York, NY, USA, 20 pages. https://doi.org/10.1145/3638529.3654216

1 INTRODUCTION

For many real-world optimisation problems, the quality of a solution depends on the actions of adversaries, competitors, or opponents. The dynamics of these game theoretic settings are often highly complex, as a player's strategy can both influence and be influenced by the strategies of opposing players. As no strategy can be seen to be optimal in isolation, optimisers commonly define the best strategies to be those belonging to a (pure or mixed) Nash equilibrium [23, 27], wherein no player is able to improve their outcome by changing their own strategy, assuming their opponents' strategies are fixed.

If every possible combination of strategies can be queried then a Nash equilibrium can be computed directly (see [28]). However this is impractical if, as is the case for many real-world settings, the strategy space is too large to be searched exhaustively. In such cases, a possible approach is to use coevolution [25, 26], where populations of strategies are evolved based on principles of natural selection. Under this regime, individuals with the best interactions against their contemporaries are used as parents for the next generation.

The question of how to effectively apply coevolution to this setting is a deeply complicated one. The success of a coevolutionary algorithm can depend on a large number of design aspects (such as population size, mutation operators, selection mechanisms, or diversity mechanisms), which have subtle but critical effects on their dynamics. Further to this, coevolutionary algorithms are often prone to pathological behaviours such as cycling or loss of gradient [12]. How to reliably obtain intuition for designing a coevolutionary algorithm which avoids these behaviours is thus an essential consideration.

For standard evolutionary algorithms, which apply in the absence of strategic interaction, this intuition can be provided by runtime analysis, where rigorous results relate algorithm design and fitness landscape to runtime distribution [7]. Accordingly, there is clear demand (see [25]) for similar results which apply to coevolution. The first rigorous runtime analysis on this front was developed by Jansen and Wiegand [17], who showed that cooperative coevolutionary algorithms do not guarantee stronger performance on separable problems over traditional evolutionary algorithms. The first runtime analysis for competitive coevolution was provided by Lehre [18], who showed that a population-based coevolutionary algorithm called PDCoEA efficiently approximates the Nash equilibrium of certain instances of Bilinear, a two-player game played on bitstrings. In the same paper, it was also shown that if the mutation rate used by PDCoEA exceeds a specific error threshold, then the algorithm is inefficient. Hevia Fajardo and Lehre [16] provided further analysis on Bilinear, showing that a $(1, \lambda)$ coevolutionary algorithm efficiently discovers the Nash equilibrium when using worst interaction as a fitness measure, but has exponential runtime when using the average over all interactions instead. In addition, Hevia Fajardo, Lehre and Lin [10] gave runtime analysis of a $(1+1)$-type coevolutionary algorithm on Bilinear, with a particular focus on the use of archives to reduce arising pathological behaviours.

GECCO '24, July 14–18, 2024, Melbourne, VIC, Australia

ACM ISBN 979-8-4007-0494-9/24/07.
https://doi.org/10.1145/3638529.3654216

Together, these results provide insight into a range of design aspects of coevolutionary algorithms. However, with the exception of the error threshold result for PDCoEA, which applies to all problems without too many global optima, all competitive runtime results are limited to Bilinear, and the added complexity of the competitive coevolutionary setting has so far prevented the development of runtime analysis for more general or complicated games. Our aim is therefore to push the scope of runtime analysis of coevolutionary algorithms to games which more generally reflect the nature of real-world games, while remaining simple enough for rigorous mathematical analysis. Many real-world games incorporate a strong sense of skill, where high skill strategies reliably win against low skill strategies, but the payoff landscape for similarly matched strategies may be flat, random-like, or highly non-transitive. For such games, the need for a coevolutionary algorithm to have access to a diverse range of solutions is apparent.

Accordingly, the contributions of this paper concern the role of populations and diversity in optimising skill-based games. In particular, we demonstrate the effectiveness of employing estimation of distribution for this purpose due to the resulting generation of diverse search points. To represent the geometry of skill-based games, we introduce a class of symmetric zero-sum games played on bitstrings with a clear Nash equilibrium. We show that coevolutionary algorithms with low population diversity cannot guarantee good performance on this class, as certain instances result in draws for strategies that are too similar to each other, thus making search heuristics based solely on local comparisons impractical. Precisely, we will prove that for a highly general class of algorithms which retain only one strategy between generations, the probability that the Nash equilibrium will be found within e^{n^c} function evaluations is at most e^{-n^c}, for some constant c. Despite this, we will also show that success on this class can indeed be assured by using estimation of distribution, proving that, with high probability, a coevolutionary instance of the univariate marginal distribution algorithm finds the Nash equilibrium in time $O(n(\log n)^2)$.

After stating our notation, the structure of the paper is as follows. In Section 1.2, we will define symmetric zero-sum games and introduce the framework used to state our runtime results. In Section 1.3, we introduce the class of symmetric zero-sum games on which we perform our analysis. After concluding the introduction by stating a number of preliminary results for use throughout the paper, Section 2 will establish the aforementioned positive runtime result for the univariate marginal distribution algorithm (opting to present this result first as it is simpler to prove). Finally, the complementary result demonstrating poor runtime on the class for algorithms without population is covered in Section 3.

1.1 Notation

Throughout, let $\mathcal{X}_n := \{0,1\}^n$ denote the set of bitstrings of length n. Given $x \in \mathcal{X}_n$, we write $|x|$ to denote the number of 1-bits of x. Given $x, y \in \mathcal{X}_n$, let $d_H(x,y)$ denote the Hamming distance between x and y. Given $x \in \mathcal{X}_n$, we write $S_r(x) = \{y \in \mathcal{X}_n : d_H(x,y) = r\}$ for the *Hamming shell* of radius r, and write $B_r(x) = \cup_{0 \leqslant r' \leqslant r} S_{r'}(x)$ for the *Hamming ball* of radius r. Let $\mathbf{1}^n \in \mathcal{X}_n$ denote the bitstring with all entries equal to 1. Given $p := (p_1, \dots, p_n) \in [0,1]^n$, we say $x \sim \mathrm{Bit}_n(p)$ if x is generated by setting the i^{th} bit equal to 1 with probability p_i (and equal to 0 with probability $1 - p_i$) independently for every bit.

Given a finite set S, we use $\mathcal{P}(S)$ to denote the set of probability distributions on S. Given $p \in \mathcal{P}(S)$, write $\mathrm{supp}(p) := \{s \in S : p(s) > 0\}$ for the *support* of p. Given $p \in \mathcal{P}(S)$ and $A \subseteq S$, we write $p(A) = \sum_{s \in A} p(s)$.

For real-valued random variables X, Y, we say that X *stochastically dominates* Y, written $X \succcurlyeq Y$, if $\mathbb{P}(X \leqslant z) \leqslant \mathbb{P}(Y \leqslant z)$ holds for all $z \in \mathbb{R}$. Indicator functions are denoted using $\mathbb{1}$. Given a sigma algebra $\mathscr{F}$, event $E \in \mathscr{F}$, and random variable X, we write $\mathbb{E}[X \,;\, E \mid \mathscr{F}]$ to mean $\mathbb{E}[X \cdot \mathbb{1}(E) \mid \mathscr{F}]$.

We use $\mathrm{Mat}_{k,k}(\mathbb{R})$ to denote the set of real-valued $k \times k$ matrices. Given a natural number n, we write $[n] = \{1, \dots, n\}$. Logarithms are given in base 2 unless stated otherwise.

1.2 Symmetric zero-sum games

A *two-player game* is defined by strategy spaces $\mathcal{X}, \mathcal{Y}$ and payoff functions $f_i : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$, $i \in [2]$, where $f_i(x,y)$ denotes the payoff awarded to player i when player 1 adopts strategy x and player 2 adopts strategy y. The game is *zero-sum* if player 1's gain is equal to player 2's loss (and vice versa) – that is, if $f_2(x,y) = -f_1(x,y)$ for every $x \in \mathcal{X}$, $y \in \mathcal{Y}$. If $\mathcal{X} = \mathcal{Y}$ and $f_1(x,y) = f_2(y,x)$ for every $x, y \in \mathcal{X}$, then the game is called *symmetric*. Symmetric games describe many naturally arising real-world interactions [11], such as firms competing for market dominance through advertising [14, 15].

Every game which is both zero-sum and symmetric can be represented by a single antisymmetric function, as follows.

Definition 1.1. *A function $f : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ is* antisymmetric *if $f(x,y) = -f(y,x)$ holds for every $x, y \in \mathcal{X}$. Given an antisymmetric function f, we refer to the game with payoff functions $f_1(x,y) = f(x,y)$ and $f_2(x,y) = -f(x,y)$ as the* symmetric zero-sum game *defined by f.*

Any classical game where the outcomes are win/draw/lose, such as Tic-Tac-Toe, Chess, or Go, can be represented by such a function by identifying $f(x,y) = 1$ with a win for player 1, $f(x,y) = -1$ with a win for player 2, and $f(x,y) = 0$ with a draw. Accordingly, all such games are symmetric and zero-sum.

For games where no single strategy is unilaterally superior, such as rock paper scissors, the best available policy is to adopt a *mixed strategy* $p \in \mathcal{P}(\mathcal{X})$, where $p(x)$ corresponds to the probability of choosing strategy x. If the game is symmetric and zero-sum, then no mixed strategy can have strictly positive expected payoff against all other mixed strategies, as the expected payoff when played against itself is zero. Therefore, the strongest possible strategy would be one which delivers non-negative expected payoff against all others. Such a mixed strategy corresponds to the solution concept of a Nash equilibrium, stated in the language of symmetric zero-sum games as follows.

Definition 1.2. *Given an antisymmetric function $f : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ on a finite set $\mathcal{X}$, we say $p \in \mathcal{P}(\mathcal{X})$ is a* Nash equilibrium *for f if*

$$\min_{q \in \mathcal{P}(\mathcal{X})} \sum_{x,y \in \mathcal{X}} p(x) q(y) f(x,y) = 0.$$

If $\mathrm{supp}(p) = \{x\}$, then we say that x is a pure Nash equilibrium *for f. Additionally, let $p_{NE}(f)$ be the unique Nash equilibrium of maximal entropy.*

Note that for symmetric zero-sum games, a strategy x^* is a pure Nash equilibrium if and only if $f(x^*, y) \geqslant 0$ for all $y \in \mathcal{X}$ (to see this, observe that if $\mathrm{supp}(p) = \{x^*\}$ and $f(x^*, y) \geqslant 0$ for all $y \in \mathcal{X}$, then for any $q \in \mathcal{P}(\mathcal{X})$ we have $\sum_{x,y\in\mathcal{X}} p(x)q(y)f(x,y) = \sum_{y\in\mathcal{X}} q(y)f(x^*,y) \geqslant 0$). For a justification of the existence and uniqueness of $p_{\mathrm{NE}}(f)$, we refer the reader to Appendix C. For a general source on Nash equilibria, see [23].

As the discovery of a Nash equilibrium represents a clear goal for game-playing, we can define the runtime of a coevolutionary algorithm for this purpose by identifying the optimum with a Nash equilibrium. The definition we adopt here applies only to games which have a pure Nash equilibrium (as will be the case for both of our results), and further work on general games may require a broader solution concept. Note that we adopt the convention standard to black box optimisation where runtime is defined as the number of times the function f is queried by the algorithm until the desired search objective is discovered (see [9]).

Definition 1.3. *Suppose that $\mathcal{A}$ is an algorithm which makes τ queries of an antisymmetric function $f : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ during each generation. Suppose that f has a unique pure Nash equilibrium x^*. Then the* runtime $T(\mathcal{A}, f)$ *of $\mathcal{A}$ on f is defined to be the random variable*

$$T(\mathcal{A}, f) = \tau \cdot \min\{t : x^* \in P_t\},$$

where $P_t \subseteq \mathcal{X}$ is the population of $\mathcal{A}$ at the start of generation t.

1.3 Deceptive OneMax games

For the remainder of this paper, we only consider games where the possible outcomes are to win, draw, or lose, and so restrict our attention to antisymmetric functions $f : \mathcal{X} \times \mathcal{X} \to \{-1, 0, 1\}$. We will now state the class of symmetric zero-sum games we consider, before discussing the definition in more detail.

Definition 1.4. *Given $x^* \in \mathcal{X}_n$, let $\mathcal{F}_n^{x^*}$ be the set of antisymmetric functions $f : \mathcal{X}_n \times \mathcal{X}_n \to \{-1, 0, 1\}$ such that $f(x_1, x_2) = 1$ whenever $d_H(x_1, x^*) < d_H(x_2, x^*)$. Let $\mathcal{F}_n = \cup_{x^*\in\mathcal{X}_n} \mathcal{F}_n^{x^*}$.*

Definition 1.5. *Given $\alpha > 0$, $m \in \mathbb{N}$, and $x^* \in \mathcal{X}_n$, let $\mathcal{G}_n^{x^*}(\alpha, m)$ be the set of antisymmetric functions $g : \mathcal{X}_n \times \mathcal{X}_n \to \{-1, 0, 1\}$, such that, for every $g \in \mathcal{G}_n^{x^*}(\alpha, m)$, there exists $f \in \mathcal{F}_n^{x^*}$ such that*

- **A1** *$g(x, y) = f(x, y)$ if $d_H(x, x^*) < \alpha n$ or $d_H(y, x^*) < \alpha n$, and*
- **A2** *for every $x \in \mathcal{X}_n$, there are at most m many $y \in \mathcal{X}_n$ for which $g(x, y) \neq f(x, y)$.*

Given $\alpha > 0$ and $m \in \mathbb{N}$, let $\mathcal{G}_n(\alpha, m) = \cup_{x^*\in\mathcal{X}_n} \mathcal{G}_n^{x^*}(\alpha, m)$.

We remark that $\mathcal{F}_n \subseteq \mathcal{G}_n(\alpha, m)$ for any $\alpha > 0$, $m \in \mathbb{N}$, and that $\mathcal{G}_n(\alpha_1, m_1) \subseteq \mathcal{G}_n(\alpha_2, m_2)$ whenever $\alpha_1 \geqslant \alpha_2$ and $m_1 \leqslant m_2$.

$\mathcal{F}_n^{x^*}$ describes games for which a player wins whenever their bitstring has a smaller Hamming distance to the unique Nash equilibrium x^*. $\mathcal{G}_n^{x^*}(\alpha, m)$ describes games for which this is only true most of the time, as whenever competing the strategies are of Hamming distance at least αn from x^* (which remains the unique Nash equilibrium due to **A1**) there may be a 'wrong winner' outcome where the more distant strategy is perceived to be stronger.

Accordingly, optimising games in $\mathcal{F}_n$ is not much more difficult than optimising the well-known OneMax function in the standard evolutionary setting (see, for example, [1]), as a clear fitness signal permeates throughout the whole strategy space (albeit possibly muddied slightly by the fact that $f \in \mathcal{F}_n$ can take arbitrary values when strategies at the same distance from x^* are played against each other). However, because we will allow m to grow exponentially with n, instances of $\mathcal{G}_n(\alpha, m)$ can be far more challenging. As an example, consider for a parameter d,

$$f_0(x, y) = \begin{cases} 1 & \text{if } |x| > |y|, \\ 0 & \text{if } |x| = |y|, \\ -1 & \text{if } |x| < |y|, \end{cases}$$

$$g_d(x, y) = \begin{cases} -f_0(x, y) & \text{if } |x|, |y| \in [\alpha n, (1-\alpha)n] \text{ and } d_H(x, y) \leqslant d, \\ f_0(x, y) & \text{otherwise,} \end{cases}$$

where we note that any fixed $\beta > 0$, we have $g_d \in \mathcal{G}_n(\alpha, (1-\beta)^{-n})$ provided $d = O(n/\log n)$. In this case, methods which rely on local search alone will be tricked into minimising $|x|$ in pursuit of better strategies. Motivated by this, we will refer to instances of $\mathcal{G}_n(\alpha, m)$ as *m-deceptive OneMax games* (with distance-to-optimum threshold αn).

Part of the motivation for $\mathcal{G}_n(\alpha, m)$ is to reflect the geometry underpinning the strategy spaces of games that are interesting and challenging to human players. The Game of Skill hypothesis [2] is stated with respect to a decomposition of the strategy space $\mathcal{X}$ into *transitive skill layers* $\mathcal{X} = A_1 \cup \cdots \cup A_m$, where strategies in higher skill levels are generally superior to strategies in lower skill levels, but the payoff landscape within individual skill levels may be highly non-transitive. A symmetric zero-sum game is then called a *Game of Skill* if the middle skill layers are very large and contain a diverse range of strategies with their own relative strengths and weaknesses, but this richness disappears as we look towards the higher and lower skill levels, where skill takes precedence of style.

Adopting Nash clustering (that is, the iterated removal of maximal entropy mixed Nash equilibria from the strategy space) as a method for obtaining the skill layers, Czarnecki et al. [2] show that if a population P_t includes a full Nash cluster A_i and then trains by seeking $x \in \mathcal{X}$ such that $f(x, y) > 0$ for every $y \in A_i$, then improvement with respect to the Nash clustering is assured. However, even though f_0 and g_d fit the game of skill description with layers $A_j := \{x \in \mathcal{X}_n : |x| = j\}$ (noting that values of $f_0(x, y)$ for $|x| = |y|$ can be varied arbitrarily to give more richness within middle layers), two limitations arise with this approach.

1. The middle Nash clusters have size $|A_{n/2}| = \Omega(2^n/\sqrt{n})$, and so it is not computationally practical to use a population size capable of covering the full cluster.
2. In the case of g_d, it is not true in general that there exists a strategy $x \in A_j$ which wins against every $y \in A_{j-1}$ (such x may exist in much higher levels than A_j, but these are difficult to discover in a single step).

In light of these observations, $\mathcal{G}_n(\alpha, m)$ constitutes a class of of symmetric zero-sum games which are challenging to optimise, despite being constructed from the simple OneMax function.

1.4 Drift theorems

Here we quote two drift theorems to invoke later in our proofs. The first is an upper tail bound for multiplicative drift [6, 20] and the second is a negative drift theorem [29].

Theorem 1.6. *Let $(X_t)_{t=0}^{\infty}$ be a stochastic process adapted to a filtration $(\mathscr{F}_t)_{t=0}^{\infty}$, taking values in a finite subset of $\{0\} \cup [x_{\min}, \infty)$ where $x_{\min} > 0$. Suppose that there exists $\delta > 0$ such that $\mathbb{E}[X_t - X_{t+1} \mid \mathscr{F}_t] \geqslant \delta X_t$ whenever $X_t > 0$. Then, if $X_0 > 0$, it holds for the first hitting time $T := \min\{t : X_t = 0\}$ that*

$$\mathbb{P}\left[T > \lceil (\ln(X_0/x_{\min}) + r)/\delta \rceil \mid \mathscr{F}_0\right] \leqslant e^{-r}.$$

Theorem 1.7. *There is an absolute constant $\overline{C} > 0$ such that the following holds. Let $(X_t)_{t=0}^{\infty}$ be a stochastic process adapted to a filtration $(\mathscr{F}_t)_{t=0}^{\infty}$. Suppose there exists an interval $[a, b] \subseteq \mathbb{R}$ and positive numbers ε, κ, r (each possibly depending on $\ell := b - a$), as well as a sequence of functions $\Delta_t := \Delta_t(X_{t+1} - X_t)$ satisfying $\Delta_t \leqslant X_{t+1} - X_t$, such that the following conditions hold for all $t \geqslant 0$.*

B1 $\mathbb{E}[\Delta_t \cdot \mathbb{1}(\Delta_t \leqslant \kappa\varepsilon) - \varepsilon\,;\, a < X_t < b \mid \mathscr{F}_t] \geqslant 0$.

B2 *If $a < X_t$ then $\mathbb{P}(\Delta_t \leqslant -jr \mid \mathscr{F}_t) \leqslant e^{-j}$ for all $j \in \mathbb{N}$.*

B3 $\lambda\ell \geqslant 2\ln(4/(\lambda\varepsilon))$ *where* $\lambda := \min\{1/(2r), \varepsilon/(17r^2), 1/(\kappa\varepsilon)\}$.

Then, if $X_0 \geqslant b$ it holds for the first hitting time $T := \min\{t : X_t \leqslant a\}$ that

$$\mathbb{P}[T \leqslant e^{\lambda\ell/4} \mid \mathscr{F}_0] \leqslant \overline{C} e^{-\lambda\ell/4}.$$

2 AN UPPER BOUND FOR THE RUNTIME OF UMDA ON DECEPTIVE ONEMAX GAMES

Instead of storing a population as an explicit set of individuals, estimation of distribution algorithms (EDAs) implicitly represent their current population as a probability distribution on the search space [24]. In each generation, a number of search points are sampled according to the current distribution, and the distribution is updated depending on the function values taken at those search points. One strength of EDAs is the high level of diversity among generated search points. Indeed, Doerr [4] showed that the compact Genetic Algorithm (cGA) can use this diversity to achieve a runtime of $O(n \log n)$ on jump functions with jump size $k = O(\log n)$, a significant improvement over the runtime of $\Omega(n^k)$ needed for most classical evolutionary algorithms. A result of Witt [30] shows that a similar speedup occurs even if the optimum is shifted. EDAs may be also uniquely relevant to game theoretic settings, as a stored probability distribution can be interpreted as a mixed strategy. Despite this connection, and other successful runtime analysis of EDAs in non-coevolutionary settings [3, 8, 29], our work constitutes the first runtime analysis for a coevolutionary EDA.

The simplest EDAs that operate on the search space $\mathcal{X}_n = \{0, 1\}^n$ make no direct attempt to correlate bits. The Univariate Marginal Distribution Algorithm (UMDA) stores the current population as a bit frequency vector $(p_1, \dots, p_n)$, where p_i denotes the probability that the i^{th} bit is equal to 1 for a sampled individual. For classical optimisation of unary functions $f : \mathcal{X}_n \to \mathbb{R}$ the most general form of UMDA (as originally stated in [22]) is given by Algorithm 1 (where we recall $\mathrm{Bit}_n(p_1, \dots, p_n)$, defined in Section 1.1, corresponds to UMDA's method for search point generation). Central to the algorithm's design is a choice of selection operator $\mathcal{S} : \mathbb{R}^\lambda \to \mathcal{P}([\lambda]^\mu)$. Given individuals $x_1, \dots, x_\lambda$, if $(j_1, \dots, j_\mu) \sim \mathcal{S}(f(x_1), \dots, f(x_\lambda))$, then $(j_1, \dots, j_\mu)$ should correspond to the indices of individuals from whom the the bit frequencies for the next generation will be derived. The most commonly adopted standard is to assume $\mathcal{S}$ deterministically selects the μ indices corresponding to the highest μ values of $f(x_j)$.

In game theoretic settings, where we have $f : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$, the fitness of an individual cannot be evaluated in isolation, and so a selection procedure based on player interaction must be adopted. For our analysis, we opt to use a selection procedure similar to a binary tournament selection. The full version of our coevolutionary instance of UMDA is then given by Algorithm 2. To eliminate the possibility of infinite runtime, most formulations of UMDA additionally ensure that the frequencies p_i cannot be equal to either 0 or 1, typically by constraining them to the interval $[1/n, 1 - 1/n]$. For the sake of simplicity, we forego this restriction, but we see no reason why the result would not hold with its inclusion.

Algorithm 1 UMDA with selection operator $\mathcal{S} : \mathbb{R}^\lambda \to \mathcal{P}([\lambda]^\mu)$

Require: Function $f : \mathcal{X}_n \to \mathbb{R}$.
Require: Algorithm parameters $\mu, \lambda \in \mathbb{N}$.
for $i \in [n]$ **do**
  Set $p_{0,i} = \frac{1}{2}$.
end for
for $t \in \mathbb{N}$ until termination criterion met **do**
  for $j \in [\lambda]$ **do**
    Sample $x_j \sim \mathrm{Bit}_n(p_{t,1}, \dots, p_{t,n})$.
  end for
  Sample $(j_1, \dots, j_\mu) \sim \mathcal{S}(f(x_1), \dots, f(x_\lambda))$.
  for $i \in [n]$ **do**
    Set $p_{t+1,i} = \frac{1}{\mu}|\{k \in [\mu] : x_{j_k}$ has a 1-bit in position $i\}|$.
  end for
end for

Algorithm 2 UMDA with binary tournament selection

Require: Antisymmetric function $f : \mathcal{X}_n \times \mathcal{X}_n \to \mathbb{R}$.
Require: Algorithm parameter $\mu \in \mathbb{N}$.
for $i \in [n]$ **do**
  Set $p_{0,i} = \frac{1}{2}$.
end for
for $t \in \mathbb{N}$ until termination criterion met **do**
  for $j \in [\mu]$ **do**
    Sample $x \sim \mathrm{Bit}_n(p_{t,1}, \dots, p_{t,n})$
    Sample $y \sim \mathrm{Bit}_n(p_{t,1}, \dots, p_{t,n})$
    if $f(x, y) > 0$ **then**
      Set $P_{t+1}(j) = x$
    else if $f(x, y) < 0$ **then**
      Set $P_{t+1}(j) = y$
    else if $f(x, y) = 0$ **then**
      Sample $P_{t+1}(j) \sim \mathrm{Unif}(\{x, y\})$
    end if
  end for
  for $i \in [n]$ **do**
    Set $p_{t+1,i} = \frac{1}{\mu}|\{j : P_{t+1}(j)$ has a 1-bit in position $i\}|$
  end for
end for

We now state the main result for this section, which shows that UMDA is able to efficiently optimise m-deceptive OneMax games, even if m grows exponentially with n.

Theorem 2.1. *Let $\mathcal{A}$ be described by Algorithm 2. There is a constant $C > 0$ and a function $n_0 : \mathbb{R}^3 \to \mathbb{R}$ such that following holds. Suppose $0 < \delta \leqslant \alpha < 1/2$, $K \geqslant 1$, and*

$$n \geqslant n_0(\alpha, \delta, K), \tag{1}$$

$$\frac{CK}{\delta}\sqrt{n}\ln n \leqslant \mu \leqslant 2e^{\delta^2 n/4}. \tag{2}$$

Then, for any $g \in \mathcal{G}_n(\alpha, (1-\alpha+\delta)^{-n})$,

$$\mathbb{P}[T(\mathcal{A}, g) \geqslant 50K\mu\sqrt{n}\log n] \leqslant n^{-K}.$$

We remark that this result also implies that a non-coevolutionary version of UMDA using binary tournament selection optimises OneMax in time $O(n(\log n)^2)$. This almost matches the $O(\lambda\sqrt{n})$ when $\lambda = \Omega(\sqrt{n}\log n)$ bound proven by Witt [29] for the standard UMDA. In fact, interpreted in the setting where the OneMax function cannot be directly queried but instead only compared for pairs of search points, the two results broadly match when accounting for the fact that $\Omega(\lambda \log \lambda)$ comparisons (rather than λ fitness evaluations) are needed to order $x_1, \ldots, x_\lambda$.

Central to the proof of Theorem 2.1 is showing that, provided no frequency has moved a significant distance away from the optimum, Algorithm 2 is extremely unlikely to generate a pair x, y such that $g(x, y)$ differs from the corresponding $\mathcal{F}_n$ function. This is handled by the following straightforward lemma, which is proven in Appendix B.1.

Lemma 2.2. *Let $\alpha, \delta > 0$, and let $g \in \mathcal{G}_n^{1^n}(\alpha, (1-\alpha+\delta)^{-n})$ and $f \in \mathcal{F}_n^{1^n}$ satisfy* **A1** *and* **A2** *(see Definition 1.5). Suppose $p = (p_1, \ldots, p_n) \in [\frac{1}{2} - \frac{\delta}{4}, 1]^n$ and that x, y are sampled independently according to $\mathrm{Bit}_n(p)$. Then, $\mathbb{P}(g(x, y) \neq f(x, y)) \leqslant e^{-\delta^2 n/8}$.*

We are now ready to prove Theorem 2.1. We will prove the result with $C = 10^5$, although this is only used directly in Appendix B.4 and a smaller value of C will likely suffice. We do not describe the function n_0 explicitly, but instead assume that n is always sufficiently large for any relevant bounds to hold.

Proof of Theorem 2.1. Let $g \in \mathcal{G}_n^{x^*}(\alpha, (1-\alpha+\delta)^{-n})$ and let $f \in \mathcal{F}_n^{x^*}$ be such that **A1** and **A2** hold. Without loss of generality, we may assume that $x^* = \mathbf{1}^n$. Let $(\mathscr{F}_t)_{t=0}^\infty$ be the filtration generated by $((p_{t,1}, \ldots, p_{t,n}))_{t=0}^\infty$. For $t \geqslant 0$ and $i \in [n]$, define

$$q_{t,i} := \mathbb{E}[p_{t+1,i} \mid \mathscr{F}_t]$$

so that $((q_{t,1}, \ldots, q_{t,n}))_{t=0}^\infty$ is a stochastic process adapted to $(\mathscr{F}_t)_{t=0}^\infty$. Equivalently, $q_{t,i}$ is the probability (in terms of $(p_{t,1}, \ldots, p_{t,n})$) that the individual $P_{t+1}(1)$ has a 1-bit in position i. Therefore, for each $t \geqslant 0$, because $P_t(1), \ldots, P_t(\mu)$ are independent and identically distributed,

$$(\mu \cdot p_{t+1,i} \mid \mathscr{F}_t) \sim \mathrm{Bin}(\mu, q_{t,i}). \tag{3}$$

The following claim, which is proven in Appendix B.2, will give us the required drift on the bit frequencies.

Claim 2.3. *If $(p_{t,1}, \ldots, p_{t,n}) \in [\frac{1}{2} - \frac{\delta}{4}, 1]^n$, then*

$$q_{t,i} \geqslant p_{t,i}\left(1 + \frac{1 - p_{t,i}}{5\sqrt{n}}\right) - 2e^{-\delta^2 n/8}.$$

For $t \geqslant 0$, define

$$X_t = \begin{cases} n - \sum_{i\in[n]} p_{t,i} & \text{if } (p_{t,1}, \ldots, p_{t,n}) \in [\frac{1}{2} - \frac{\delta}{4}, 1]^n \\ & \text{and } \sum_{i\in[n]} p_{t,i} \leqslant n-1, \\ 0 & \text{otherwise,} \end{cases} \tag{4}$$

so that $(X_t)_{t=0}^\infty$ is a stochastic process adapted to $(\mathscr{F}_t)_{t=0}^\infty$, taking values in a finite subset of $\{0\} \cup [1, \infty)$. We now define the following hitting times.

$$T_0 = \min\{t : X_t = 0\},$$
$$T_{\text{good}} = \min\{t : \textstyle\sum_{i\in[n]} p_{t,i} > n-1\},$$
$$T_{\text{bad}} = \min\{t : (p_{t,1}, \ldots, p_{t,n}) \notin [\tfrac{1}{2} - \tfrac{\delta}{4}, 1]^n\}.$$

Note that

$$T_0 \overset{(4)}{=} \min\{T_{\text{good}}, T_{\text{bad}}\}. \tag{5}$$

If for some generation t we find $\sum_{i\in[n]} p_{t,i} > n-1$, then the bitstrings $P_t(1), \ldots, P_t(\mu)$ share more than $\mu(n-1)$ 1-bits between them, and so we must have $P_t(j) = \mathbf{1}^n$ for some j. In particular, $\mathcal{A}$ must have sampled the pure Nash equilibrium $\mathbf{1}^n$ within time T_{good}, and so because $\mathcal{A}$ makes μ queries of f per generation, we have $T(\mathcal{A}, g) \leqslant \mu T_{\text{good}}$. Therefore,

$$\begin{aligned}\mathbb{P}[T(\mathcal{A}, g) \geqslant 50K\mu\sqrt{n}\ln n] &\leqslant \mathbb{P}[T_{\text{good}} \geqslant 50K\sqrt{n}\log n] \\ &\overset{(5)}{\leqslant} \mathbb{P}[T_0 \geqslant 50K\sqrt{n}\log n] + \mathbb{P}[T_{\text{bad}} \leqslant 50K\sqrt{n}\log n].\end{aligned}$$

Thus, all that remains is to prove the following claims.

Claim 2.4. $\mathbb{P}[T_0 \geqslant 50K\sqrt{n}\ln n] \leqslant \frac{1}{2}n^{-K}$.

Claim 2.5. $\mathbb{P}[T_{\text{bad}} \leqslant 50K\sqrt{n}\log n] \leqslant \frac{1}{2}n^{-K}$.

Claim 2.4 follows from a fairly straightforward application of Claim 2.3 to Theorem 1.6. While Claim 2.5 is proven using Claim 2.3 with Theorem 1.7, the application is far less direct and requires some careful handling of the relevant stochastic processes. For this, we use a coupling argument which may be of independent interest. For proofs of these claims, we refer the reader to Appendices B.3 and B.4. □

3 A LOWER BOUND FOR THE RUNTIME OF SINGLE-INDIVIDUAL ALGORITHMS ON DECEPTIVE ONEMAX GAMES

As a complementary result to Theorem 2.1, in this section we will prove that a broad class of single-individual algorithms, which retain only one search point as the current population in between generations, cannot guarantee a polynomial runtime on all instances of $\mathcal{G}_n(\alpha, (1-\beta)^{-n})$. No assumption is made on the initial distribution used to sample the algorithm's first search point. In terms of mutation, the only assumption is that mutation is unbiased [19], adopting the characterisation of Lemma 1 of [5] where a number $r \in \{0, \ldots, n\}$ is sampled according to a probability distribution s, and then a set of r bits is selected uniformly at random to flip. (For simplicity, we choose to restrict r to $\{0, \ldots, \lfloor n/2 \rfloor\}$, but remark that our results still follow in the absence of this restriction with some additional details.)

Definition 3.1. *Given $s \in \mathcal{P}(\{0,\ldots,\lfloor n/2\rfloor\})$, let $\mathcal{M}_s$ be the mutation operator such that, if $y \sim \mathcal{M}_s(x)$, then for every $r \geqslant 0$ and $z \in S_r(x)$,*

$$\mathbb{P}(y = z) = \frac{s(r)}{|S_r(x)|} = \frac{s(r)}{\binom{n}{r}}.$$

Our memory-restricted model is described by Algorithm 3. Generation t begins with the current individual x_t, from which offspring $y_1, \ldots, y_\mu$ are generated using $\mathcal{M}_s$. Then, based solely on the values of $f(x, y)$ for $x, y \in \{y_1, \ldots, y_\mu, x_t\}$, the algorithm selects (perhaps probabilistically) an x_{t+1} from $\{y_1, \ldots, y_\mu, x_t\}$. By identifying $y_{\mu+1}$ with x_t, this selection operator can be described formally as a map from the set of all possible $(\mu+1)\times(\mu+1)$ payoff matrices to the set of distributions on the index set $[\mu+1]$, and so we adopt the following notation.

Definition 3.2. *Given a function $f : \mathcal{X}\times\mathcal{X} \to \mathbb{R}$ and elements $x_1, \ldots, x_k \in \mathcal{X}$, we use $A_f(x_1, \ldots, x_k)$ to denote the $k\times k$ real-valued matrix A with entries given by $A_{i,j} = f(x_i, x_j)$.*

Algorithm 3 Single individual CoEA

Require: Antisymmetric function $f : \mathcal{X}_n \times \mathcal{X}_n \to \mathbb{R}$
Require: Offspring size $\mu \in \mathbb{N}$
Require: Initial distribution $p_{\text{init}} \in \mathcal{P}(\mathcal{X}_n)$
Require: Unbiased mutation distribution $s \in \mathcal{P}(\{0,\ldots,\lfloor n/2\rfloor\})$
Require: Selection operator $\mathcal{S} : \mathrm{Mat}_{\mu+1,\mu+1}(\mathbb{R}) \to \mathcal{P}([\mu+1])$

```
Sample x_0 according to p_init
for t ∈ ℕ until termination criterion met do
    for i ∈ [μ] do
        Sample y_i ~ M_s(x_t).
    end for
    Set y_{μ+1} = x_t
    Sample j ~ S(A_f(y_1, ..., y_{μ+1}))
    Set x_{t+1} = y_j
end for
```

We now state the main result for this section. In consideration of Definition 1.3, we assume in the following that sampling from $\mathcal{S}$ requires at least $\lceil \mu/2\rceil$ queries of f, as this is the smallest number needed to include every new search point in at least one query.

Theorem 3.3. *There exists $c > 0$ and a function $n_0 : \mathbb{R}^2 \to \mathbb{N}$ such that the following holds. If $\alpha, \beta \in (0, 1/2)$ and $n \geqslant n_0(\alpha, \beta)$, then there exists $g \in \mathcal{G}_n(\alpha, (1-\beta)^{-n})$ such that, if $\mathcal{A}$ is described by Algorithm 3, then*

$$\mathbb{P}[T(\mathcal{A}, g) \leqslant e^{n^c}] \leqslant e^{-n^c}.$$

The instance of $\mathcal{G}_n(\alpha, (1-\beta)^n)$ we will use to prove Theorem 3.3 is similar to the function g_d described in Section 1.3, where a reliable fitness signal is generally unavailable for inputs x, y with Hamming distance $d_H(x, y) \leqslant d = \Theta(n/\log n)$. The absence of a local fitness signal makes it difficult to exploit unbiased mutations of small Hamming distance, whereas unbiased mutations of large Hamming distance are vulnerable to genetic drift, thus making such a function challenging for Algorithm 3.

3.1 Dynamics of the initial distribution

The following lemma will be used to show that for any initial distribution p_{init}, we can choose x^* such that p_{init} is no more likely to generate a search point close to x^* than if we had used the uniform distribution on $\mathcal{X}_n$ in place of p_{init}. Its simple proof is given in Appendix B.5.

Lemma 3.4. *For any $p \in \mathcal{P}(\mathcal{X}_n)$ there exists $x^* \in \mathcal{X}_n$ such that, if x is sampled according to p and $0 \leqslant m \leqslant n/2$, then*

$$\mathbb{P}(d_H(x, x^*) \leqslant m) \leqslant 2^{(H(m/n)-1)n},$$

where $H(q) := -q\log q - (1-q)\log(1-q)$ is the binary entropy function.

3.2 Dynamics of the mutation operator

The following lemma establishes some useful properties common to all unbiased mutation operators $\mathcal{M}_s$. These properties together quantitatively describe the phenomenon that unbiased mutations cause drift towards bitstrings with an even distribution of 0-bits and 1-bits. Roughly speaking, the first property states that if $y \sim \mathcal{M}_s(x)$, then with high probability either $d_H(x, y)$ is small or $|y|$ is significantly closer to $\frac{1}{2}n$ than $|x|$ is. The second property states that if $y \sim \mathcal{M}_s(x)$, then the random variable $|y|$ has an exponential upper tail. Finally, the third property states, with a precise bounding factor, that if $|x| \gg \frac{1}{2}n$ and $y \sim \mathcal{M}_s(x)$, then $|y|$ is more likely to be smaller than $|x|$ rather than larger. For the statement, we recall that $B_r(x)$ is used to denote the Hamming ball of radius r. The proof is deferred to Appendix B.6.

Lemma 3.5. *Given $\varepsilon, \eta_0 > 0$, the following holds for any $s \in \mathcal{P}(\{0,\ldots,\lfloor n/2\rfloor\})$ provided n is sufficiently large. Given $x \in \mathcal{X}_n$, let $p_x \in \mathcal{P}(\mathcal{X}_n)$ be the probability mass function corresponding to the unbiased mutation operator $\mathcal{M}_s(x)$. Set $\eta = \eta_0/\log n$ and, given j, write $A_j = \{y \in \mathcal{X}_n : |y| = j\}$ and $A_{\geqslant j} = \cup_{j' \geqslant j} A_{j'}$. The following properties then hold.*

C1 *If $|x| < (\frac{1}{2} + 2\varepsilon + 3\eta)n$ and $d = 100\eta n/\varepsilon$, then*

$$p_x\left(A_{\geqslant(\frac{1}{2}+2\varepsilon)n} \setminus B_{d/2}(x)\right) \leqslant e^{-\sqrt{n}}.$$

C2 *For any $x \in \mathcal{X}_n$, if $m = \max\{|x|, (\frac{1}{2}+2\varepsilon)n\}$ and $j \geqslant 0$ then*

$$p_x\left(A_{\geqslant m+j}\right) \leqslant e^{-8\varepsilon j}.$$

C3 *If $(\frac{1}{2} + 2\varepsilon + \eta)n < |x| < (\frac{1}{2} + 2\varepsilon + 3\eta)n$ and $0 < k \leqslant \eta n$ then*

$$p_x\left(A_{|x|+k}\right) \leqslant (1-4\varepsilon)^k \cdot p_x\left(A_{|x|-k}\right).$$

3.3 Dynamics of the selection operator

Roughly speaking, the following lemma (which is proven in Appendix B.7) will be used to show that if $D \subseteq \mathcal{X}_n$ is a 'deceptive region', then the dynamics of Algorithm 3 will look like a random walk inside D.

Lemma 3.6. *Let $\mathcal{X}$ be a finite set, and let $f : \mathcal{X}\times\mathcal{X} \to \mathbb{R}$ be an antisymmetric function. Let $p \in \mathcal{P}(\mathcal{X})$, $\mu \in \mathbb{N}$, and let $\mathcal{S}$ be a map from $\mathrm{Mat}_{\mu+1,\mu+1}(\mathbb{R})$ to $\mathcal{P}([\mu+1])$. Suppose there are constants $f_0, f_1 \in \mathbb{R}$ and a non-empty subset $D \subseteq \mathcal{X}$ such that $f(x, y) = f_0$ whenever $x, y \in D$ and $f(x, y) = f_1$ whenever $x \in D$ and $y \in \mathrm{supp}(p) \setminus D$. Let $\overline{x} \in \mathcal{X}$ be fixed.*

Let y be the random variable defined by sampling $y_1, \ldots, y_\mu$ independently according to p, then sampling $j \sim \mathcal{S}(A_f(y_1, \ldots, y_\mu, \overline{x}))$, and finally setting $y = y_j$ (where we identify $y_{\mu+1} = \overline{x}$). Then there exists a constant $c \in [0, 1]$ such that $\mathbb{P}(y = x) = c \cdot p(x)$ for every $x \in D \setminus \{\overline{x}\}$.

3.4 Proof of Theorem 3.3

We will take $c = 1/8$ in the following proof of Theorem 3.3. Even though a higher value of c is implied by the proof, we opt to minimise the amount of detailed analysis of exact constants. For the same reason, we do not describe the function n_0 explicitly, but instead assume that n is always sufficiently large for any appropriate bounds to hold.

Proof of Theorem 3.3. Set

$$d = \frac{n \log(1/(1-\beta))}{2 \log n}, \qquad \varepsilon = \frac{\frac{1}{2} - \alpha}{3},$$
$$\eta = \frac{\varepsilon d}{100 n} = \frac{\varepsilon \log(1/(1-\beta))}{200 \log n}, \qquad \gamma = e^{-\varepsilon n^{1/3}}.$$

By Lemma 3.4, there exists some $x^* \in \mathcal{X}_n$ such that

$$\mathbb{P}(d_H(x_0, x^*) \leqslant (\tfrac{1}{2} - 2\varepsilon)n) \leqslant 2^{(H(\frac{1}{2} - 2\varepsilon) - 1)n} \leqslant \gamma^{10}.$$

By relabelling bits, we may assume without loss of generality that $x^* = \mathbf{1}^n$.

Let

$$A = \{x \in \mathcal{X}_n : |x| < (\tfrac{1}{2} + 2\varepsilon)n\},$$
$$B = \{x \in \mathcal{X}_n : (\tfrac{1}{2} + 2\varepsilon)n \leqslant |x| < (\tfrac{1}{2} + 2\varepsilon + 3\eta)n\},$$
$$C = \{x \in \mathcal{X}_n : (\tfrac{1}{2} + 2\varepsilon + 3\eta)n \leqslant |x|\}.$$

Note that

$$\mathbb{P}(x_0 \notin A) \leqslant \gamma^{10}. \tag{6}$$

Let $f \in \mathcal{F}_n^{\mathbf{1}^n}$ be the function given by

$$f(x, y) = \begin{cases} 1 & \text{if } |x| > |y|, \\ 0 & \text{if } |x| = |y|, \\ -1 & \text{if } |x| < |y|. \end{cases}$$

Let $g : \mathcal{X}_n \times \mathcal{X}_n \to \{-1, 0, 1\}$ be the function given by

$$g(x, y) = \begin{cases} 0 & \text{if } x, y \in B \text{ and } d_H(x, y) \leqslant d, \\ f(x, y) & \text{otherwise.} \end{cases} \tag{7}$$

Note that for any $x \in \mathcal{X}_n$,

$$|B_d(x)| = \sum_{r=0}^{d} |S_r(x)| \leqslant \sum_{r=0}^{d} n^r \leqslant (d+1) \cdot n^d$$
$$\leqslant n^{2d} = (1-\beta)^{-n},$$

and so $g \in \mathcal{G}_n^{\mathbf{1}^n}(\alpha, (1-\beta)^{-n})$.

Let $(x_t)_{t=0}^{\infty}$ denote the individuals generated by a run of Algorithm 3 on g. Let $T_C = \min\{t : x_t \in C\}$ so that, because each generation requires at least $\lceil \mu/2 \rceil$ function evaluations, we have $T(\mathcal{A}, g) \geqslant \lceil \mu/2 \rceil \cdot T_C$. We will show that $\mathbb{P}[\lceil \mu/2 \rceil \cdot T_C \leqslant e^{n^{1/8}}] \leqslant e^{-n^{1/8}}$, thus proving the desired result. Note that if $\mu > 1/\gamma$ then

$$\mathbb{P}[\lceil \mu/2 \rceil \cdot T_C \leqslant e^{n^{1/8}}] \leqslant \mathbb{P}[\lceil \mu/2 \rceil \cdot T_C \leqslant 1/(2\gamma)] = \mathbb{P}[T_C = 0]$$
$$= \mathbb{P}(x_0 \in C) \leqslant \mathbb{P}(x_0 \notin A) \overset{(6)}{\leqslant} \gamma^{10} \leqslant e^{-n^{1/8}},$$

and so we may additionally proceed under the assumption that

$$\mu \leqslant 1/\gamma. \tag{8}$$

It is clear from the algorithm description that $(x_t)_{t=0}^{\infty}$ is a Markov chain. In fact, we can also show the following.

Claim 3.7. *$(|x_t|)_{t=0}^{\infty}$ is a Markov chain taking values in the state space $S := \{0, \ldots, n\}$.*

Sketch proof of Claim 3.7. This follows from the fact that the dynamics of Algorithm 3, as well as the values taken by g, are unaffected by permuting bit positions. For a complete proof, see Appendix B.8. ⊡

Let us denote the transition probabilities of the chain $(|x_t|)_{t=0}^{\infty}$ as

$$q_{i,j} := \mathbb{P}(|x_{t+1}| = j \mid |x_t| = i). \tag{9}$$

Define $C(\varepsilon) = 8\varepsilon + 2 \cdot \max_{k \geqslant 0} \{k(1-4\varepsilon)^k\}$ and

$$h(k) = \min\{k, C(\varepsilon)\}. \tag{10}$$

In the following claims we collect two properties that will be used in the drift analysis later. While the proofs of these claims are deferred to Appendix B.9, we remark that Claim 3.8 follows from a simple application of **C2**, whereas the proof of Claim 3.9 involves a more detailed application of **C3** together with Lemma 3.6.

Claim 3.8. *If $0 \leqslant i \leqslant n$ and $i' = \max\{i, (\frac{1}{2} + 2\varepsilon)n\}$, then for any $j > 0$,*

$$\sum_{k=j}^{n-i'} q_{i,i'+k} \leqslant e^{\varepsilon(n^{1/3} - 8j)}.$$

Sketch proof of Claim 3.8. Apply **C2** with a union bound. For a complete proof, see Appendix B.9. ⊡

Claim 3.9. *If $(\frac{1}{2} + 2\varepsilon + \eta)n < i < (\frac{1}{2} + 2\varepsilon + 2\eta)n$ then*

$$\sum_{k=i-n}^{i} h(k) \cdot q_{i,i-k} \geqslant 2\varepsilon(1 - q_{i,i}) - \gamma^8.$$

Sketch proof of Claim 3.9. Suppose that x_t satisfies $|x_t| = i$. $\sum_{k=i-n}^{i} h(k) \cdot q_{i,i-k}$ is then the expected value of $h(|x_t| - |x_{t+1}|)$. Using **C1**, we can approximate $q_{i,i-k}$ by assuming that the mutants of x_t all lie in $A \cup B_{d/2}(x_t)$. Under this regime, Lemma 3.6 implies that if $|x_{t+1}|$ departs from $|x_t|$ (which introduces the factor of $(1 - q_{i,i})$), but does not fall down to A (in which case $h(|x_t| - |x_{t+1}|)$ is large), then the distribution of x_{t+1} looks like a random mutation of x_t, in which case we expect $|x_{t+1}|$ to be smaller than $|x_t|$ on average due to **C3**. For a complete proof, see Appendix B.9. ⊡

Next, recalling $S = \{0, \ldots, n\}$, define

$$V = \{\, i \in S : 1 - q_{i,i} \leqslant e^{-\varepsilon n^{1/3}} \text{ and } i < (\tfrac{1}{2} + 2\varepsilon + 2\eta)n\}, \tag{11}$$
$$W = \{\, i \in S : i \geqslant (\tfrac{1}{2} + 2\varepsilon + 2\eta)n\}.$$

Let $(Y_t)_{t=0}^{\infty}$ be the Markov chain with the same initial distribution as $(|x_t|)_{t=0}^{\infty}$, but with transition probabilities given by

$$P_Y(i,j) = \begin{cases} q_{i,j} & \text{if } i \in V \cup W, \\ \frac{q_{i,j}}{1-q_{i,i}} & \text{if } i \notin V \cup W \text{ and } i \neq j, \\ 0 & \text{if } i \notin V \cup W \text{ and } i = j. \end{cases} \tag{12}$$

Claim 3.10. *If $R \subseteq S$, then the first hitting time* $\min\{t : Y_t \in R\}$ *is stochastically dominated by* $\min\{t : |x_t| \in R\}$.

Sketch proof of Claim 3.10. $(Y_t)_{t=0}^{\infty}$ can be simulated by considering a run of $(|x_t|)_{t=0}^{\infty}$ and skipping any t for which $|x_t| = |x_{t-1}| = i$ for some $i \notin V \cup W$. In this way, $(Y_t)_{t=0}^{\infty}$ is simply an accelerated version of $(|x_t|)_{t=0}^{\infty}$. For a complete proof, see Appendix B.10. ⊡

Let us define the following hitting times for Y.

$$T_Y^{\text{hit}}(V) = \min\{t : Y_t \in V\}, \qquad T_Y^{\text{hit}}(W) = \min\{t : Y_t \in W\}.$$

We are now ready to define the stochastic process to which drift analysis will be applied. Let $(\mathscr{G}_t)_{t=0}^{\infty}$ denote the filtration generated by $(Y_t)_{t=0}^{\infty}$. Let $(Z_t)_{t\geqslant 0}$ be the stochastic process taking values in $\{0,\ldots,(\frac{1}{2}-2\varepsilon)n\} \cup \{n\}$ defined by

$$Z_t = \begin{cases} n - \max\{Y_t, (\frac{1}{2}+2\varepsilon)n\} & \text{if } t < T_Y^{\text{hit}}(V), \\ n & \text{if } t \geqslant T_Y^{\text{hit}}(V), \end{cases}$$

so that $(Z_t)_{t=0}^{\infty}$ is adapted to $(\mathscr{G}_t)_{t=0}^{\infty}$. Define

$$T^* := \min\{t : Z_t \leqslant (\tfrac{1}{2} - 2\varepsilon - 2\eta)n\}.$$

We now have the following claim.

Claim 3.11. $\mathbb{P}[T^* \leqslant e^{n^{1/8}}] \leqslant \frac{1}{2}e^{-n^{1/8}}$.

Sketch proof of Claim 3.11. The result follows from an application of Theorem 1.7 with $a = (\frac{1}{2} - 2\varepsilon - 2\eta)n$, $b = (\frac{1}{2} - 2\varepsilon - \eta)n$, $\ell = b - a = \eta n$, $\kappa = C(\varepsilon)/\varepsilon$, $r = n^{1/3}$, and

$$\Delta_t = h(Z_{t+1} - Z_t) = \min\{Z_{t+1} - Z_t, C(\varepsilon)\}.$$

In the application, Claim 3.9 is used to verify **B1**, and Claim 3.8 is used to verify **B2**. For a complete proof, see Appendix B.11. ⊡

Let us define the first departure time of $(Y_t)_{t\geqslant 0}$ from V as $T_Y^{\text{dep}}(V) = \min\{t : Y_{t-1} \in V, Y_t \notin V\}$, and note that for any $\tau \geqslant 0$,

$$\begin{aligned} \mathbb{P}[T_Y^{\text{dep}}(V) \leqslant \tau] &\leqslant \sum_{t=1}^{\tau} \mathbb{P}(Y_{t-1} \in V \wedge Y_t \notin V) \\ &\leqslant \tau \cdot \sup_{i \in V} \mathbb{P}(Y_t \notin V \mid Y_{t-1} = i) \\ &\leqslant \tau \cdot \sup_{i \in V} \mathbb{P}(Y_t \neq Y_{t-1} \mid Y_{t-1} = i) \qquad (13) \\ &\overset{(11)}{\leqslant} \tau \cdot e^{-\varepsilon n^{1/3}}. \qquad (14) \end{aligned}$$

If $Y_t \in W$ then $Y_t \geqslant (\frac{1}{2}+2\varepsilon+2\eta)n$, and so either $Z_t \leqslant (\frac{1}{2}-2\varepsilon-2\eta)n$ or $t \geqslant T_Y^{\text{hit}}(V)$. In particular, $T_Y^{\text{hit}}(W) \geqslant \min\{T^*, T_Y^{\text{hit}}(V)\}$. However, if $T_Y^{\text{hit}}(W) \geqslant T_Y^{\text{hit}}(V)$, then in fact we have $T_Y^{\text{hit}}(W) > T_Y^{\text{dep}}(V)$, as $W \cap V = \emptyset$. Therefore, we can deduce that

$$T_Y^{\text{hit}}(W) \geqslant \min\{T^*, T_Y^{\text{dep}}(V)\}. \tag{15}$$

By applying Claim 3.10 with $R = W$, $T_Y^{\text{hit}}(W)$ is stochastically dominated by $\min\{t : |x_t| \geqslant (\frac{1}{2} + 2\varepsilon + 2\eta)n\}$, which is in turn stochastically dominated by T_C. Combining these observations, we can conclude that

$$\begin{aligned} \mathbb{P}[\lceil \mu/2 \rceil \cdot T_C \leqslant e^{n^{1/8}}] &\leqslant \mathbb{P}[T_C \leqslant e^{n^{1/8}}] \leqslant \mathbb{P}[T_Y^{\text{hit}}(W) \leqslant e^{n^{1/8}}] \\ &\overset{(15)}{\leqslant} \mathbb{P}[T^* \leqslant e^{n^{1/8}}] + \mathbb{P}[T_Y^{\text{dep}}(V) \leqslant e^{n^{1/8}}] \\ &\overset{\text{Claim 3.11}}{\leqslant} \tfrac{1}{2}e^{-n^{1/8}} + \mathbb{P}[T_Y^{\text{dep}}(V) \leqslant e^{n^{1/8}}] \\ &\overset{(14)}{\leqslant} \tfrac{1}{2}e^{-n^{1/8}} + e^{n^{1/8}} \cdot e^{-\varepsilon n^{1/3}} \leqslant e^{-n^{1/8}}, \end{aligned}$$

thus proving the theorem. □

4 CONCLUDING REMARKS

We have shown that a coevolutionary instance of UMDA is able to efficiently discover the Nash equilibrium for a large class of symmetric zero-sum games on which single-individual algorithms cannot guarantee polynomial runtime. A difficult subclass of these games are those that are locally flat, in the sense that strategies that are similar to each other will only result in draws. The proof for single-individual algorithms relies closely on the fact that, regardless of the selection operator, a local search performed on a locally flat game resembles a random walk with steps generated by the mutation operator. On the other hand, UMDA overcomes this challenge by using estimation of distribution to generate a much richer diversity of search points for comparison.

Even if we require distinct strategies to never draw (either by restricting the definition of $\mathcal{G}_n(\alpha, m)$ or instead by determining the winner probabilistically in such cases), both of our results still follow with some straightforward modification of the proofs. This may be noteworthy as games where similar opponents exhibit random-like payoffs may be more realistic than a locally flat payoff landscape.

Two natural open questions arise from our results. The first asks how far the class $\mathcal{G}_n(\alpha, m)$ can be generalised. While Definition 1.5 is general enough to allow for a large range of payoff landscapes within and around individual skill levels, the overall geometry is still fundamentally underpinned by OneMax. It would be interesting to know how the behaviour of coevolutionary algorithms would be affected we used a different pseudo-Boolean function or a different representation of skill-based games.

We can also ask for a more precise description of exactly what level or type of population diversity would suffice to efficiently optimise instances of $\mathcal{G}_n(\alpha, m)$. In terms of their ability to generate diverse search points, EDAs and single-individual algorithms lie at two extreme ends of the spectrum, inviting the question of what behaviour would we find for population based algorithms which do not utilise estimation of distribution. As one possible direction, we conjecture that, in the absence of diversity promoting features such as crossover or archives, population based coevolutionary algorithms offer no substantial improvement over single-individual algorithms on our problem class.

ACKNOWLEDGMENTS

This research was supported by a Turing AI Fellowship (EPSRC grant ref EP/V025562/1).

REFERENCES

[1] Maxim Buzdalov and Carola Doerr. 2020. Optimal Mutation Rates for the $(1 + \lambda)$ EA on OneMax. In *Parallel Problem Solving from Nature – PPSN XVI*. Springer International Publishing, Cham, 574–587. https://doi.org/10.1007/978-3-030-58115-2_40

[2] Wojciech Marian Czarnecki, Gauthier Gidel, Brendan Tracey, Karl Tuyls, Shayegan Omidshafiei, David Balduzzi, and Max Jaderberg. 2020. Real world games look like spinning tops. In *Proceedings of the 34th International Conference on Neural Information Processing Systems* (Vancouver, BC, Canada) *(NIPS'20)*. Curran Associates Inc., Red Hook, NY, USA, Article 1463, 12 pages.

[3] Duc-Cuong Dang and Per Kristian Lehre. 2015. Simplified Runtime Analysis of Estimation of Distribution Algorithms. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation* (Madrid, Spain) *(GECCO '15)*. Association for Computing Machinery, New York, NY, USA, 513–518. https://doi.org/10.1145/2739480.2754814

[4] Benjamin Doerr. 2021. The Runtime of the Compact Genetic Algorithm on Jump Functions. *Algorithmica* 83, 10 (2021), 3059–3107. https://doi.org/10.1007/s00453-020-00780-w

[5] Benjamin Doerr, Carola Doerr, and Jing Yang. 2020. Optimal parameter choices via precise black-box analysis. *Theoretical Computer Science* 801 (2020), 1–34. https://doi.org/10.1016/j.tcs.2019.06.014

[6] Benjamin Doerr and Leslie Ann Goldberg. 2013. Adaptive drift analysis. *Algorithmica* 65 (2013), 1–27. https://doi.org/10.1007/s00453-011-9585-3

[7] Benjamin Doerr and Frank Neumann. 2020. *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization*. Springer. https://doi.org/10.1007/978-3-030-29414-4

[8] Stefan Droste. 2006. A rigorous analysis of the compact genetic algorithm for linear functions. *Natural Computing* 5 (2006), 257–283. https://doi.org/10.1007/s11047-006-9001-0

[9] Stefan Droste, Thomas Jansen, and Ingo Wegener. 2006. Upper and Lower Bounds for Randomized Search Heuristics in Black-Box Optimization. *Theory of Computing Systems* 39, 4 (2006), 525–544. https://doi.org/10.1007/s00224-004-1177-z

[10] Mario Alejandro Hevia Fajardo, Per Kristian Lehre, and Shishen Lin. 2023. Runtime Analysis of a Co-Evolutionary Algorithm: Overcoming Negative Drift in Maximin-Optimisation. In *Proceedings of the 17th ACM/SIGEVO Conference on Foundations of Genetic Algorithms* (Potsdam, Germany) *(FOGA '23)*. Association for Computing Machinery, New York, NY, USA, 73–83. https://doi.org/10.1145/3594805.3607132

[11] Mahdieh Shahrabi Farahani and Majid Sheikhmohammady. 2014. A review on symmetric games: theory, comparison and applications. *International Journal of Applied Operational Research - An Open Access Journal* 4 (2014), 91–106. https://api.semanticscholar.org/CorpusID:199476913

[12] Sevan Gregory Ficici. 2004. *Solution concepts in coevolutionary algorithms*. Ph. D. Dissertation. Brandeis University, USA.

[13] David Galvin. 2014. Three tutorial lectures on entropy and counting. arXiv:1406.7872

[14] Gene M. Grossman and Carl Shapiro. 1984. Informative Advertising with Differentiated Products. *The Review of Economic Studies* 51, 1 (1984), 63–81. https://doi.org/10.2307/2297705

[15] Andreas Hefti. 2017. Equilibria in symmetric games: Theory and applications. *Theoretical Economics* 12, 3 (2017), 979–1002. https://doi.org/10.3982/TE2151 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.3982/TE2151

[16] Mario Alejandro Hevia Fajardo and Per Kristian Lehre. 2023. How Fitness Aggregation Methods Affect the Performance of Competitive CoEAs on Bilinear Problems. In *Proceedings of the Genetic and Evolutionary Computation Conference* (Lisbon, Portugal) *(GECCO '23)*. Association for Computing Machinery, New York, NY, USA, 1593–1601. https://doi.org/10.1145/3583131.3590506

[17] Thomas Jansen and R. Paul Wiegand. 2004. The Cooperative Coevolutionary (1+1) EA. *Evolutionary Computation* 12, 4 (2004), 405–434. https://doi.org/10.1162/1063656043138905

[18] Per Kristian Lehre. 2022. Runtime analysis of competitive co-evolutionary algorithms for maximin optimisation of a bilinear function. In *Proceedings of the Genetic and Evolutionary Computation Conference* (Boston, Massachusetts) *(GECCO '22)*. Association for Computing Machinery, New York, NY, USA, 1408–1416. https://doi.org/10.1145/3512290.3528853

[19] Per Kristian Lehre and Carsten Witt. 2012. Black-Box Search by Unbiased Variation. *Algorithmica* 64, 4 (2012), 623–642. https://doi.org/10.1007/s00453-012-9616-8

[20] Johannes Lengler. 2020. Drift Analysis. In *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization*. Springer International Publishing, Cham, 89–131. https://doi.org/10.1007/978-3-030-29414-4_2

[21] Colin McDiarmid. 1989. On the method of bounded differences. In *Surveys in Combinatorics, 1989: Invited Papers at the Twelfth British Combinatorial Conference (London Mathematical Society Lecture Note Series)*. Cambridge University Press, 148–188. https://doi.org/10.1017/CBO9781107359949.008

[22] Heinz Mühlenbein and Gerhard Paaß. 1996. From recombination of genes to the estimation of distributions I. Binary parameters. In *Parallel Problem Solving from Nature (PPSN IV)*. Springer Berlin Heidelberg, Berlin, Heidelberg, 178–187. https://doi.org/10.1007/3-540-61723-X_982

[23] Y Narahari. 2014. *Game Theory and Mechanism Design*. World Scientific/Indian Institute of Science. https://doi.org/10.1142/8902

[24] Martin Pelikan, Mark Hauschild, and Fernando G. Lobo. 2015. Estimation of Distribution Algorithms. In *Springer Handbook of Computational Intelligence*. Springer, 899–928. https://doi.org/10.1007/978-3-662-43505-2_45

[25] Elena Popovici, Anthony Bucci, R. Paul Wiegand, and Edwin D. De Jong. 2012. *Coevolutionary Principles*. Springer Berlin Heidelberg, Berlin, Heidelberg, 987–1033. https://doi.org/10.1007/978-3-540-92910-9_31

[26] Christopher D. Rosin and Richard K. Belew. 1997. New methods for competitive coevolution. *Evolutionary Computation* 5, 1 (1997), 1–29. https://doi.org/10.1162/evco.1997.5.1.1

[27] Mohammad Karim Sohrabi and Hossein Azgomi. 2020. A Survey on the Combined Use of Optimization Methods and Game Theory. *Archives of Computational Methods in Engineering* 27, 1 (2020), 59–80. https://doi.org/10.1007/s11831-018-9300-5

[28] Bernhard Von Stengel. 2002. Chapter 45 Computing equilibria for two-person games. In *Handbook of Game Theory with Economic Applications*. Vol. 3. Elsevier, 1723–1759. https://doi.org/10.1016/S1574-0005(02)03008-4

[29] Carsten Witt. 2019. Upper bounds on the running time of the univariate marginal distribution algorithm on onemax. *Algorithmica* 81 (2019), 632–667. https://doi.org/10.1007/s00453-018-0463-0

[30] Carsten Witt. 2023. How majority-vote crossover and estimation-of-distribution algorithms cope with fitness valleys. *Theoretical Computer Science* 940 (2023), 18–42. https://doi.org/10.1016/j.tcs.2022.08.014

A CONCENTRATION INEQUALITIES

In several places throughout the appendex, we will need to show that certain random variables are unlikely to deviate significantly from their expectation. To assist with this, we will use McDiarmid's inequality [21].

Theorem A.1. *Suppose $f : \mathcal{X}_1 \times \ldots \times \mathcal{X}_n \to \mathbb{R}$ has the property that substituting the value of the i^{th} coordinate changes the value of f by at most c_i. Suppose that $X_1, \ldots, X_n$ are independent random variables where $X_i \in \mathcal{X}_i$ for each $i \in [n]$. Then, for any $t > 0$,*

$$\mathbb{P}(f(X_1, \ldots, X_n) \geqslant \mathbb{E}[f(X_1, \ldots, X_n)] + t) \leqslant \exp\left(-\frac{2t^2}{\sum_{i=1}^n c_i^2}\right),$$

$$\mathbb{P}(f(X_1, \ldots, X_n) \leqslant \mathbb{E}[f(X_1, \ldots, X_n)] - t) \leqslant \exp\left(-\frac{2t^2}{\sum_{i=1}^n c_i^2}\right).$$

In most cases, we will need to have concentration for sums of Bernoulli variables. The following corollary of Theorem A.1 will be convenient for such cases (although we remark that a Chernoff bound would also suffice).

Corollary A.2. *Let $X_1, \ldots, X_n$ be independent random variables taking values in $\{0, 1\}$, and let $X = \sum_{i \in [n]} X_i$. Then, for any $t > 0$,*

$$\mathbb{P}(X \geqslant \mathbb{E}[X] + t) \leqslant \exp(-2t^2/n),$$

$$\mathbb{P}(X \leqslant \mathbb{E}[X] - t) \leqslant \exp(-2t^2/n).$$

Proof. Apply Theorem A.1 with $f(x_1, \ldots, x_n) = \sum_{i \in [n]} x_i$ and $c_1 = \ldots = c_n = 1$. □

B DEFERRED PROOFS

B.1 Proof of Lemma 2.2

In the following proof of Lemma 2.2, we make use of the AM-GM inequality, which asserts that

$$\left(\prod_{i \in [n]} x_i\right)^{1/n} \leqslant \frac{1}{n} \sum_{i \in [n]} x_i \tag{16}$$

holds for any non-negative real numbers $x_1, \dots, x_n$.

Proof of Lemma 2.2. First, if $\sum_{i\in[n]} p_i \geqslant (1-\alpha+\delta/4)n$, then by applying Corollary A.2 to the random variable $|x|$ (and noting that $x^* = \mathbf{1}^n$),

$$\begin{aligned}\mathbb{P}(g(x,y)\neq f(x,y)) &\overset{\textbf{A1}}{\leqslant} \mathbb{P}(|x|\leqslant(1-\alpha)n)\\ &\leqslant \mathbb{P}(|x|\leqslant \mathbb{E}[|x|]-\delta n/4)\\ &\leqslant \exp(-2(\delta n/4)^2/n) = e^{-\delta^2 n/8}.\end{aligned}$$

Thus, we may assume instead that $\sum_{i\in[n]} p_i < (1-\alpha+\delta/4)n$. But then, for any $z \in \mathcal{X}_n$,

$$\begin{aligned}\mathbb{P}(y=z) &\leqslant \prod_{i\in[n]} \max\{p_i, 1-p_i\} \leqslant \prod_{i\in[n]} (p_i+\delta/2)\\ &\overset{(16)}{\leqslant} \left(\tfrac{1}{n}\textstyle\sum_{i\in[n]}(p_i+\delta/2)\right)^n \leqslant (1-\alpha+\tfrac{3\delta}{4})^n,\end{aligned}$$

where the second inequality follows because

$$1-p_i \leqslant 1-(\tfrac{1}{2}-\tfrac{\delta}{4}) = (\tfrac{1}{2}-\tfrac{\delta}{4})+\tfrac{\delta}{2} \leqslant p_i+\tfrac{\delta}{2}$$

for every $i\in[n]$. Therefore,

$$\mathbb{P}(g(x,y)\neq f(x,y)) \overset{\textbf{A2}}{\leqslant} \left(\frac{1-\alpha+\frac{3\delta}{4}}{1-\alpha+\delta}\right)^n \leqslant (1-\delta/8)^n \leqslant e^{-\delta n/8},$$

as required. □

B.2 Proof of Claim 2.3

We require the following simple property of sums of independent Bernoulli variables, which we prove before continuing to the proof of Claim 2.3.

Lemma B.1. *Let $X_1,\dots,X_n,Y_1,\dots,Y_n$ be independent random variables taking values in $\{0,1\}$ such that $\mathbb{P}(X_i=1)=\mathbb{P}(Y_i=1)$ for every $i\in[n]$. Let $X=\sum_{i\in[n]}X_i$ and $Y=\sum_{i\in[n]}Y_i$. Then, $\mathbb{P}(X=Y)\geqslant \frac{1}{4\sqrt{n}+6}$.*

Proof. For each $k\in\mathbb{N}$ let $q_k=\mathbb{P}(X=k)$. Let $m_1=\lfloor \mathbb{E}[X]-\sqrt{n}\rfloor$ and $m_2=\lceil \mathbb{E}[X]+\sqrt{n}\rceil$. By Corollary A.2

$$\mathbb{P}(|X-\mathbb{E}[X]|\geqslant\sqrt{n})\leqslant 2e^{-2}. \tag{17}$$

By Jensen's inequality (applied with the function $f(x)=x^2$),

$$\frac{\sum_{k=m_1}^{m_2} q_k^2}{m_2-m_1+1} \geqslant \left(\frac{\sum_{k=m_1}^{m_2} q_k}{m_2-m_1+1}\right)^2. \tag{18}$$

Therefore, because X and Y are independent and identically distributed,

$$\begin{aligned}\mathbb{P}(X=Y) &= \sum_{k\in\mathbb{N}}\mathbb{P}(X=k\wedge Y=k) = \sum_{k\in\mathbb{N}} q_k^2 \geqslant \sum_{k=m_1}^{m_2} q_k^2\\ &\overset{(18)}{\geqslant} \frac{(\sum_{k=m_1}^{m_2} q_k)^2}{m_2-m_1+1} \geqslant \frac{\mathbb{P}(|X-\mathbb{E}[X]|<\sqrt{n})^2}{2\sqrt{n}+3}\\ &\overset{(17)}{\geqslant} \frac{(1-2e^{-2})^2}{2\sqrt{n}+3} \geqslant \frac{1}{4\sqrt{n}+6},\end{aligned}$$

as required. □

For conciseness, we use $\mathbb{P}(\cdot)$ in place of $\mathbb{P}(\cdot \mid \mathscr{F}_t)$ in the following proof of Claim 2.3.

Proof of Claim 2.3. For an arbitrary $j\in[\mu]$, we will analyse the probability that $P_{t+1}(j)$ has a 1-bit in position i, where $P_{t+1}(j)$ is sampled according to the process described in lines 6-14 of Algorithm 2. Let x and y be the bitstrings sampled at the beginning of this procedure, so that x and y are sampled independently according to $\mathrm{Bit}_n(p_{t,1},\dots,p_{t,n})$. For $a,b\in\{0,1\}$, let A_{ab} be the event that x has an a-bit in position i and y has a b-bit in position i. By Lemma 2.2,

$$\mathbb{P}(A_{10}\wedge f(x,y)=1)\leqslant \mathbb{P}(A_{10}\wedge g(x,y)=1)+e^{-\delta^2 n/8}. \tag{19}$$

Furthermore, by the symmetry of the selection process and the fact that $g(x,y)=-g(y,x)$, we have for any $k\in\{-1,0,1\}$ that

$$\begin{aligned}\mathbb{P}(A_{10}\wedge g(x,y)=k) &= \mathbb{P}(A_{01}\wedge g(y,x)=k)\\ &= \mathbb{P}(A_{01}\wedge g(x,y)=-k).\end{aligned} \tag{20}$$

In addition,

$$\begin{aligned}\mathbb{P}(f(x,y)=1\mid A_{10}) &\geqslant \mathbb{P}(|x|>|y|\mid A_{10})\\ &= \mathbb{P}\left(\textstyle\sum_{j\neq i} x^{(j)} \geqslant \sum_{j\neq i} y^{(j)}\right)\\ &= \frac{1}{2}\left(1+\mathbb{P}\left(\textstyle\sum_{j\neq i} x^{(j)} = \sum_{j\neq i} y^{(j)}\right)\right)\\ &\overset{\text{Lemma B.1}}{\geqslant} \frac{1}{2}\left(1+\frac{1}{4\sqrt{n-1}+6}\right)\\ &\overset{(1)}{\geqslant} \frac{1}{2}\left(1+\frac{1}{5\sqrt{n}}\right).\end{aligned} \tag{21}$$

If A_{11} occurs, then $P_{t+1}(j)$ will have a 1-bit in position i with probability 1. If A_{10} occurs, then $P_{t+1}(j)$ will have a 1-bit in position i with probability 1 if $g(x,y)=1$, probability $\frac{1}{2}$ if $g(x,y)=0$, and probability 0 if $g(x,y)=-1$. Similarly, if A_{01} occurs, then $P_{t+1}(j)$ will have a 1-bit in position i with probability 1 if $g(x,y)=-1$, probability $\frac{1}{2}$ if $g(x,y)=0$, and probability 0 if $g(x,y)=1$. Therefore,

$$\begin{aligned}q_{t,i} &= \mathbb{P}(A_{11})+\mathbb{P}(A_{10}\wedge g(x,y)=1)+\mathbb{P}(A_{01}\wedge g(x,y)=-1)\\ &\quad+\frac{1}{2}\mathbb{P}(A_{10}\wedge g(x,y)=0)+\frac{1}{2}\mathbb{P}(A_{01}\wedge g(x,y)=0)\\ &\overset{(20)}{=} \mathbb{P}(A_{11})+2\cdot\mathbb{P}(A_{10}\wedge g(x,y)=1)+\mathbb{P}(A_{10}\wedge g(x,y)=0)\\ &\geqslant \mathbb{P}(A_{11})+2\cdot\mathbb{P}(A_{10}\wedge g(x,y)=1)\\ &\overset{(19)}{\geqslant} \mathbb{P}(A_{11})+2\cdot\mathbb{P}(A_{10}\wedge f(x,y)=1)-2e^{-\delta^2 n/8}\\ &= \mathbb{P}(A_{11})+2\cdot\mathbb{P}(A_{10})\cdot\mathbb{P}(f(x,y)=1\mid A_{10})-2e^{-\delta^2 n/8}\\ &\overset{(21)}{\geqslant} p_{t,i}^2+p_{t,i}(1-p_{t,i})(1+\tfrac{1}{5\sqrt{n}})-2e^{-\delta^2 n/8}\\ &= p_{t,i}\left(1+\frac{1-p_{t,i}}{5\sqrt{n}}\right)-2e^{-\delta^2 n/8},\end{aligned}$$

as required. □

B.3 Proof of Claim 2.4

We require the following straightforward lemma, which we prove before continuing to the proof of Claim 2.4

Lemma B.2. *Let $a,b\in\mathbb{R}$ and $p_1,\dots,p_n\in[a,b]$. Then*

$$\sum_{i\in[n]} p_i^2 \leqslant (b+a)\sum_{i\in[n]} p_i - abn.$$

Proof. Let $S = \sum_{i \in [n]} p_i$. Then,

$$\begin{aligned}\sum_{i\in[n]} p_i^2 &= \sum_{i\in[n]} (a + (p_i - a))^2 \\ &= \sum_{i\in[n]} (a^2 + 2a(p_i - a) + (p_i - a)^2) \\ &= 2aS - a^2 n + \sum_{i\in[n]} (p_i - a)^2 \\ &\leqslant 2aS - a^2 n + (S - an)(b - a) \\ &= (b + a)S - abn,\end{aligned}$$

as required. □

Proof of Claim 2.4. Recall that $(X_t)_{t=0}^\infty$ is a stochastic process adapted to $(\mathscr{F}_t)_{t=0}^\infty$ taking values in a finite subset of $\{0\} \cup [1, \infty)$. If $X_t > 0$, then $(p_{t,1}, \dots, p_{t,n}) \in [\frac{1}{2} - \frac{\delta}{4}, 1]^n$, $\sum_{i\in[n]} p_{t,i} \leqslant n - 1$, and

$$X_t \overset{(4)}{=} n - \sum_{i\in[n]} p_{t,i}. \tag{22}$$

It also holds that

$$X_{t+1} \overset{(4)}{\leqslant} n - \sum_{i\in[n]} p_{t+1,i}. \tag{23}$$

Because $\delta < 1/2$ is an assumption of Theorem 2.1, we can apply Lemma B.2 with $a = \frac{3}{8}$, $b = 1$, and $S = n - X_t$ to obtain

$$\sum_{i\in[n]} p_{t,i}^2 \leqslant \tfrac{11}{8}(n - X_t) - \tfrac{3}{8}n = n - X_t - \tfrac{3}{8}X_t. \tag{24}$$

Therefore,

$$\begin{aligned}\mathbb{E}[X_t - X_{t+1} \mid \mathscr{F}_t] &\overset{(22)}{=} n - \sum_{i\in[n]} p_{t,i} - \mathbb{E}[X_{t+1} \mid \mathscr{F}_t] \\ &\overset{(23)}{\geqslant} n - \sum_{i\in[n]} p_{t,i} - n + \sum_{i\in[n]} \mathbb{E}[p_{t+1,i} \mid \mathscr{F}_t] \\ &\overset{(3)}{\geqslant} \sum_{i\in[n]} (q_{t,i} - p_{t,i}) \\ &\overset{\text{Claim 2.3}}{\geqslant} \sum_{i\in[n]} \left(\frac{p_{t,i}(1 - p_{t,i})}{5\sqrt{n}} - 2e^{-\delta^2 n/8}\right) \\ &= \frac{1}{5\sqrt{n}}\left(n - X_t - \textstyle\sum_{i\in[n]} p_{t,i}^2\right) - 2ne^{-\delta^2 n/8} \\ &\overset{(24)}{\geqslant} \frac{3}{40\sqrt{n}} X_t - 2ne^{-\delta^2 n/8} \geqslant \frac{1}{20\sqrt{n}} X_t,\end{aligned}$$

where in the final inequality, we have used that

$$2ne^{-\delta^2 n/8} \leqslant 1/(40\sqrt{n}) \leqslant X_t/(40\sqrt{n}).$$

Therefore, by applying Theorem 1.6 with $\delta = 1/(20\sqrt{n})$, $x_{\min} = 1$, $r = 2K \ln n$, and noting that $X_0 = n/2$, we can compute that for any $K \geqslant 1$,

$$\begin{aligned}\mathbb{P}[T_0 \geqslant 50K\sqrt{n}\ln n] &\leqslant \mathbb{P}[T_0 > \lceil 20\sqrt{n}(\ln(n/2) + 2K\ln n)\rceil] \\ &\leqslant e^{-2K\ln n} = n^{-2K} \leqslant \tfrac{1}{2} n^{-K},\end{aligned}$$

as required. ⊡

B.4 Proof of Claim 2.5

At first glance, a proof strategy for deriving a lower tail bound on the random variable

$$T_{\text{bad}} := \min\{t : (p_{t,1}, \dots, p_{t,n}) \notin [\tfrac{1}{2} - \tfrac{\delta}{4}, 1]^n\}$$

might be to apply a negative drift theorem to each of the bit frequencies $(p_{t,i})_{t=0}^\infty$ and then take a union bound. Indeed, it seems reasonable that Claim 2.3 (together with (3)) could be used to verify conditions **B1**-**B3** in Theorem 1.7 (just as it was used to verify the conditions of Theorem 1.6 when proving an upper tail bound on T_0). However, two main complications arise with this approach.

(a) Claim 2.3 only applies if $(p_{t,1}, \dots, p_{t,n}) \in [\frac{1}{2} - \frac{\delta}{4}, 1]^n$, and so obtaining a lower bound on $\mathbb{E}[p_{t+1,i}\, ;\, a < p_{t,i} < b \mid \mathscr{F}_t]$ in isolation is impossible without also accounting for the behaviour of the other bit frequencies.
(b) If the actual value of $q_{t,i}$ is significantly larger than the lower bound provided by Claim 2.3, then there may not be a suitable choice of $\Delta_t := \Delta_t(p_{t+1,i} - p_{t,i})$ and κ in Theorem 1.7 for **B1** to hold.

Handling (a) is relatively straightforward – instead of examining $(p_{t,i})_{t=0}^\infty$, we consider the process $(Y_t)_{t=0}^\infty$ defined by

$$Y_{t,i} = \begin{cases}\mu \cdot \max\{p_{t,i}, (\frac{1}{2} - \frac{\delta}{4})\} & \text{if } t \leqslant T_{\text{bad}}, \\ \mu & \text{otherwise.}\end{cases}$$

so that Claim 2.3 always applies when $(\frac{1}{2} - \frac{\delta}{4})\mu < Y_{t,i} < \mu$ (note here that scaling by μ is purely for ease of notation). Handling (b) is considerably more subtle – for this we will introduce a new Markov chain $(Z_t)_{t=0}^\infty$ much more amenable to the application of the negative drift theorem, and use $(Z_t)_{t=0}^\infty$ to 'bound $(Y_{t,i})_{t=0}^\infty$ from below', in a very precise sense based on a coupling which we introduce in the next section.

B.4.1 A coupling result. In order to formally establish the relationship between $(Y_{t,i})_{t=0}^\infty$ and $(Z_t)_{t=0}^\infty$, we must first extend our notation of stochastic domination slightly. Recall from Section 1.1 that X is said to stochastically dominate Y, written $X \succcurlyeq Y$, if $\mathbb{P}(X \leqslant z) \leqslant \mathbb{P}(Y \leqslant z)$ holds for all $z \in \mathbb{R}$. Because stochastic domination depends only on the distribution functions on $\mathbb{R}$ induced by X and Y, this relation can just as easily apply to real-valued random variables which are not defined on the same probability spaces, as follows.

Definition B.3. *Suppose $(\Omega_1, \mathscr{F}_1, \mathbb{P}_1)$ and $(\Omega_2, \mathscr{F}_2, \mathbb{P}_2)$ are probability spaces, and that $X_1 : \Omega_1 \to \mathbb{R}$ and $X_2 : \Omega_2 \to \mathbb{R}$ are real-valued random variables. Then X_1 stochastically dominates X_2, written $X_1 \succcurlyeq X_2$, if $\mathbb{P}_1(X_1 \leqslant x) \leqslant \mathbb{P}_2(X_2 \leqslant x)$ holds for all $x \in \mathbb{R}$.*

One notable instance of Definition B.3 will be when the random variables are subject to conditioning, as follows.

Definition B.4. *Given a random variable X defined on a probability space $(\Omega, \mathscr{F}, \mathbb{P})$ and an event $E \in \mathscr{F}$, we write $(X \mid E)$ to denote an instance of X defined on the probability space $(\Omega, \mathscr{F}, \mathbb{P}(\,\cdot \mid E))$.*

Accordingly, if X_1 and X_2 are real-valued random variables defined on probability spaces $(\Omega_1, \mathscr{F}_1, \mathbb{P}_1)$ and $(\Omega_2, \mathscr{F}_2, \mathbb{P}_2)$, and $E_1 \in \mathscr{F}_1$ and $E_2 \in \mathscr{F}_2$ are events, then we write $(X_1 \mid E_1) \succcurlyeq (X_2 \mid E_2)$ if $\mathbb{P}_1(X_1 \leqslant x \mid E_1) \leqslant \mathbb{P}_2(X_2 \leqslant x \mid E_2)$ holds for all $x \in \mathbb{R}$.

We are now ready to introduce couplings of random variables, which is the main tool for this section. In the following, $X \overset{d}{=} Y$ is used to denote that X and Y are *equal in distribution.*

Definition B.5. *Suppose X_1 and X_2 are random variables defined on probability spaces $(\Omega_1, \mathscr{F}_1, \mathbb{P}_1)$ and $(\Omega_2, \mathscr{F}_2, \mathbb{P}_2)$. A* coupling *of X_1 and X_2 is a pair $(\tilde{X}_1, \tilde{X}_2)$ of random variables defined on a new probability space $(\Omega, \mathscr{F}, \mathbb{P})$ such that $\tilde{X}_1 \overset{d}{=} X_1$ and $\tilde{X}_2 \overset{d}{=} X_2$.*

If there is a coupling $(\tilde{X}_1, \tilde{X}_2)$ of X_1 and X_2 such that $\mathbb{P}(\tilde{X}_1 \geqslant \tilde{X}_2) = 1$, then we can observe that for any $x \in \mathbb{R}$,

$$\mathbb{P}_1(X_1 \leqslant x) = \mathbb{P}(\tilde{X}_1 \leqslant x) \leqslant \mathbb{P}(\tilde{X}_2 \leqslant x) = \mathbb{P}_2(X_2 \leqslant x),$$

and so $X_1 \succcurlyeq X_2$. In fact it is well-known that the converse also holds, and so we have the following (see [7, Theorem 1.8.10]).

Theorem B.6. *For real-valued random variables X_1 and X_2, X_1 stochastically dominates X_2 if and only if there is a coupling $(\tilde{X}_1, \tilde{X}_2)$ of X_1 and X_2 such that $\mathbb{P}(\tilde{X}_1 \geqslant \tilde{X}_2) = 1$.*

In such cases, $(\tilde{X}_1, \tilde{X}_2)$ is called a *monotone coupling* of X_1 and X_2. In the following lemma, we extend the existence of a monotone coupling to stochastic processes. The subsequent corollary then relates this extension to applications which involve hitting times. (Note that going forward, if the underlying probability spaces are clear from context, we will omit the subscripts from $\mathbb{P}_1$ and $\mathbb{P}_2$, as has been the convention prior to this section.)

Lemma B.7. *Suppose $(Y_t)_{t=0}^{\infty}$ and $(Z_t)_{t=0}^{\infty}$ are stochastic processes taking values in a finite set $S \subseteq \mathbb{R}$ which satisfy the following property.*

D *For every $t \geqslant 0$, if $y_0, \ldots, y_{t-1}, z_0 \ldots, z_{t-1} \in S$ are such that $y_s \geqslant z_s$ for every $0 \leqslant s < t$, then*

$$(Y_t \mid Y_0 = y_0, \ldots, Y_{t-1} = y_{t-1}) \succcurlyeq (Z_t \mid Z_0 = z_0, \ldots, Z_{t-1} = z_{t-1}).$$

Then there is a coupling $((\tilde{Y}_t)_{t=0}^{\infty}, (\tilde{Z}_t)_{t=0}^{\infty})$ of $(Y_t)_{t=0}^{\infty}$ and $(Z_t)_{t=0}^{\infty}$ such that $\mathbb{P}(\wedge_{t\geqslant 0}(\tilde{Y}_t \geqslant \tilde{Z}_t)) = 1$.

Proof. For every $t \geqslant 0$ and $y_0, \ldots, y_{t-1}, z_0, \ldots, z_{t-1} \in S$, let $(\hat{Y}_t(y_0, \ldots, y_{t-1}), \hat{Z}_t(z_0, \ldots, z_{t-1}))$ be a coupling of $(Y_t \mid Y_0 = y_0, \ldots, Y_{t-1} = y_{t-1})$ and $(Z_t \mid Z_0 = z_0, \ldots, Z_{t-1} = z_{t-1})$, chosen such that

(i) by **D**, the coupling is monotone whenever $y_s \geqslant z_s$ for every $0 \leqslant s < t$, and
(ii) each coupling is independent of every other coupling.

(We remark that this family of couplings includes a monotone coupling $(\hat{Y}_0(), \hat{Z}_0())$ of Y_0 and Z_0.) Now, for each $t \geqslant 0$, define inductively

$$\begin{aligned} \tilde{Y}_t &= \hat{Y}_t(\tilde{Y}_0, \ldots, \tilde{Y}_{t-1}), \\ \tilde{Z}_t &= \hat{Z}_t(\tilde{Z}_0, \ldots, \tilde{Z}_{t-1}). \end{aligned} \tag{25}$$

First, let us verify that $(\tilde{Y}_t)_{t=0}^{\infty}$ and $(Y_t)_{t=0}^{\infty}$ are equal in distribution. Because, by (ii), $\{\hat{Y}_t(y_0, \ldots, y_{t-1}) : t \geqslant 0 \text{ and } y_0, \ldots, y_{t-1} \in S\}$ is an independent family of random variables,

$$\begin{aligned} \mathbb{P}(\tilde{Y}_0 = y_0 \wedge \ldots \wedge \tilde{Y}_t = y_t) &= \prod_{s=0}^{t} \mathbb{P}(\tilde{Y}_s = y_s \mid \wedge_{r<s}(\tilde{Y}_r = y_r)) \\ &= \prod_{s=0}^{t} \mathbb{P}(\hat{Y}_s(y_0, \ldots, y_{s-1}) = y_s). \end{aligned} \tag{26}$$

Additionally, because $\hat{Y}_s(y_0, \ldots, y_{s-1})$ and $(Y_s \mid \wedge_{r<s}(Y_r = y_r))$ are equal in distribution, we always have

$$\mathbb{P}(\hat{Y}_s(y_0, \ldots, y_{s-1}) = y_s) = \mathbb{P}(Y_s = y_s \mid \wedge_{r<s}(Y_r = y_r)). \tag{27}$$

Thus, it holds for every $t \geqslant 0$ and $y_0, \ldots, y_{t-1} \in S$ that

$$\begin{aligned} \mathbb{P}(\tilde{Y}_0 = y_0 \wedge \ldots \wedge \tilde{Y}_t = y_t) &\overset{(26),(27)}{=} \prod_{s=0}^{t} \mathbb{P}(Y_s = y_s \mid \wedge_{r<s}(Y_r = y_r)) \\ &= \mathbb{P}(Y_0 = y_0 \wedge \ldots \wedge Y_t = y_t), \end{aligned}$$

and hence $(\tilde{Y}_t)_{t=0}^{\infty}$ and $(Y_t)_{t=0}^{\infty}$ are equal in distribution. Similarly, $(\tilde{Z}_t)_{t=0}^{\infty}$ and $(Z_t)_{t=0}^{\infty}$ are equal in distribution, and so $((\tilde{Y}_t)_{t=0}^{\infty}, (\tilde{Z}_t)_{t=0}^{\infty})$ is indeed a coupling of $(Y_t)_{t=0}^{\infty}$ and $(Z_t)_{t=0}^{\infty}$.

Finally, let us verify that $\mathbb{P}(\wedge_{t\geqslant 0}(\tilde{Y}_t \geqslant \tilde{Z}_t)) = 1$. To do this, we first claim by induction that

$$\mathbb{P}(\wedge_{t\leqslant \tau}(\tilde{Y}_t \geqslant \tilde{Z}_t)) = 1 \tag{28}$$

holds for every $\tau \geqslant 0$. Indeed, the case $\tau = 0$ holds because, by (i), $(\tilde{Y}_0, \tilde{Z}_0) := (\hat{Y}_0(), \hat{Z}_0())$ is a monotone coupling of Y_0 and Z_0. For the case $\tau > 0$, first set $Q = \{(y, z) \in S^2 : y \geqslant z\}$ and note that for any $(y_0, z_0), \ldots, (y_{\tau-1}, z_{\tau-1}) \in Q$,

$$\begin{aligned} &\mathbb{P}(\tilde{Y}_\tau \geqslant \tilde{Z}_\tau \mid \wedge_{t<\tau}((\tilde{Y}_t, \tilde{Z}_t) = (y_t, z_t))) \\ &\overset{(25)}{=} \mathbb{P}(\hat{Y}_\tau(y_0, \ldots, y_{\tau-1}) \geqslant \hat{Z}_\tau(z_0, \ldots, z_{\tau-1}) \mid \wedge_{t<\tau}((\tilde{Y}_t, \tilde{Z}_t) = (y_t, z_t))) \\ &\overset{(ii)}{=} \mathbb{P}(\hat{Y}_\tau(y_0, \ldots, y_{\tau-1}) \geqslant \hat{Z}_\tau(z_0, \ldots, z_{\tau-1})) \overset{(i)}{=} 1. \end{aligned} \tag{29}$$

Therefore, applying the inductive hypothesis,

$$\begin{aligned} &\mathbb{P}(\wedge_{t\leqslant \tau}(\tilde{Y}_t \geqslant \tilde{Z}_t)) \\ &\overset{(29)}{=} \sum_{(y_0,z_0),\ldots,(y_{\tau-1},z_{\tau-1})\in Q} \mathbb{P}(\wedge_{t<\tau}((\tilde{Y}_t, \tilde{Z}_t) = (y_t, z_t))) \\ &= \mathbb{P}(\wedge_{t\leqslant \tau-1}(\tilde{Y}_t \geqslant \tilde{Z}_t)) = 1, \end{aligned}$$

and so (28) holds for all $\tau \geqslant 0$. Therefore, we have

$$\mathbb{P}(\wedge_{t\geqslant 0}(\tilde{Y}_t \geqslant \tilde{Z}_t)) = \lim_{\tau\to\infty} \mathbb{P}(\wedge_{t\leqslant \tau}(\tilde{Y}_t \geqslant \tilde{Z}_t)) \overset{(28)}{=} \lim_{\tau\to\infty} 1 = 1,$$

as required. □

Corollary B.8. *Suppose $(Y_t)_{t=0}^{\infty}$ and $(Z_t)_{t=0}^{\infty}$ are stochastic processes taking values in a finite set $S \subseteq \mathbb{R}$ which satisfy property **D**. Then, for any $a \in \mathbb{R}$, the first hitting times $T_Y := \min\{t : Y_t \leqslant a\}$ and $T_Z := \min\{t : Z_t \leqslant a\}$ satisfy $T_Y \succcurlyeq T_Z$.*

Proof. Using Lemma B.7, let $((\tilde{Y}_t)_{t=0}^{\infty}, (\tilde{Z}_t)_{t=0}^{\infty})$ be a coupling of $(Y_t)_{t=0}^{\infty}$ and $(Z_t)_{t=0}^{\infty}$ such that $\mathbb{P}(\wedge_{t\geqslant 0}(\tilde{Y}_t \geqslant \tilde{Z}_t)) = 1$. We then have, for any $\tau \geqslant 0$, that

$$\begin{aligned} \mathbb{P}[T_Y \leqslant \tau] &= \mathbb{P}(Y_0 \leqslant a \vee \ldots \vee Y_\tau \leqslant a) \\ &= \mathbb{P}(\tilde{Y}_0 \leqslant a \vee \ldots \vee \tilde{Y}_\tau \leqslant a) \\ &\leqslant \mathbb{P}(\tilde{Z}_0 \leqslant a \vee \ldots \vee \tilde{Z}_\tau \leqslant a) \\ &= \mathbb{P}(Z_0 \leqslant a \vee \ldots \vee Z_\tau \leqslant a) \\ &= \mathbb{P}[T_Z \leqslant \tau], \end{aligned}$$

as required. □

B.4.2 Drift analysis. In the following lemma, we introduce the Markov chain $(Z_t)_{t=0}^\infty$ and prove a bound on the relevant hitting time using drift analysis.

Lemma B.9. *Given constants $0 < \delta < 1/2$ and $K \geqslant 1$, the following holds for any sufficiently large n and μ satisfying (2). Let $(Z_t)_{t=0}^\infty$ be the Markov chain defined by*

$$Z_0 = \mu/2,$$
$$Z_{t+1} \sim Bin(\mu, q(Z_t/\mu)),$$

where

$$q(p) = p\left(1 + \frac{1-p}{5\sqrt{n}}\right) - 2e^{-\delta^2 n/8}.$$

Then it holds for the first hitting time $T = \min\{t : Z_t \leqslant (\frac{1}{2} - \frac{\delta}{4})\mu\}$ that

$$\mathbb{P}[T_Z \leqslant 50K\sqrt{n}\log n] \leqslant \tfrac{1}{2}n^{-2K}.$$

Proof. Set $a = (\frac{1}{2} - \frac{\delta}{4})\mu$, $b = \frac{1}{2}\mu$, $\ell = b - a = \frac{\delta}{4}\mu$, $r = \sqrt{2\mu}$, $\varepsilon = \mu/(50\sqrt{n})$, and $\kappa = \delta\mu/(40K\varepsilon\ln n)$. Let $(\mathscr{G}_t)_{t=0}^\infty$ denote the filtration generated by $(Z_t)_{t=0}^\infty$. We will verify the conditions **B1**-**B3** for $(Z_t)_{t=0}^\infty$ with $\Delta_t = Z_{t+1} - Z_t$. To assist with this, first note that because $q(p) - p$ is a quadratic function of p attaining its maximum at $p = \frac{1}{2} = \frac{b}{\mu}$, it holds whenever $a < Z_t$ that

$$\mathbb{E}[\Delta_t \mid \mathscr{G}_t] \leqslant \sup_{p \in (\frac{a}{\mu}, 1]} \mu(q(p) - p) \leqslant \mu/(20\sqrt{n}), \tag{30}$$

$$\mathbb{E}[\Delta_t \mid \mathscr{G}_t] \geqslant \inf_{p \in (\frac{a}{\mu}, 1]} \mu(q(p) - p) \geqslant -2\mu e^{-\delta^2 n/8}. \tag{31}$$

Next, because $\delta \leqslant 1/2$,

$$\begin{aligned}\mathbb{E}[\Delta_t - \tfrac{\mu}{25\sqrt{n}}\,;\, a < Z_t < b \mid \mathscr{G}_t] &\geqslant \inf_{p \in (\frac{a}{\mu}, \frac{b}{\mu})} \mu(q(p) - p) - \tfrac{\mu}{25\sqrt{n}} \\ &= \mu(q(a/\mu) - \tfrac{a}{\mu}) - \tfrac{\mu}{25\sqrt{n}} \\ &= \mu\left(\frac{\frac{1}{4} - \frac{\delta^2}{16}}{5\sqrt{n}} - 2e^{-\delta^2 n/8}\right) - \frac{\mu}{25\sqrt{n}} \geqslant 0 \end{aligned} \tag{32}$$

In addition, by applying Corollary A.2, we have for any $t > 0$ that

$$\mathbb{P}(\Delta_t \geqslant \mathbb{E}[\Delta_t] + t) \leqslant \exp\left(-\tfrac{2t^2}{\mu}\right), \tag{33}$$

$$\mathbb{P}(\Delta_t \leqslant \mathbb{E}[\Delta_t] - t) \leqslant \exp\left(-\tfrac{2t^2}{\mu}\right). \tag{34}$$

To see that **B1** holds, first observe that if $a < Z_t$ then, because $K \geqslant 1$ and n is assumed to be sufficiently large,

$$\begin{aligned}\mathbb{P}(\Delta_t > \kappa\varepsilon \mid \mathscr{G}_t) &= \mathbb{P}\left(\Delta_t > \tfrac{\delta\mu}{40K\ln n} \mid \mathscr{G}_t\right) \\ &= \mathbb{P}\left(\Delta_t > \tfrac{\delta\mu}{80K\ln n} + \tfrac{\delta\mu}{80K\ln n} \mid \mathscr{G}_t\right) \\ &\leqslant \mathbb{P}\left(\Delta_t \geqslant \tfrac{\mu}{20\sqrt{n}} + \tfrac{\delta\mu}{80K\ln n} \mid \mathscr{G}_t\right) \\ &\overset{(30)}{\leqslant} \mathbb{P}\left(\Delta_t \geqslant \mathbb{E}[\Delta_t] + \tfrac{\delta\mu}{80K\ln n} \mid \mathscr{G}_t\right) \\ &\overset{(33)}{\leqslant} \exp\left(-\tfrac{2\delta^2\mu}{(80K\ln n)^2}\right) \overset{(2)}{\leqslant} \exp\left(-\tfrac{2C\delta\sqrt{n}}{80^2K\ln n}\right) \\ &\leqslant 1/(50\sqrt{n}).\end{aligned}$$

In particular,

$$\mathbb{E}[\tfrac{\mu}{50\sqrt{n}} - \mu \cdot \mathbb{1}(\Delta_t > \kappa\varepsilon)\,;\, a < Z_t < b \mid \mathscr{G}_t] \geqslant 0. \tag{35}$$

Note also that, because $\Delta_t \leqslant \mu$,

$$\begin{aligned}\Delta_t \cdot \mathbb{1}(\Delta_t \leqslant \kappa\varepsilon) - \varepsilon &= \Delta_t \cdot \mathbb{1}(\Delta_t \leqslant \kappa\varepsilon) - \tfrac{\mu}{25\sqrt{n}} + \tfrac{\mu}{50\sqrt{n}} \\ &\geqslant \Delta_t - \tfrac{\mu}{25\sqrt{n}} + \tfrac{\mu}{50\sqrt{n}} - \mu \cdot \mathbb{1}(\Delta_t > \kappa\varepsilon)\end{aligned} \tag{36}$$

Therefore,

$$\begin{aligned}&\mathbb{E}[\Delta_t \cdot \mathbb{1}(\Delta_t \leqslant \kappa\varepsilon)\,;\, a < Z_t < b \mid \mathscr{G}_t] \\ &\overset{(36)}{\geqslant} \mathbb{E}[\Delta_t - \tfrac{\mu}{25\sqrt{n}} + \tfrac{\mu}{50\sqrt{n}} - \mu \cdot \mathbb{1}(\Delta_t > \kappa\varepsilon)\,;\, a < Z_t < b \mid \mathscr{G}_t] \\ &\overset{(32),(35)}{\geqslant} 0.\end{aligned}$$

To see that **B2** holds, first note that

$$\frac{2\mu}{r} = \sqrt{2\mu} \overset{(2)}{\leqslant} 2e^{\delta^2 n/8},$$

and hence $2\mu e^{-\delta^2 n/8} \leqslant jr/2$ holds for any $j \geqslant 1$. Therefore, if $a < Z_t$ and $j \in \mathbb{N}$,

$$\begin{aligned}\mathbb{P}(\Delta_t \leqslant -jr \mid \mathscr{G}_t) &\overset{(31)}{\leqslant} \mathbb{P}(\Delta_t \leqslant \mathbb{E}[\Delta_t] - jr/2 \mid \mathscr{G}_t) \\ &\overset{(34)}{\leqslant} \exp\left(-j^2r^2/2\mu\right) \leqslant e^{-j}.\end{aligned}$$

Finally, to verify **B3**, recall that $C = 10^5 \geqslant 1700 \cdot 40$ and set

$$\begin{aligned}\lambda :&= \min\{1/(2r), \varepsilon/(17r^2), 1/(\kappa\varepsilon)\} \\ &= \min\{1/(2\sqrt{2\mu}), 1/(1700\sqrt{n}), 40K\ln n/(\delta\mu)\} \overset{(2)}{=} \frac{40K\ln n}{\delta\mu}.\end{aligned}$$

Observe now that

$$\begin{aligned}\lambda\ell &= 10K\ln n \\ &\geqslant 2\ln\left(\frac{5\delta\sqrt{n}}{K\ln n}\right) = 2\ln\left(\frac{4}{\lambda\varepsilon}\right).\end{aligned} \tag{37}$$

Therefore, by Theorem 1.7 we have

$$\begin{aligned}\mathbb{P}[T_Z \leqslant 50K\sqrt{n}\log n] &\leqslant \mathbb{P}[T_Z \leqslant n^{5K/2}] \overset{(37)}{=} \mathbb{P}[T_Z \leqslant e^{\lambda\ell/4}] \\ &\leqslant \overline{C} \cdot e^{-\lambda\ell/4} = \overline{C} \cdot n^{-5K/2} \overset{(1)}{\leqslant} \tfrac{1}{2}n^{-2K},\end{aligned} \tag{38}$$

as required. □

B.4.3 Proof of Claim 2.5.

Proof of Claim 2.5. For each $i \in [n]$ and $t \geqslant 0$, let

$$Y_{t,i} = \begin{cases}\mu \cdot \max\{p_{t,i}, (\frac{1}{2} - \frac{\delta}{4})\} & \text{if } t \leqslant T_{\text{bad}}, \\ \mu & \text{otherwise.}\end{cases}$$

so that, for each $i \in [n]$, $(Y_{t,i})_{t=0}^\infty$ is a stochastic process adapted to $(\mathscr{F}_t)_{t=0}^\infty$. Note that

$$\mu \cdot p_{t,i} \leqslant Y_{t,i} \tag{39}$$

always holds. For each $i \in [n]$, define

$$T^i_{\text{bad}} = \min\{t : Y_{t,i} \leqslant \mu(\tfrac{1}{2} - \tfrac{\delta}{4})\},$$

so that, recalling $T_{\text{bad}} = \min\{t : (p_{t,1}, \ldots, p_{t,n}) \notin [\frac{1}{2} - \frac{\delta}{4}, 1]^n\}$,

$$T_{\text{bad}} \geqslant \min_{i \in [n]} T^i_{\text{bad}}. \tag{40}$$

Let $(Z_t)_{t=0}^{\infty}$ be the Markov chain defined in Lemma B.9. We will show that, for any $i \in [n]$, $t \geqslant 0$ and $y_0, \ldots, y_{t-1}, z_0, \ldots, z_{t-1} \in [0, \mu]$ such that $z_s \leqslant y_s$ for every $0 \leqslant s < t$, we have

$$(Z_t \mid Z_0 = z_0, \ldots, Z_{t-1} = z_{t-1}) \preccurlyeq (Y_{t,i} \mid Y_{0,i} = y_0, \ldots, Y_{t-1,i} = y_{t-1}). \tag{41}$$

To see this, first observe that

$$\begin{aligned}(Z_t \mid Z_0 = z_0, \ldots, Z_{t-1} = z_{t-1}) &= (Z_t \mid Z_{t-1} = z_{t-1})\\ &= \mathrm{Bin}(\mu, q(z_{t-1}/\mu)) \preccurlyeq \mathrm{Bin}(\mu, q(y_{t-1}/\mu)).\end{aligned}$$

From here, there are then two cases. First, if $Y_{t-1,i} = y_{t-1} < \mu$ and $t-1 < T_{\mathrm{bad}}$, then $(p_{t-1,1}, \ldots, p_{t-1,n}) \in [\frac{1}{2} - \frac{\delta}{4}, 1]^n$. In this case, $p_{t-1,i} = y_{t-1}/\mu$, and hence

$$\begin{aligned}\mathrm{Bin}(\mu, q(y_{t-1}/\mu)) = \mathrm{Bin}(\mu, q(p_{t-1,i})) &\overset{\text{Claim 2.3}}{\preccurlyeq} \mathrm{Bin}(\mu, q_{t-1,i})\\ &\overset{(3)}{\preccurlyeq} (\mu \cdot p_{t,i} \mid Y_{0,i} = y_0, \ldots, Y_{t-1,i} = y_{t-1})\\ &\overset{(39)}{\preccurlyeq} (Y_{t,i} \mid Y_{0,i} = y_0, \ldots, Y_{t-1,i} = y_{t-1}).\end{aligned}$$

On the other hand, if $Y_{t-1,i} = y_{t-1,i} = \mu$ or $t-1 \geqslant T_{\mathrm{bad}}$, then either $p_{t-1,i} = 1$ or $t > T_{\mathrm{bad}}$, and so $Y_{t,i}$ is necessarily equal to μ. Thus,

$$\mathrm{Bin}(\mu, q(y_{t-1}/\mu)) \preccurlyeq \mu = (Y_{t,i} \mid Y_{0,i} = y_0, \ldots, Y_{t-1,i} = y_{t-1}).$$

In either case, (41) holds. Therefore, by Corollary B.8, it holds for the first hitting time $T_Z = \min\{t : Z_t \leqslant (\frac{1}{2} - \frac{\delta}{4})\mu\}$ that

$$T^i_{\mathrm{bad}} \geqslant T_Z \quad \text{for each } i \in [n]. \tag{42}$$

Therefore, by applying Lemma B.9,

$$\begin{aligned}\mathbb{P}[T_{\mathrm{bad}} \leqslant 50K\sqrt{n}\log n] &\overset{(40)}{\leqslant} \mathbb{P}[\min_{i\in[n]} T^i_{\mathrm{bad}} \leqslant 50K\sqrt{n}\log n]\\ &\leqslant \sum_{i\in[n]} \mathbb{P}[T^i_{\mathrm{bad}} \leqslant 50K\sqrt{n}\log n]\\ &\overset{(42)}{\leqslant} n \cdot \mathbb{P}[T_Z \leqslant 50K\sqrt{n}\log n]\\ &\overset{(38)}{\leqslant} \tfrac{1}{2}n^{-2K+1} \leqslant \tfrac{1}{2}n^{-K}\end{aligned}$$

as required. □

B.5 Proof of Lemma 3.4

Proof of Lemma 3.4. Let $b = \sum_{k=0}^{m}\binom{n}{k}$, so that b is the size of a Hamming ball of radius m. Note that the entropy bound on the size of a Hamming ball (see [13, Theorem 3.1]) gives us that $b \leqslant 2^{H(m/n)\cdot n}$. We therefore have

$$\begin{aligned}\frac{1}{2^n}\sum_{z\in\mathcal{X}_n}\mathbb{P}(d_H(x,z) \leqslant m) &= \frac{1}{2^n}\sum_{z\in\mathcal{X}_n}\sum_{y\in B_m(z)} p(y)\\ &= \frac{1}{2^n}\sum_{z,y\in\mathcal{X}_n} p(y)\cdot \mathbb{1}(d_H(y,z) \leqslant m)\\ &= \frac{b}{2^n}\sum_{y\in\mathcal{X}_n} p(y) = \frac{b}{2^n} \leqslant 2^{(H(m/n)-1)n}.\end{aligned}$$

Thus, there is some $x^* \in \mathcal{X}_n$ satisfying $\mathbb{P}(d_H(x, x^*) \leqslant m) \leqslant 2^{(H(m/n)-1)n}$. □

B.6 Proof of Lemma 3.5

The proof of Lemma 3.5 depends on the following bound on binomial coefficients.

Lemma B.10. *Suppose $a, b, x, z \in \mathbb{N}$ are such that $x - z$ is even, $x - z \geqslant 0$, and $\frac{1}{2}(x+z) \leqslant a \leqslant b$. Then*

$$\frac{\binom{a}{\frac{1}{2}(x+z)}\binom{b}{\frac{1}{2}(x-z)}}{\binom{a}{\frac{1}{2}(x-z)}\binom{b}{\frac{1}{2}(x+z)}} \leqslant \left(\frac{2a+z}{2b+z}\right)^z. \tag{43}$$

Proof. Note that the fact that all binomial coefficients in (43) are well defined follows from the conditions of the lemma. Note also that $2a - x - z \leqslant 2b - x - z$, and hence for any $0 \leqslant i \leqslant z$,

$$\frac{2a-x-z+2i}{2b-x-z+2i} \leqslant \frac{2a-x-z+2i+(x+2z-2i)}{2b-x-z+2i+(x+2z-2i)} = \frac{2a+z}{2b+z}. \tag{44}$$

Next, observe that for any $c, u, v \geqslant 0$ we have

$$\begin{aligned}\frac{\binom{c}{u+v}}{\binom{c}{u-v}} = \frac{(u-v)!(c-u+v)!}{(u+v)!(c-u-v)!} &= \prod_{i=1}^{2v}\left(\frac{c-u-v+i}{u-v+i}\right)\\ &= \prod_{i=1}^{2v}\left(\frac{2c-2u-2v+2i}{2u-2v+2i}\right).\end{aligned} \tag{45}$$

Therefore,

$$\begin{aligned}\frac{\binom{a}{\frac{1}{2}x+\frac{1}{2}z}\binom{b}{\frac{1}{2}x-\frac{1}{2}z}}{\binom{a}{\frac{1}{2}x-\frac{1}{2}z}\binom{b}{\frac{1}{2}x+\frac{1}{2}z}} &\overset{(45)}{=} \prod_{i=1}^{z}\left[\left(\frac{2a-x-z+2i}{x-z+2i}\right)\left(\frac{x-z+2i}{2b-x-z+2i}\right)\right]\\ &\leqslant \prod_{i=1}^{z}\left(\frac{2a-x-z+2i}{2b-x-z+2i}\right) \overset{(44)}{\leqslant} \left(\frac{2a+z}{2b+z}\right)^z,\end{aligned}$$

as required. □

Proof of Lemma 3.5. Let us write $[0, n/2]$ as a shorthand for $\{0, \ldots, \lfloor n/2\rfloor\}$. Let $y_0(x), y_1(x), \ldots, y_{\lfloor n/2\rfloor}(x)$ be independent random variables taking values in $\mathcal{X}_n$, such that $y_r(x)$ is distributed according to $\mathrm{Unif}(S_r(x))$ for each $r \in [0, n/2]$. Given x, let $R \sim s([0, n/2])$ and $y(x) = y_R(x)$ so that $y(x)$ has probability mass function p_x. Note that for any set $A \subseteq \mathcal{X}_n$,

$$p_x(A) \leqslant \sup_{r\in[0,n/2]} \mathbb{P}(y_r(x) \in A). \tag{46}$$

Note that, for any $r \geqslant 0$, if we flip a uniformly random subset of r bits of x, then the probability that a given bit is flipped is r/n. In particular, the expected number of 0-bits of x that are flipped is $\frac{r}{n}(n - |x|)$ and the expected number of 1-bits that are not flipped is $(1 - \frac{r}{n})|x|$, and so we have

$$\mathbb{E}[|y_r(x)|] = (1 - \tfrac{r}{n})|x| + \tfrac{r}{n}(n - |x|) = |x| + r - \tfrac{2r}{n}|x|. \tag{47}$$

Note that by Theorem A.1, it holds for any $r \in [0, n/2]$ and $t > 0$ that

$$\begin{aligned}\mathbb{P}(|y_r(x)| \geqslant \mathbb{E}[|y_r(x)|] + t) &\leqslant \exp(-2t^2/r) &(48)\\ &\leqslant \exp(-4t^2/n). &(49)\end{aligned}$$

We will now verify the properties **C1**-**C3** in turn.

C1: Suppose that $|x| < (\frac{1}{2} + 2\varepsilon + 3\eta)n$. If $r \leqslant d/2$, then

$$\mathbb{P}(y_r(x) \in A_{\geqslant(\frac{1}{2}+2\varepsilon)n} \setminus B_{d/2}(x)) = 0 \leqslant e^{-\sqrt{n}}.$$

On the other hand, if $d/2 \leqslant r \leqslant n/2$ then $\varepsilon r \geqslant 50\eta n$, and so

$$\mathbb{E}[|y_r(x)|] \overset{(47)}{=} (1 - \tfrac{2r}{n})|x| + r \leqslant (1 - \tfrac{2r}{n})(\tfrac{1}{2} + 2\varepsilon + 3\eta)n + r$$
$$\leqslant (\tfrac{1}{2} + 2\varepsilon + 3\eta)n - 4\varepsilon r \leqslant (\tfrac{1}{2} + 2\varepsilon - 100\eta)n. \tag{50}$$

Therefore,

$$\mathbb{P}(y_r(x) \in A_{\geqslant(\frac{1}{2}+2\varepsilon)n} \setminus B_{d/2}(x)) \leqslant \mathbb{P}(|y_r(x)| \geqslant (\tfrac{1}{2} + 2\varepsilon)n)$$
$$\overset{(50)}{\leqslant} \mathbb{P}(|y_r(x)| \geqslant \mathbb{E}[|y_r(x)|] + 100\eta n)$$
$$\overset{(49)}{\leqslant} \exp\left(-\frac{4(100\eta n)^2}{n}\right) \leqslant \exp\left(-\frac{4(100\eta_0)^2 n}{(\log n)^2}\right) \leqslant e^{-\sqrt{n}}.$$

Putting the two cases together, we have

$$p_x\left(A_{\geqslant(\frac{1}{2}+2\varepsilon)n} \setminus B_{d/2}(x)\right) \overset{(46)}{\leqslant} e^{-\sqrt{n}},$$

and so **C1** holds.

C2: Suppose first that $|x| \leqslant (\frac{1}{2} + \varepsilon)n$ and $r \in [0, n/2]$. In this case we have

$$\mathbb{E}[|y_r(x)|] \overset{(47)}{=} (1 - \tfrac{2r}{n})|x| + r$$
$$\leqslant (1 - \tfrac{2r}{n})(\tfrac{1}{2} + \varepsilon)n + r \leqslant (\tfrac{1}{2} + \varepsilon)n. \tag{51}$$

Thus we have for any $j \geqslant 0$,

$$\mathbb{P}(y_r(x) \in A_{\geqslant m+j}) \leqslant \mathbb{P}(|y_r(x)| \geqslant (\tfrac{1}{2} + 2\varepsilon)n + j)$$
$$\overset{(51)}{\leqslant} \mathbb{P}(|y_r(x)| \geqslant \mathbb{E}[|y_r(x)|] + (\varepsilon n + j))$$
$$\overset{(49)}{\leqslant} \exp\left(-\frac{4(\varepsilon n + j)^2}{n}\right) \leqslant e^{-8\varepsilon j}.$$

On the other hand, if $|x| \geqslant (\frac{1}{2} + 2\varepsilon)n$, then for any $r \in [0, n/2]$,

$$\mathbb{E}[|y_r(x)|] \overset{(47)}{=} |x| + r - \tfrac{2r}{n}|x|$$
$$\leqslant |x| + r - \tfrac{2r}{n}(\tfrac{1}{2} + 2\varepsilon)n \leqslant |x| - 4\varepsilon r, \tag{52}$$

and hence,

$$\mathbb{P}(y_r(x) \in A_{\geqslant m+j}) \leqslant \mathbb{P}(|y_r(x)| \geqslant |x| + j)$$
$$\overset{(52)}{\leqslant} \mathbb{P}(|y_r(x)| \geqslant \mathbb{E}[|y_r(x)|] + (4\varepsilon r + j))$$
$$\overset{(48)}{\leqslant} \exp\left(-\frac{2(4\varepsilon r + j)^2}{r}\right) \leqslant e^{-8\varepsilon j}.$$

Putting the two cases together, we have

$$p_x\left(A_{\geqslant m+j}\right) \overset{(46)}{\leqslant} e^{-8\varepsilon j},$$

and so **C2** holds.

C3: Let us first consider $|S_r(x) \cap A_{|x|+j}|$ for arbitrary $j \in \mathbb{Z}$. If y is obtained from x by flipping r_0 0-bits and r_1 1-bits, then we have $|y| = |x| + r_0 - r_1$. In particular, if $y \in S_r(x) \cap A_{|x|+j}$, then we can obtain y from x by flipping r_0 0-bits and r_1 1-bits, where r_0 and r_1 satisfy

$$r_0 + r_1 = r \qquad \text{(because } y \in S_r(x)\text{)},$$
$$r_0 - r_1 = j \qquad \text{(because } y \in A_{|x|+j}\text{)}.$$

These equations are solved by setting $r_0 = \frac{1}{2}(r+j)$ and $r_1 = \frac{1}{2}(r-j)$. Because x has $n - |x|$ 0-bits and $|x|$ 1-bits to flip, we can make two observations. First, $S_r(x) \cap A_{|x|+j}$ is non-empty if and only if $\frac{1}{2}(r+j)$ is a non-negative integer satisfying $\frac{1}{2}(r+j) \leqslant n - |x|$ and $\frac{1}{2}(r-j)$ is a non-negative integer satisfying $\frac{1}{2}(r-j) \leqslant |x|$. Second, in the case where $S_r(x) \cap A_{|x|+j}$ is non-empty, we can calculate its size exactly by counting the number of ways of choosing a set of $\frac{1}{2}(r+j)$ 0-bits to flip and $\frac{1}{2}(r-j)$ 1-bits to flip, as follows.

$$|S_r(x) \cap A_{|x|+j}| = \binom{n-|x|}{\frac{1}{2}(r+j)}\binom{|x|}{\frac{1}{2}(r-j)}. \tag{53}$$

With these observations in mind, we are now in a position to verify **C3**. Let $0 < k \leqslant \eta n$ be fixed. First, if $r \in [0, n/2]$ is any number such that $S_r(x) \cap A_{|x|+k}$ is empty, then we have

$$0 = |S_r(x) \cap A_{|x|+k}| \leqslant (1 - 4\varepsilon)^k \cdot |S_r(x) \cap A_{|x|+k}|.$$

On the other hand, if $r \in [0, n/2]$ is any number such that $S_r(x) \cap A_{|x|+k}$ is non-empty, then $r + k$ and $r - k$ are both non-negative and even and $\frac{1}{2}(r+k) \leqslant n - |x|$. Moreover, because $|x| \geqslant \frac{1}{2}n$, it holds that

$$\tfrac{1}{2}(r-k) \leqslant \tfrac{1}{2}(r+k) \leqslant n - |x| \leqslant |x|.$$

Therefore $S_r(x) \cap A_{|x|-k}$ is also non-empty, and we may compute using Lemma B.10 that

$$\frac{|S_r(x) \cap A_{|x|+k}|}{|S_r(x) \cap A_{|x|-k}|} \overset{(53)}{=} \frac{\binom{n-|x|}{\frac{1}{2}(r+k)}\binom{|x|}{\frac{1}{2}(r-k)}}{\binom{n-|x|}{\frac{1}{2}(r-k)}\binom{|x|}{\frac{1}{2}(r+k)}} \overset{(43)}{\leqslant} \left(\frac{2(n-|x|)+k}{2|x|+k}\right)^k$$
$$= \left(\frac{n - |x| + \frac{k}{2}}{|x| + \frac{k}{2}}\right)^k \leqslant \left(\frac{n + \frac{k}{2} - |x|}{|x|}\right)^k$$
$$\leqslant \left(\frac{(1+\eta/2)n - |x|}{|x|}\right)^k \leqslant \left(\frac{1+\varepsilon/2}{\frac{1}{2}+3\varepsilon} - 1\right)^k$$
$$\leqslant (1 - 4\varepsilon)^k.$$

In either case, we have for any $r \in [0, n/2]$ that

$$|S_r(x) \cap A_{|x|+k}| \leqslant (1 - 4\varepsilon)^k \cdot |S_r(x) \cap A_{|x|-k}|. \tag{54}$$

Therefore,

$$p_x(A_{|x|+k}) = \mathbb{P}(|y(x)| = |x| + k)$$
$$= \sum_{r \in [0,n/2]} \mathbb{P}(R = r) \cdot \mathbb{P}(|y_r(x)| = |x| + k)$$
$$= \sum_{r \in [0,n/2]} \mathbb{P}(R = r) \cdot \left(\frac{|S_r(x) \cap A_{|x|+k}|}{|S_r(x)|}\right)$$
$$\overset{(54)}{\leqslant} (1 - 4\varepsilon)^k \cdot \sum_{r \in [0,n/2]} \mathbb{P}(R = r) \cdot \left(\frac{|S_r(x) \cap A_{|x|-k}|}{|S_r(x)|}\right)$$
$$= (1 - 4\varepsilon)^k \cdot \sum_{r \in [0,n/2]} \mathbb{P}(R = r) \cdot \mathbb{P}(|y_r(x)| = |x| - k)$$
$$= (1 - 4\varepsilon)^k \cdot p_x(A_{|x|-k}),$$

and so **C3** holds. □

B.7 Proof of Lemma 3.6

Proof of Lemma 3.6. Let $\hat{x} \in D$ be fixed and write $S = \mathrm{supp}(p)$. Given $A \in \mathrm{Mat}_{\mu+1,\mu+1}(\mathbb{R})$, let $j(A)$ be a random variable over $[\mu+1]$ distributed according to $\mathcal{S}(A)$. Recall from the statement of the lemma that $j \sim \mathcal{S}(A_f(y_1, \ldots, y_\mu, \overline{x}))$ where $y_1, \ldots, y_\mu$ are sampled independently according to p. We first claim that there exist constants $c_1, \ldots, c_\mu$ such that $\mathbb{P}(j = i \wedge y_i = x) = c_i \cdot p(x)$ for every $i \in [\mu]$ and $x \in D \setminus \{\overline{x}\}$. Indeed, if $x \in D$ then it holds for any $z \in S$ that $f(x, z) = f(\hat{x}, z)$ and, using also that f is antisymmetric, $f(z, x) = f(z, \hat{x})$. Hence,

$$A_f(x, z_2, \ldots, z_\mu, \overline{x}) = A_f(\hat{x}, z_2, \ldots, z_\mu, \overline{x}). \tag{55}$$

holds for every $z_2, \ldots, z_\mu \in S$. Therefore, if $x \in D$,

$$\begin{aligned}
&\mathbb{P}(j = 1 \wedge y_1 = x)\\
&= \sum_{z_2,\ldots,z_\mu \in S} \left(p(x) \cdot \textstyle\prod_{i=2}^{\mu} p(z_i)\right) \cdot \mathbb{P}(j(A_f(x, z_2, \ldots, z_\mu, \overline{x})) = 1)\\
&\overset{(55)}{=} \sum_{z_2,\ldots,z_\mu \in S} \left(p(x) \cdot \textstyle\prod_{i=2}^{\mu} p(z_i)\right) \cdot \mathbb{P}(j(A_f(\hat{x}, z_2, \ldots, z_\mu, \overline{x})) = 1).
\end{aligned}$$

So, if one defines

$$c_1 := \sum_{z_2,\ldots,z_\mu \in S} \textstyle\prod_{i=2}^{\mu} p(z_i) \cdot \mathbb{P}(j(A_f(\hat{x}, z_2, \ldots, z_\mu, \overline{x})) = 1),$$

then $\mathbb{P}(j = 1 \wedge y_1 = x) = c_1 \cdot p(x)$ holds for every $x \in D$ (where c_1 has no dependence on x). A similar argument shows that such a c_i exists for all other $i \in [\mu]$. Therefore, if $c = \sum_{i \in [\mu]} c_i$, then

$$\mathbb{P}(y = x) = \sum_{i \in [\mu]} \mathbb{P}(j = i \wedge y_i = x) = \sum_{i \in [\mu]} c_i \cdot p(x) = c \cdot p(x),$$

holds for every $x \in D$, as required. □

B.8 Proof of Claim 3.7

Claim 3.7 is an immediate consequence of the following lemma.

Lemma B.11. *If $|x| = |x'|$, then for any $k \in \{0, \ldots, n\}$,*

$$\mathbb{P}(|x_{t+1}| = k \mid x_t = x) = \mathbb{P}(|x_{t+1}| = k \mid x_t = x'). \tag{56}$$

Proof. Let $y_1^{(t+1)}, \ldots, y_\mu^{(t+1)}$ be the mutants of x_t sampled according to $\mathcal{M}_s(x_t)$, and let j_{t+1} denote the index sampled according to $\mathcal{S}(A_f(y_1^{(t+1)}, \ldots, y_\mu^{(t+1)}, x_t))$ in generation t of Algorithm 3, so that $x_{t+1} = y_{j_{t+1}}^{(t+1)}$ whenever $j_{t+1} \leqslant \mu$. Given $A \in \mathrm{Mat}_{\mu+1,\mu+1}$, let $q_1(A)$ denote the probability that $j = 1$ if j is sampled according to $\mathcal{S}(A)$. Given $y \in \mathcal{X}_n$, let $p_y \in \mathcal{P}(\mathcal{X}_n)$ be the probability mass function corresponding to the unbiased mutation operator $\mathcal{M}_s(y)$.

Let us define the function $h_1 : \mathcal{X}_n \times \mathcal{X}_n \to \mathbb{R}$ by

$$h_1(x, y) = \sum_{z_2,\ldots,z_\mu \in \mathcal{X}_n} \left(\textstyle\prod_{i=2}^{\mu} p_x(z_i)\right) \cdot q_1(A_g(y, z_2, \ldots, z_\mu, x)).$$

Note that, for any $x, y \in \mathcal{X}_n$,

$$\mathbb{P}(j_t = 1 \wedge y_1^{(t+1)} = y \mid x_t = x) = p_x(y) \cdot h_1(x, y).$$

To prove the lemma, let $x, x' \in \mathcal{X}_n$ be arbitrary with $|x| = |x'|$ and recall that we need to show that (56) holds. Let $\sigma : \mathcal{X}_n \to \mathcal{X}_n$ be a bijection which rearranges bit positions according to a fixed permutation which satisfies $\sigma(x) = x'$. Note that $|y| = |\sigma(y)|$ and $d_H(y, z) = d_H(\sigma(y), \sigma(z))$ holds for any $y, z \in \mathcal{X}_n$. By considering (7), $g(y, z)$ depends only on $|y|$, $|z|$, and $d_H(y, z)$, and so it follows that $g(y, z) = g(\sigma(y), \sigma(z))$ for any $y, z \in \mathcal{X}_n$. Hence,

$$A_g(y, z_2, \ldots, z_\mu, x) = A_g(\sigma(y), \sigma(z_2), \ldots, \sigma(z_\mu), x') \tag{57}$$

for any $y, z_2, \ldots, z_\mu \in \mathcal{X}_n$. Additionally, recalling from Definition 3.1 that

$$p_y(z) = \frac{s(d_H(y, z))}{\binom{n}{d_H(y,z)}},$$

we have

$$p_x(z) = p_{x'}(\sigma(z)) \tag{58}$$

for any $z \in \mathcal{X}_n$. Therefore, for any $y \in \mathcal{X}_n$,

$$\begin{aligned}
h_1(x, y) &= \sum_{z_2,\ldots,z_\mu \in \mathcal{X}_n} \left(\textstyle\prod_{i=2}^{\mu} p_x(z_i)\right) \cdot q_1(A_g(y, z_2, \ldots, z_\mu, x))\\
&\overset{(57),(58)}{=} \sum_{z_2,\ldots,z_\mu \in \mathcal{X}_n} \left(\textstyle\prod_{i=2}^{\mu} p_{x'}(\sigma(z_i))\right) \cdot q_1(A_g(\sigma(y), \sigma(z_2), \ldots, \sigma(z_\mu), x')\\
&= \sum_{z_2,\ldots,z_\mu \in \mathcal{X}_n} \left(\textstyle\prod_{i=2}^{\mu} p_{x'}(z_i)\right) \cdot q_1(A_g(\sigma(y), z_2, \ldots, z_\mu, x')\\
&= h_1(x', \sigma(y)).
\end{aligned} \tag{59}$$

Given k, let $A_k := \{y \in \mathcal{X}_n : |y| = k\}$. We can now deduce that for any k, because the restriction of σ to A_k is a bijection,

$$\begin{aligned}
\mathbb{P}(j_t = 1 \wedge |y_1^{(t+1)}| = k \mid x_t = x) &= \sum_{y \in A_k} p_x(y) \cdot h_1(x, y)\\
&\overset{(58),(59)}{=} \sum_{y \in A_k} p_{x'}(\sigma(y)) \cdot h_1(x', \sigma(y))\\
&= \sum_{y \in A_k} p_{x'}(y) \cdot h_1(x', y)\\
&= \mathbb{P}(j_t = 1 \wedge |y_1^{(t+1)}| = k \mid x_t = x').
\end{aligned}$$

Arguing similarly shows that in fact

$$\mathbb{P}(j_t = i \wedge |y_i^{(t+1)}| = k \mid x_t = x) = \mathbb{P}(j_t = i \wedge |y_i^{(t+1)}| = k \mid x_t = x')$$

holds for all $i \in [\mu]$. Thus, for any $k \neq |x|$,

$$\begin{aligned}
\mathbb{P}(|x_{t+1}| = k \mid x_t = x) &= \sum_{i \in [\mu]} \mathbb{P}(j_t = i \wedge |y_i^{(t+1)}| = k \mid x_t = x)\\
&= \sum_{i \in [\mu]} \mathbb{P}(j_t = i \wedge |y_i^{(t+1)}| = k \mid x_t = x')\\
&= \mathbb{P}(|x_{t+1}| = k \mid x_t = x').
\end{aligned} \tag{60}$$

and so (56) holds whenever $k \neq |x|$. The first line of the above calculation does not hold for $k = |x|$, as $\mathbb{P}(j_t = \mu + 1 \mid x_t = x)$ must also be accounted for. Nonetheless, by noting that

$$\mathbb{P}(|x_{t+1}| = |x| \mid x_t = x) = 1 - \sum_{k \neq |x|} \mathbb{P}(|x_{t+1}| = k \mid x_t = x),$$

we then have that (56) holds for $k = |x|$ as well, thus proving the lemma. □

B.9 Proof of Claims 3.8 and 3.9

To prove these claims, let us first readopt some of the notation from Lemma 3.5. Let $p_x \in \mathcal{P}(\mathcal{X}_n)$ be the probability measure corresponding to sampling according to $\mathcal{M}_s(x)$. Define also $\hat{p}_x$ to be the probability measure corresponding to sampling according to $\mathcal{M}_s(x)$ conditioned on the set $A \cup B_{d/2}(x)$ (where we recall that

$A = \{x \in \mathcal{X}_n : |x| < (\frac{1}{2} + 2\varepsilon)n\})$. Note that, for every $x \in \mathcal{X}_n$ and $y \in A \cup B_{d/2}(x)$ we have

$$p_x(y) = \hat{p}_x(y) \cdot p_x(A \cup B_{d/2}(x)). \tag{61}$$

Given j, write $A_j = \{y \in \mathcal{X}_n : |y| = j\}$ and $A_{\geqslant j} = \cup_{j' \geqslant j} A_{j'}$. Let us also use $y_1^{(t+1)}, \dots, y_\mu^{(t+1)}$ to denote the mutants of x_t sampled according to $\mathcal{M}_s(x_t)$ in generation t of Algorithm 3.

The proof of Claim 3.8 is very straightforward, and so we give this quickly now.

Proof of Claim 3.8. Let $x \in \mathcal{X}_n$ satisfy $|x| = i$ and note that $i' = \max\{|x|, (\frac{1}{2} + 2\varepsilon)n\}$. Note also that i' may be identified with m in property **C2** of Lemma 3.5. By applying a union bound,

$$\begin{aligned}\sum_{k=j}^{n-i'} q_{i,i'+k} &\overset{(9)}{\leqslant} \mathbb{P}(|x_{t+1}| \geqslant i' + j \mid x_t = x)\\ &\leqslant \sum_{i \in [\mu]} \mathbb{P}(|y_i^{(t+1)}| \geqslant i' + j \mid x_t = x)\\ &= \mu \cdot p_x(A_{\geqslant i'+j}) \overset{\mathbf{C2}}{\leqslant} \mu \cdot e^{-8\varepsilon j} \overset{(8)}{\leqslant} e^{\varepsilon(n^{1/3} - 8j)},\end{aligned}$$

as required. □

To prove Claim 3.9, we will at one point assume that all $y_i^{(t+1)}$ generated are in the set $A \cup B_{d/2}(x_t)$, so that Lemma 3.6 can be applied (with $D = B_{d/2}(x_t) \setminus A$). For this reason, let G_{t+1} be the event that $y_i^{(t+1)} \in A \cup B_{d/2}(x_t)$ for every $i \in [\mu]$. If $x \in A \cup B$, we have

$$\begin{aligned}\mathbb{P}(G_{t+1} \mid x_t = x) &= \prod_{i \in [n]} \mathbb{P}(y_i^{(t+1)} \in A \cup B_{d/2}(x))\\ &\overset{\mathbf{C1}}{\geqslant} (1 - e^{-\sqrt{n}})^\mu \geqslant 1 - \mu e^{-\sqrt{n}} \overset{(8)}{\geqslant} 1 - \gamma^{10}.\end{aligned} \tag{62}$$

In fact, we have the following lemma (where we recall that $S = \{0, \dots, n\}$).

Lemma B.12. *There exist constants $\hat{q}_{i,j}$ for $i, j \in S$ such that*

$$\hat{q}_{i,j} = \mathbb{P}(|x_{t+1}| = j \wedge G_{t+1} \mid x_t = x)$$

whenever $|x| = i$.

Proof. It suffices to show that if $|x| = |x'|$, then for any $j \in \{0, \dots, n\}$,

$$\mathbb{P}(|x_{t+1}| = j \wedge G_{t+1} \mid x_t = x) = \mathbb{P}(|x_{t+1}| = j \wedge G_{t+1} \mid x_t = x').$$

Accordingly, this result follows almost identically to Lemma B.11. In the proof, one only needs to change the sum in the definition of h_1 to be over $z_2, \dots, z_\mu \in A \cup B_{d/2}(x)$ and redefine $A_k := \{y \in A \cup B_{d/2}(x) : |y| = k\}$. □

Proof of Claim 3.9. Let $\hat{q}_{i,j}$ be as defined in Lemma B.12. Noting that $\sum_{j=0}^n \hat{q}_{i,j} = \mathbb{P}(G_{t+1} \mid |x_t| = i)$, we have the following simple properties of $q_{i,j}$ and $\hat{q}_{i,j}$.

$$q_{i,j} \geqslant \hat{q}_{i,j} \qquad \text{for every } i, j \in S, \tag{63}$$

$$\sum_{j=0}^n q_{i,j} = 1 \qquad \text{for every } i \in S, \tag{64}$$

$$\sum_{j=0}^n \hat{q}_{i,j} \overset{(62)}{\geqslant} 1 - \gamma^{10} \qquad \text{if } i < (\tfrac{1}{2} + 2\varepsilon + 3\eta)n. \tag{65}$$

We will use $\hat{q}_{i,j}$ as an approximation to $q_{i,j}$ which is more amenable to calculation. First, we will establish that if $0 < k \leqslant \eta n$,

$$\hat{q}_{i,i+k} \leqslant (1 - 4\varepsilon)^k \cdot \hat{q}_{i,i-k}. \tag{66}$$

Indeed, suppose that $x \in \mathcal{X}_n$ satisfies $|x| = i$, so that for any $j \in S$ we have

$$\hat{q}_{i,j} = \mathbb{P}(|x_{t+1}| = j \wedge G_{t+1} \mid x_t = x).$$

Let

$$D := B_{d/2}(x) \setminus A = \mathrm{supp}(\hat{p}_x) \cap B \subseteq B,$$

and note that $d_H(x, y) \leqslant d$ whenever $x, y \in D$, and also that $\mathrm{supp}(\hat{p}_x) \setminus D \subseteq A$. If $y, z \in D$ then $y, z \in B$ and $d_H(y, z) \leqslant d_H(y, x) + d_H(x, z) \leqslant d$, and so by referring to the definition of g given by (7), we have $g(y, z) = 0$. On the other hand, if $y \in D$ and $z \in \mathrm{supp}(\hat{p}_x) \setminus D$ then $y \in B$ and $z \in A$ (and hence $|y| > |z|$), and so we have $g(y, z) = 1$. Therefore, by applying Lemma 3.6 (with $p = \hat{p}_x$, $f_0 = 0$, $f_1 = 1$, $\overline{x} = x$, and identifying f with g), there is a constant c_x such that for any $x' \in D \setminus \{x\}$,

$$\mathbb{P}(x_{t+1} = x' \mid x_t = x \wedge G_{t+1}) = c_x \cdot \hat{p}_x(x'). \tag{67}$$

From this it follows that for any $E \subseteq D \setminus \{x\}$,

$$\begin{aligned}&\mathbb{P}((x_{t+1} \in E) \wedge G_{t+1} \mid x_t = x)\\ &\quad= \mathbb{P}(G_{t+1} \mid x_t = x) \cdot \mathbb{P}(x_{t+1} \in E \mid x_t = x \wedge G_{t+1})\\ &\quad= \mathbb{P}(G_{t+1} \mid x_t = x) \cdot \sum_{x' \in E} \mathbb{P}(x_{t+1} = x' \mid x_t = x \wedge G_{t+1})\\ &\quad\overset{(67)}{=} \mathbb{P}(G_{t+1} \mid x_t = x) \cdot \sum_{x' \in E} c_x \cdot \hat{p}_x(x')\\ &\quad= \mathbb{P}(G_{t+1} \mid x_t = x) \cdot c_x \cdot \hat{p}_x(E).\end{aligned} \tag{68}$$

In fact, we can say more generally that for any $E \subseteq B \setminus \{x\}$, because $E \cap D = E \cap \mathrm{supp}(\hat{p}_x)$,

$$\begin{aligned}\mathbb{P}((x_{t+1} \in E) \wedge G_{t+1} \mid x_t = x)&\\ &= \mathbb{P}(x_{t+1} \in E \cap D \wedge G_{t+1} \mid x_t = x)\\ &\overset{(68)}{=} \mathbb{P}(G_{t+1} \mid x_t = x) \cdot c_x \cdot \hat{p}_x(E \cap D)\\ &= \mathbb{P}(G_{t+1} \mid x_t = x) \cdot c_x \cdot \hat{p}_x(E)\\ &\overset{(61)}{=} \left(\frac{\mathbb{P}(G_{t+1} \mid x_t = x) \cdot c_x}{p_x(A \cup B_{d/2}(x))}\right) \cdot p_x(E).\end{aligned} \tag{69}$$

Recalling that $B = \{x' \in \mathcal{X}_n : (\frac{1}{2} + 2\varepsilon)n \leqslant |x'| < (\frac{1}{2} + 2\varepsilon + 3\eta)n\}$ and $(\frac{1}{2} + 2\varepsilon + \eta)n < |x| < (\frac{1}{2} + 2\varepsilon + 2\eta)n$, we have that $A_{|x|+k} \subseteq B \setminus \{x\}$ and $A_{|x|-k} \subseteq B \setminus \{x\}$ whenever $0 < k \leqslant \eta n$. In particular, if we define

$$\hat{c}_x = \left(\frac{\mathbb{P}(G_{t+1} \mid x_t = x) \cdot c_x}{p_x(A \cup B_{d/2}(x))}\right),$$

then, if $0 < k \leqslant \eta n$,

$$\begin{aligned}
&\mathbb{P}(|x_{t+1}| = |x| + k \wedge G_{t+1} \mid x_t = x) \\
&\quad = \mathbb{P}((x_{t+1} \in A_{|x|+k}) \wedge G_{t+1} \mid x_{t+1} = x) \\
&\quad \overset{(69)}{=} \hat{c}_x \cdot p_x(A_{|x|+k}) \overset{\textbf{C3}}{\leqslant} (1-4\varepsilon)^k \cdot \hat{c}_x \cdot p_x(A_{|x|-k}) \\
&\quad \overset{(69)}{=} (1-4\varepsilon)^k \cdot \mathbb{P}((x_{t+1} \in A_{|x|-k}) \wedge G_{t+1} \mid x_t = x) \\
&\quad = (1-4\varepsilon)^k \cdot \mathbb{P}(|x_{t+1}| = |x| - k \wedge G_{t+1} \mid x_t = x),
\end{aligned}$$

and so (66) holds.

Because $C(\varepsilon) = 8\varepsilon + 2 \cdot \max_{k \geqslant 0} \{k(1-4\varepsilon)^k\}$, we have

$$C(\varepsilon) \geqslant 8\varepsilon, \tag{70}$$

and, for any $k \geqslant 0$,

$$k(1-4\varepsilon)^k \leqslant \tfrac{C(\varepsilon)}{2}. \tag{71}$$

Next, we will use (66) to establish that if $0 < k \leqslant C(\varepsilon)$,

$$h(k) \cdot \hat{q}_{i,i-k} + h(-k) \cdot \hat{q}_{i,i+k} \geqslant 2\varepsilon \cdot (\hat{q}_{i,i-k} + \hat{q}_{i,i+k}). \tag{72}$$

First note that (66) certainly implies that $\hat{q}_{i,i-k} \geqslant \hat{q}_{i,i+k}$ whenever $0 < k \leqslant \eta n$, and hence

$$4\varepsilon \cdot \hat{q}_{i,i-k} \geqslant 2\varepsilon \cdot (\hat{q}_{i,i-k} + \hat{q}_{i,i+k}). \tag{73}$$

Now, if $0 < k \leqslant C(\varepsilon)$ then

$$\begin{aligned}
&h(k) \cdot \hat{q}_{i,i-k} + h(-k) \cdot \hat{q}_{i,i+k} \overset{(10)}{=} k \cdot (\hat{q}_{i,i-k} - \hat{q}_{i,i+k}) \\
&\quad \overset{(66)}{\geqslant} k \cdot (\hat{q}_{i,i-k} - (1-4\varepsilon)^k \cdot \hat{q}_{i,i-k}) \geqslant 4\varepsilon \cdot \hat{q}_{i,i-k} \\
&\quad \overset{(73)}{\geqslant} 2\varepsilon \cdot (\hat{q}_{i,i-k} + \hat{q}_{i,i+k}).
\end{aligned}$$

On the other hand, if $C(\varepsilon) < k \leqslant \eta n$ then

$$\begin{aligned}
&h(k) \cdot \hat{q}_{i,i-k} + h(-k) \cdot \hat{q}_{i,i+k} \overset{(10)}{=} C(\varepsilon) \cdot \hat{q}_{i,i-k} - k \cdot \hat{q}_{i,i+k} \\
&\quad \overset{(66)}{\geqslant} C(\varepsilon) \cdot \hat{q}_{i,i-k} - k(1-4\varepsilon)^k \cdot \hat{q}_{i,i-k} \overset{(71)}{\geqslant} \tfrac{C(\varepsilon)}{2} \cdot \hat{q}_{i,i-k} \\
&\quad \overset{(70)}{\geqslant} \varepsilon \cdot \hat{q}_{i,i-k} \overset{(73)}{\geqslant} 2\varepsilon \cdot (\hat{q}_{i,i-k} + \hat{q}_{i,i+k}).
\end{aligned}$$

In either case, (72) holds.

Our attention can now turn to the sum $\sum_{k=i-n}^{i} h(k) \cdot \hat{q}_{i,i-k}$. First, note that $\max\{i, (\frac{1}{2} + 2\varepsilon)n\} = i$, and so applying Claim 3.8 with $j = \eta n$ yields

$$\sum_{k=\eta n+1}^{n-i} q_{i,i+k} \leqslant \sum_{k=\eta n}^{n-i} q_{i,i+k} \leqslant e^{\varepsilon(n^{1/3} - 8\eta n)}. \tag{74}$$

Therefore,

$$\begin{aligned}
\sum_{k=i-n}^{-\eta n-1} (h(k) - 2\varepsilon) \cdot \hat{q}_{i,i-k} &\overset{(63),(10)}{\geqslant} -n \sum_{k=i-n}^{-\eta n-1} q_{i,i-k} \\
&= -n \sum_{k=\eta n+1}^{n-i} q_{i,i+k} \overset{(74)}{\geqslant} -n \cdot e^{\varepsilon(n^{1/3} - 8\eta n)} \geqslant -2\varepsilon\gamma^{10},
\end{aligned}$$

and so

$$\sum_{k=i-n}^{\eta n-1} h(k) \cdot \hat{q}_{i,i-k} \geqslant 2\varepsilon \left(\sum_{k=i-n}^{\eta n-1} \hat{q}_{i,i-k} - \gamma^{10} \right). \tag{75}$$

Next, note that

$$\begin{aligned}
\sum_{k=\eta n+1}^{i} h(k) \cdot \hat{q}_{i,i-k} &\overset{(10)}{\geqslant} \min\{\eta n + 1, C(\varepsilon)\} \sum_{k=\eta n+1}^{i} \hat{q}_{i,i-k} \\
&\overset{(70)}{\geqslant} 2\varepsilon \sum_{k=\eta n+1}^{i} \hat{q}_{i,i-k}
\end{aligned} \tag{76}$$

and

$$\begin{aligned}
\sum_{k=-\eta n}^{\eta n} h(k) \cdot \hat{q}_{i,i-k} &= h(0) \cdot \hat{q}_{i,i} + \sum_{k=1}^{\eta n} (h(k) \cdot \hat{q}_{i,i-k} + h(-k) \cdot \hat{q}_{i,i+k}) \\
&\overset{(72)}{\geqslant} 0 + 2\varepsilon \sum_{k=1}^{\eta n} (\hat{q}_{i,i-k} + \hat{q}_{i,i+k}) \\
&= 2\varepsilon \left(\sum_{k=-\eta n}^{\eta n} \hat{q}_{i,i-k} - \hat{q}_{i,i} \right).
\end{aligned} \tag{77}$$

Combining these observations, we deduce that

$$\begin{aligned}
&\sum_{k=i-n}^{i} h(k) \cdot \hat{q}_{i,i-k} \\
&= \sum_{k=i-n}^{-\eta n-1} h(k) \cdot \hat{q}_{i,i-k} + \sum_{k=-\eta n}^{\eta n} h(k) \cdot \hat{q}_{i,i-k} + \sum_{k=\eta n+1}^{i} h(k) \cdot \hat{q}_{i,i-k} \\
&\overset{(75),(76),(77)}{\geqslant} 2\varepsilon \left(\sum_{k=i-n}^{i} \hat{q}_{i,i-k} - \gamma^{10} - \hat{q}_{i,i} \right) \\
&\overset{(63),(65)}{\geqslant} 2\varepsilon(1 - 2\gamma^{10} - q_{i,i}) \geqslant 2\varepsilon(1 - q_{i,i} - \gamma^9).
\end{aligned} \tag{78}$$

Finally, because (63) implies that $q_{i,i-k} - \hat{q}_{i,i-k}$ is always non-negative,

$$\begin{aligned}
\sum_{k=i-n}^{i} h(k) \cdot q_{i,i-k} &= \sum_{k=i-n}^{i} h(k) \cdot [\hat{q}_{i,i-k} + q_{i,i-k} - \hat{q}_{i,i-k}] \\
&\geqslant \sum_{k=i-n}^{i} h(k) \cdot \hat{q}_{i,i-k} - \sum_{k=i-n}^{i} n \cdot (q_{i,i-k} - \hat{q}_{i,i-k}) \\
&\overset{(78)}{\geqslant} 2\varepsilon(1 - q_{i,i} - \gamma^9) - \sum_{k=i-n}^{i} n \cdot (q_{i,i-k} - \hat{q}_{i,i-k}) \\
&\overset{(64),(65)}{\geqslant} 2\varepsilon(1 - q_{i,i} - \gamma^9) - n \cdot (1 - (1 - \gamma^9)) \\
&= 2\varepsilon(1 - q_{i,i}) - (2\varepsilon\gamma^9 + n\gamma^9) \geqslant 2\varepsilon(1 - q_{i,i}) - \gamma^8,
\end{aligned} \tag{79}$$

as required. ⊡

B.10 Proof of Claim 3.10

Proof of Claim 3.10. Fix $R \subseteq S$ and write $T_Y := \min\{t : Y_t \in R\}$ and $T_X := \{t : |x_t| \in R\}$. First, we will show by induction on t that, for any $t \geqslant 0$ and $i \in S$,

$$\mathbb{P}[T_Y \leqslant t \mid Y_0 = i] \geqslant \mathbb{P}[T_X \leqslant t \mid |x_0| = i]. \tag{80}$$

Indeed, (80) is true for $t = 0$, as in that case both sides are equal to $\mathbb{1}(i \in R)$. Additionally, if we assume that (80) holds for some fixed

$t \geqslant 0$, then if $i \in V \cup W$,

$$\begin{aligned}\mathbb{P}[T_Y \leqslant t+1 \mid Y_0 = i] &= \sum_{j\in S} P_Y(i,j)\cdot \mathbb{P}[T_Y \leqslant t \mid Y_0 = j]\\ &\overset{(12)}{=} \sum_{j\in S} q_{i,j}\cdot \mathbb{P}[T_Y \leqslant t \mid Y_0 = j]\\ &\overset{(80)}{\geqslant} \sum_{j\in S} q_{i,j}\cdot \mathbb{P}[T_X \leqslant t \mid |x_0| = j]\\ &= \mathbb{P}[T_X \leqslant t+1 \mid |x_0| = i],\end{aligned}$$

whereas if $i \notin V \cup W$,

$$\begin{aligned}&\mathbb{P}[T_Y \leqslant t+1 \mid Y_0 = i] = \sum_{j\in S} P_Y(i,j)\cdot \mathbb{P}[T_Y \leqslant t \mid Y_0 = j]\\ &\overset{(12)}{=} \sum_{j\in S\setminus\{i\}} \left(\tfrac{q_{i,j}}{1-q_{i,i}}\right)\cdot \mathbb{P}[T_Y \leqslant t \mid Y_0 = j]\\ &\overset{(80)}{\geqslant} \tfrac{1}{1-q_{i,i}} \sum_{j\in S\setminus\{i\}} q_{i,j}\cdot \mathbb{P}[T_X \leqslant t \mid |x_0| = j]\\ &= \tfrac{1}{1-q_{i,i}}(\mathbb{P}[T_X \leqslant t+1 \mid |x_0| = i] - q_{i,i}\cdot \mathbb{P}[T_X \leqslant t \mid |x_0| = i])\\ &\geqslant \tfrac{1}{1-q_{i,i}}(\mathbb{P}[T_X \leqslant t+1 \mid |x_0| = i] - q_{i,i}\cdot \mathbb{P}[T_X \leqslant t+1 \mid |x_0| = i])\\ &= \mathbb{P}[T_X \leqslant t+1 \mid |x_0| = i].\end{aligned}$$

Therefore, (80) holds for all $t \geqslant 0$ and $i \in S$. But then, because the $(Y_t)_{t=0}^{\infty}$ and $(|x_t|)_{t=0}^{\infty}$ have the same initial distribution, we have for any $t \geqslant 0$ that

$$\begin{aligned}\mathbb{P}[T_Y \leqslant t] &= \sum_{i\in S} \mathbb{P}[T_Y \leqslant t \mid Y_0 = i]\cdot \mathbb{P}(Y_0 = i)\\ &\overset{(80)}{\geqslant} \sum_{i\in S} \mathbb{P}[T_X \leqslant t \mid |x_0| = i]\cdot \mathbb{P}(|x_0| = i)\\ &= \mathbb{P}[T_X \leqslant t],\end{aligned}$$

and so $T_Y \preceq T_X$, as required. ⊡

B.11 Proof of Claim 3.11

Proof of Claim 3.11. Let

$$\Delta_t = h(Z_{t+1} - Z_t) = \min\{Z_{t+1} - Z_t, C(\varepsilon)\}.$$

In order to apply Theorem 1.7, we will verify that conditions **B1**-**B3** hold for Δ_t with $a = (\frac{1}{2} - 2\varepsilon - 2\eta)n$, $b = (\frac{1}{2} - 2\varepsilon - \eta)n$, $\ell = b - a = \eta n$, $\kappa = C(\varepsilon)/\varepsilon$, and $r = n^{1/3}$.

Before proceeding to the verification of **B1**-**B3**, we first remark that, from the definition of V given by (11),

$$1 - q_{i,i} > e^{-\varepsilon n^{1/3}} \text{ whenever } i \in [0, (\tfrac{1}{2} + 2\varepsilon + 2\eta)n) \setminus V. \quad (81)$$

To see that **B1** holds, observe that if $a < Z_t < b$ then $Z_t = n - Y_t$ and $Y_t = i$ for some $i \in ((\frac{1}{2} + 2\varepsilon + \eta)n, (\frac{1}{2} + 2\varepsilon + 2\eta)n) \setminus V$. In this case, if $Y_{t+1} \in V$ then

$$\Delta_t = \min\{n - (n - i), C(\varepsilon)\} = C(\varepsilon),$$

whereas if $Y_{t+1} \notin V$ then

$$\begin{aligned}\Delta_t &= \min\{i - \max\{Y_{t+1}, (\tfrac{1}{2} + 2\varepsilon)n\}, C(\varepsilon)\}\\ &= \min\{\min\{i - Y_{t+1}, i - (\tfrac{1}{2} + 2\varepsilon)n\}, C(\varepsilon)\}\\ &\geqslant \min\{\min\{i - Y_{t+1}, \eta n\}, C(\varepsilon)\}\\ &= \min\{i - Y_{t+1}, \min\{\eta n, C(\varepsilon)\}\} = \min\{i - Y_{t+1}, C(\varepsilon)\}.\end{aligned}$$

Thus we have $\min\{i - Y_{t+1}, C(\varepsilon)\} \leqslant \Delta_t \leqslant C(\varepsilon)$, and hence if $a < Z_t = n - i < b$,

$$\begin{aligned}&\mathbb{E}[\Delta_t \cdot \mathbb{1}(\Delta_t \leqslant \kappa\varepsilon) \mid \mathscr{G}_t] = \mathbb{E}[\Delta_t \mid \mathscr{G}_t]\\ &\geqslant \sum_{k=i-n}^{i} P_Y(i, i-k)\cdot \min\{k, C(\varepsilon)\}\\ &= (1 - q_{i,i})^{-1}\cdot \sum_{k=i-n}^{i} h(k)\cdot q_{i,i-k}\\ &\overset{\text{Claim 3.9}}{\geqslant} 2\varepsilon - \frac{\gamma^8}{1 - q_{i,i}} \overset{(81)}{>} 2\varepsilon - \gamma^8 e^{\varepsilon n^{1/3}} = 2\varepsilon - \gamma^7 \geqslant \varepsilon.\end{aligned}$$

Therefore, **B1** holds.

To see that **B2** holds, first note that if $t \geqslant T_Y^{\text{hit}}(V)$ then $\mathbb{P}(\Delta_t \leqslant -k \mid \mathscr{G}_t) = 0$ holds for any $k > 0$ by definition. So suppose instead that $t < T_Y^{\text{hit}}(V)$ and $a < Z_t < n$. Set $i = Y_t$ and $i' = \max\{i, (\frac{1}{2} + 2\varepsilon)n\}$. Because $(\frac{1}{2} - 2\varepsilon - 2\eta)n < Z_t < n$, it must hold that $i \in (0, (\frac{1}{2} + 2\varepsilon + 2\eta)n) \setminus V$ and $Z_t = n - i'$. Therefore, we have for any $j \in \mathbb{N}$ that

$$\begin{aligned}\mathbb{P}(\Delta_t \leqslant -jr \mid \mathscr{G}_t) &= \sum_{k\geqslant jr} \mathbb{P}(Z_{t+1} - Z_t = -k \mid \mathscr{G}_t)\\ &= \sum_{k\geqslant jr} \mathbb{P}(Z_{t+1} = n - i' - k \mid \mathscr{G}_t)\\ &= \sum_{k\geqslant jr} \mathbb{P}(Y_{t+1} = i' + k \wedge i' + k \notin V \mid \mathscr{G}_t)\\ &\leqslant \sum_{k\geqslant jr} \mathbb{P}(Y_{t+1} = i' + k \mid \mathscr{G}_t) = \sum_{k\geqslant jr} P_Y(i, i' + k)\\ &\overset{(12)}{=} \sum_{k=jr}^{n-i'} \frac{q_{i,i'+k}}{1 - q_{i,i}} \overset{\text{Claim 3.8}}{\leqslant} \frac{e^{\varepsilon(n^{1/3} - 8jr)}}{1 - q_{i,i}}\\ &\overset{(81)}{\leqslant} e^{\varepsilon(2n^{1/3} - 8jr)} = e^{\varepsilon r(2 - 8j)} \leqslant e^{-j},\end{aligned}$$

and so **B2** holds.

To verify **B3**, set

$$\begin{aligned}\lambda :&= \min\{1/(2r), \varepsilon/(17r^2), 1/(\kappa\varepsilon)\}\\ &= \min\{1/(2n^{1/3}), \varepsilon/(17n^{2/3}), 1/2C(\varepsilon)\} = \frac{\varepsilon}{17n^{2/3}}\end{aligned}$$

and observe that, because $\eta = \frac{\varepsilon \log(1/(1-\beta))}{200 \log n}$,

$$\begin{aligned}\lambda\ell &= \frac{\varepsilon\eta}{17} n^{1/3} = \frac{\varepsilon^2 \log(1/(1-\beta))}{3400 \log n} n^{1/3}\\ &\geqslant 2\ln\left(\frac{68n^{2/3}}{\varepsilon^2}\right) = 2\ln\left(\frac{4}{\lambda\varepsilon}\right).\end{aligned}$$

Therefore, by Theorem 1.7, we have for the first hitting time $T^* := \min\{t \geqslant 0 : Z_t \leqslant (\frac{1}{2} - 2\varepsilon - 2\eta)n\}$ that, if $Z_0 \geqslant b$,

$$\mathbb{P}[T^* \leqslant e^{\lambda\ell/4} \mid \mathscr{G}_0] \leqslant \overline{C}e^{-\lambda\ell/4}. \tag{82}$$

We also have

$$\begin{aligned}\mathbb{P}(Z_0 < b) &= \mathbb{P}(|x_0| > n - b \wedge |x_0| \notin V) \leqslant \mathbb{P}(|x_0| > n - b) \\ &= \mathbb{P}(|x_0| > (\tfrac{1}{2} + 2\varepsilon + \eta)n) \leqslant \mathbb{P}(x_0 \notin A) \overset{(6)}{\leqslant} \gamma^{10}. \end{aligned} \tag{83}$$

and hence,

$$\begin{aligned}\mathbb{P}[T^* \leqslant e^{n^{1/8}}] \leqslant \mathbb{P}[T^* \leqslant e^{\lambda\ell/4}] &\overset{(82)}{\leqslant} \overline{C}e^{-\lambda\ell/4} + \mathbb{P}(Z_0 < b) \\ &\overset{(83)}{\leqslant} \overline{C}e^{-\lambda\ell/4} + \gamma^{10} \leqslant \tfrac{1}{2}e^{-n^{1/8}},\end{aligned}$$

as required. ⊡

C A NOTE ON NASH EQUILIBRIA

The entropy of a probability distribution p over a finite set $\mathcal{X}$ is defined to be

$$H(p) = -\sum_{x \in \mathcal{X}} p(x) \log p(x).$$

When formulating Definition 1.2, we claimed the existence of a unique Nash equilibrium of maximal entropy. Here we quickly justify this claim for symmetric zero-sum games. In the subsequent proof, the topology on $\mathcal{P}(\mathcal{X})$ is that induced by the standard topology on $\mathbb{R}^{|\mathcal{X}|}$, when identifying $\mathcal{P}(\mathcal{X})$ with the probability simplex $\Delta(\mathcal{X}) := \{z \in [0,1]^{|\mathcal{X}|} : \sum_{i=1}^{|\mathcal{X}|} z_i = 1\}$.

Proposition C.1. *Let $f : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ be an antisymmetric function on a finite set $\mathcal{X}$. Then there is a unique Nash equilibrium for f of maximal entropy.*

Proof. The function $H : \mathcal{P}(\mathcal{X}) \to \mathbb{R}$ is strictly concave, and so attains a unique maximum on every subset of $\mathcal{P}(\mathcal{X})$ that is non-empty, compact, and convex. Therefore, it is enough to show that the set of Nash equilibria for f has these three properties.

Recalling the discussion surrounding Definition 1.2, note that the following conditions are equivalent for $p \in \mathcal{P}(\mathcal{X})$.

E1 p is a Nash equilibrium for f.
E2 $\min_{q \in \mathcal{P}(\mathcal{X})} \sum_{x,y \in \mathcal{X}} p(x)q(y)f(x,y) = 0$.
E3 $\sum_{x,y \in \mathcal{X}} p(x)q(y)f(x,y) \geqslant 0$ for every $q \in \mathcal{P}(\mathcal{X})$.

Let $N \subseteq \mathcal{P}(\mathcal{X})$ be the set of Nash equilibria for f. By Nash's theorem (see [23, Theorem 10.4]), N is non-empty. Let $h : \mathcal{P}(\mathcal{X}) \to \mathbb{R}$ be the continuous function

$$h(p) = \min_{q \in \mathcal{P}(\mathcal{X})} \sum_{x,y \in \mathcal{X}} p(x)q(y)f(x,y).$$

By **E2**, $N = h^{-1}(\{0\})$. Therefore, N is a closed subset of the compact topological space $\mathcal{P}(\mathcal{X})$, and hence N is compact. If $p_1, p_2 \in \mathcal{P}(\mathcal{X})$ and $s \in [0,1]$, then for any $q \in \mathcal{P}(\mathcal{X})$,

$$\begin{aligned}&\sum_{x,y \in \mathcal{X}} (sp_1 + (1-s)p_2)(x)q(y)f(x,y) \\ &\quad = s \sum_{x,y \in \mathcal{X}} p_1(x)q(y)f(x,y) + (1-s) \sum_{x,y \in \mathcal{X}} p_2(x)q(y)f(x,y) \\ &\quad \overset{\textbf{E3}}{\geqslant} s \cdot 0 + (1-s) \cdot 0 = 0.\end{aligned}$$

Therefore N is also convex, and hence the result holds. □