TO REEVALUATE OR NOT: A THEORETICAL STUDY OF NSGA-II UNDER NOISE

Anonymous authors
Paper under double-blind review

ABSTRACT

Multi-objective evolutionary algorithms (MOEAs) are widely used to optimize problems with multiple objectives. In real-world applications, however, solution evaluations may be corrupted by noise, for example, when using a simulator. A straightforward response is to reevaluate solutions. This raises a key question: is it more beneficial to reevaluate solutions in every generation, or retain their initial noisy fitness values throughout the search? We investigate this question theoretically for NSGA-II, one of the most widely used MOEAs. On the common benchmark LOTZ under one-bit prior noise, we show that under noise levels of $\omega(\sqrt{\log n/n})$ and $O(1/\log n)$, retaining initial fitness values yields polynomial expected runtime, whereas reevaluation requires superpolynomial expected runtime, with all other algorithm settings unchanged. The proven detrimental effect of reevaluation echoes known results for simple MOEAs. Further, we construct a new benchmark PairedCOCZ and prove that for noise levels below but sufficiently close to one, reevaluation is helpful: it leads to polynomial expected runtime, compared with at least exponential expected runtime without reevaluation. These results not only provide theoretical comparison of evaluation strategies for practical MOEAs under noise for the first time, but also challenge the prevailing belief that reevaluation is always detrimental. We complement the theory with empirical studies.

1 INTRODUCTION

Many optimization tasks involve objectives that cannot all be improved simultaneously. The aim is therefore to find a set of Pareto-optimal trade-offs, rather than a single best solution (Deb, 2001; Coello Coello et al., 2007). Multi-objective evolutionary algorithms (MOEAs) approach such a task by evolving a population of candidate solutions, allowing different parts of the Pareto front to be explored in one run. Algorithms such as the NSGA-II (Deb et al., 2002), MOEA/D (Zhang & Li, 2007) and SMS-EMOA (Beume et al., 2007) have become widely used solvers, with applications across engineering and other fields (Coello Coello & Lamont, 2004; Zhou et al., 2019).

The theoretical understanding of MOEAs has fallen behind their practical adoption, but runtime analysis has made increasing progress in recent years. Early studies focused on simple MOEAs such as SEMO and GSEMO (Giel, 2003; Laumanns et al., 2004). More recent work analyses the NSGA-II (Zheng & Doerr, 2023; Doerr & Qu, 2023b) and its crowding-distance, selection and tie-breaking mechanisms (Zheng & Doerr, 2022; Bian & Qian, 2022; Doerr et al., 2025a). Further results establish the benefits of stochastic population updates (Bian et al., 2025), external archives (Bian et al., 2024; Ren et al., 2026), and dynamic population sizes (Doerr et al., 2026c), sharpen population-dynamics bounds for GSEMO and SPEA2 (Doerr et al., 2025b; 2026b), and identify limitations of crowding or density estimation (Doerr et al., 2026a; Dang et al., 2026b). Analyses of the information used by selection also help explain why practical MOEAs can outperform simpler dominance-based algorithms (Dang et al., 2026a).

Most of these theories, however, assume exact objective evaluations. In practice, measurement errors and stochastic simulations can make the same solution receive different fitness values on different evaluations (Beyer, 2000; Goh & Tan, 2006), as in simulation-driven engineering design (Li et al., 2005; Sun et al., 2018) and logistics with uncertain demand (Tirkolaee et al., 2020). This can mislead selection: a useful trade-off solution can appear dominated, while an inferior solution can appear

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Table 1: Proved noise regimes and expected-runtime bounds for NSGA-II under one-bit prior noise with noise probability p. Bounds hold under the assumptions of Theorems 1–4. On PairedCOCZ, r is the fixed number of bit pairs and $\mu = \Theta(\log n)$.

Problem	Algorithm	Reeval.	Polynomial runtime		Superpolynomial runtime	
			p upper bound	$\mathbb{E}[T]$	p lower bound	$\mathbb{E}[T]$
LOTZ	NSGA-II	No	$O(1/\log n)$	$O(n^4)$	—	
		Yes	$O(1/n^2)$	$O(n^2)$	$\omega(\sqrt{\log n/n})$	$2^{\Omega(p^2 n)}$
	NSGA-II: noise regime		No reevaluation		Reevaluation	
PairedCOCZ	$p = 1-q,\ 0 < q \le 2^{-2n}$		$2^{n-o(n)} \le \mathbb{E}[T] < \infty$		$O(n^2 \log^{4r-3} n)$	
	$p = 1$		$\mathbb{E}[T] = \infty$		$O(n^2 \log^{4r-3} n)$	

worth retaining. For algorithms that keep individuals across generations, this raises a basic choice: reevaluate their fitness for each selection step, or retain the values obtained when each individual was created. Reevaluation offers a chance to correct past errors, but also exposes previously discovered solutions to fresh noise.

Research on runtime analyses of MOEAs for noisy problems remains scarce. Under one-bit prior noise, Dinot et al. (2023) analysed SEMO on OneMinMax and proved an expected $O(n^2 \log n)$ iteration bound for $p \le \alpha/n$ without reevaluation and for $p \le \beta/n^2$ with population reevaluation, where $\alpha, \beta > 0$ are sufficiently small constants. With reevaluation, the expected runtime is superpolynomial when $\omega(\log n/n^2) = p \le \lambda/n$ for a fixed constant $\lambda \in (0, 1)$, implying that reevaluation is detrimental. A separate line of work asks whether the choice of algorithm improves robustness to noise. Dang et al. (2023) showed that NSGA-II can outperform GSEMO under additive Bernoulli posterior noise. Note that their comparison concerns different algorithms rather than the decision to reevaluate within a single algorithm. In contrast, in this paper, we investigate whether the detrimental effect of reevaluation observed for SEMO persists in the widely used practical MOEA, NSGA-II.

A number of studies have considered runtime analyses in noisy single-objective evolutionary optimization, examining the robustness of various strategies to noisy settings, including reevaluation and selection (Qian et al., 2018b; Sudholt, 2021; Antipov & Doerr, 2025), sampling (Qian et al., 2018a; Doerr & Sutton, 2019), and increasing the population size (Qian et al., 2021; Antipov et al., 2024). For example, Antipov & Doerr (2025) showed that retaining creation-time fitness values (i.e., not using reevaluation) can greatly improve noise tolerance on the LeadingOnes problem (Droste et al., 2002). However, these single-objective results do not settle the multi-objective question: reaching one optimum in single-objective optimization differs from acquiring and simultaneously retaining representatives of many incomparable Pareto-optimal solutions in multi-objective optimization.

In this paper, we compare reevaluation with no reevaluation for NSGA-II under one-bit prior noise. Within each comparison, the population size and variation operators remain unchanged. Our analysis yields two contrasting results:

- We establish a polynomial-versus-superpolynomial runtime difference between evaluation strategies for the widely used NSGA-II. On LOTZ, with population size $\mu = 2n + 2$ and fixed crossover probability $p_c \in (0, 1)$, no reevaluation leads to $O(n^4)$ expected runtime when the noise level (i.e., the probability of noise occurrence) $p \le \delta\, p_c(1 - p_c)/\log n$, for a sufficiently small constant $\delta > 0$. Within this range, when $p = \omega(\sqrt{\log n/n})$, the expected number of generations required with reevaluation is $\exp(\Omega(p^2 n))$ from uniform initialization (Corollary 1). Thus, the detrimental effect previously found for SEMO persists despite the NSGA-II's multiple copies and crowding-distance selection.
- We introduce a new benchmark PairedCOCZ to show that reevaluation can instead be beneficial in some cases. For a constant-size Pareto front and a population of logarithmic size with a sufficiently large constant factor, reevaluation leads to $n^2\,\mathrm{polylog}(n)$ expected runtime when the noise level $p = 1 - q$ with $0 < q \le 2^{-2n}$. Without reevaluation, the expected runtime from uniform initialization is finite but at least $2^{n-o(n)}$. At $p = 1$, no

reevaluation has infinite expected runtime unless the initial population already covers the entire Pareto front. In contrast, reevaluation still has polynomial expected runtime. These results show that the previous negative findings do not imply that avoiding reevaluation is always preferable.

Table 1 summarizes both comparisons. Runtime T counts the number of generations until the population contains a true Pareto-optimal representative of every front value. On LOTZ, repeated noisy evaluations can cause useful solutions to be discarded, whereas retaining their earlier fitness values helps preserve them. On PairedCOCZ, retained incorrect fitness values can keep inferior solutions in the population, whereas reevaluation allows them to be replaced. We complement the theoretical results with experiments on LOTZ and PairedCOCZ.

2 PRELIMINARIES

2.1 LOTZ AND PARETO DOMINANCE

Let us assume that objectives are maximized. For fitness vectors, $(a, b) \succeq (a', b')$ means $a \geq a'$ and $b \geq b'$. Strict dominance $\succ$ additionally requires at least one of these inequalities to be strict.

For $\boldsymbol{x} = (x_1, \ldots, x_n) \in \{0,1\}^n$, the LeadingOnesTrailingZeroes (LOTZ) problem, one of the most widely used benchmarks in theoretical analysis of MOEAs (Laumanns et al., 2004; Qian et al., 2013; Dang et al., 2023), has true fitness vector $g(\boldsymbol{x}) = (\mathrm{LO}(\boldsymbol{x}), \mathrm{TZ}(\boldsymbol{x}))$, where

$$\mathrm{LO}(\boldsymbol{x}) = \sum_{i=1}^{n} \prod_{j=1}^{i} x_j, \qquad \mathrm{TZ}(\boldsymbol{x}) = \sum_{i=1}^{n} \prod_{j=i}^{n} (1 - x_j).$$

These count the number of leading ones and trailing zeros, respectively. The Pareto set is

$$F = \{\pi_k = 1^k 0^{n-k} : 0 \leq k \leq n\}, \qquad g(\pi_k) = (k, n-k).$$

We call k a *slot*. Obtaining the entire Pareto front requires all $n+1$ strings π_k to coexist in the current population.

2.2 ONE-BIT PRIOR NOISE AND THE TWO EVALUATION STRATEGIES

Under one-bit prior noise of level p, each evaluation of a solution $\boldsymbol{y}$ returns the noisy fitness vector

$$\tilde{g}(\boldsymbol{y}) = \begin{cases} g(\boldsymbol{y}), & \text{with probability } 1-p, \\ g(\boldsymbol{y} \oplus e_I), & \text{with probability } p, \end{cases}$$

where I is uniform on $\{1, \ldots, n\}$, e_I is the I-th unit vector, and $\oplus$ denotes bitwise XOR, so $\boldsymbol{y} \oplus e_I$ flips exactly bit I of $\boldsymbol{y}$. Both objectives are evaluated on the same string; the stored solution $\boldsymbol{y}$ remains unchanged. Noise draws are independent across evaluations. An evaluation is called *corrupted* when the bit flip occurs, even if the objective vector is unchanged.

Under the *reevaluation strategy*, all 2μ individuals in the combined parent and offspring population are evaluated again before each survival-selection step, where μ is the population size. Under the *no-reevaluation strategy*, each individual is evaluated only when it is created; this fitness value is kept unchanged and used whenever that individual is compared with others.

2.3 THE NSGA-II

We analyze NSGA-II (Deb et al., 2002) with binary tournament selection, one-point crossover and bit-wise mutation, following Bian & Qian (2022); Doerr & Qu (2023a). On LOTZ, we use $\mu = 2n + 2$, the minimum sufficient size in their analysis. Algorithm 1 presents the algorithm procedure, including both evaluation strategies and survival selection.

Each tournament samples two individuals uniformly with replacement. The one with the lower non-domination rank is selected; if their ranks are equal, the one with the larger crowding distance is selected. If both are equal, each individual is selected with probability $1/2$. Crossover cuts

Algorithm 1 NSGA-II under one-bit prior noise

Input: objectives g, even population size μ, noise level p, crossover rate p_c, strategy S (reeval or no-reeval).
Output: μ solutions in $\{0,1\}^n$, allowing duplicates.

1: $P_0 \leftarrow \mu$ independent uniform strings in $\{0,1\}^n$; $t \leftarrow 0$;
2: evaluate each individual in P_0 independently and store its fitness;
3: compute ranks and crowding on P_0 using lines 14–15;
4: **while** not terminated **do**
5: $\quad Q_t \leftarrow \emptyset$;
6: $\quad$ **for** $i = 1$ **to** $\mu/2$ **do**
7: $\quad\quad$ select $\boldsymbol{x}$ and $\boldsymbol{y}$ from P_t by two independent binary tournaments;
8: $\quad\quad$ apply one-point crossover to $\boldsymbol{x}$ and $\boldsymbol{y}$ to obtain $\boldsymbol{x}'$ and $\boldsymbol{y}'$, with probability p_c;
9: $\quad\quad$ apply bit-wise mutation on $\boldsymbol{x}'$ and $\boldsymbol{y}'$ to generate $\boldsymbol{x}''$ and $\boldsymbol{y}''$, respectively;
10: $\quad\quad$ add $\boldsymbol{x}''$ and $\boldsymbol{y}''$ to Q_t
11: $\quad$ **end for**
12: $\quad R_t \leftarrow P_t \cup Q_t$ (parents first);
13: $\quad$ evaluate R_t if $S =$ reeval, else only Q_t; update each evaluated individual's stored fitness;
14: $\quad$ form non-dominated layers $L_1, L_2, \ldots$ from stored fitness; keep order and store ranks;
15: $\quad$ compute normalized crowding distance in each layer;
16: $\quad P_{t+1} \leftarrow \emptyset$; $i \leftarrow 1$;
17: $\quad$ **while** $|P_{t+1}| + |L_i| < \mu$ **do**
18: $\quad\quad P_{t+1} \leftarrow P_{t+1} \cup L_i$; $i \leftarrow i + 1$
19: $\quad$ **end while**
20: $\quad$ add the $\mu - |P_{t+1}|$ highest-crowding copies of L_i to P_{t+1} (ties uniform);
21: $\quad$ retain survivors' fitness values, ranks and crowding distances; $t \leftarrow t + 1$
22: **end while**
23: **return** P_t

are uniform in $\{1, \ldots, n-1\}$; without crossover, $\boldsymbol{x}' = \boldsymbol{x}$ and $\boldsymbol{y}' = \boldsymbol{y}$. Mutation flips each bit independently with probability $1/n$. Each evaluation uses independent noise, and $\cup$ retains repeated individuals.

Parents are placed before offspring, and this order is preserved within each non-dominated layer. Stable ascending sorts by both objectives preserve the order of individuals with identical fitness vectors, as in Bian & Qian (2022). We use parents first and no shuffling. For each objective, the first and last individuals receive infinite crowding distance. For an interior individual, we divide its neighbours' objective-value difference by the layer's largest value minus its smallest; a zero denominator contributes zero. Its crowding distance is the sum of these normalized differences for both objectives. Survivors keep their distances for the next generation's tournaments.

On LOTZ, $T = \min\{t : F \subseteq P_t\}$ counts the number of generations needed to obtain F. Initialization is uniform unless stated otherwise; every initial population is evaluated and sorted.

3 Theoretical Results on LOTZ

Throughout this section, Algorithm 1 runs on LOTZ with $\mu = 2n + 2$. Constants δ, C, κ are positive and absolute; asymptotics are for $n \to \infty$. Let $\mathcal{F}_t$ be the history before generation t. The appendices give the proofs and explicit choices of constants.

3.1 No Reevaluation

Theorem 1 (no reevaluation, upper bound)**.** *Consider NSGA-II (Algorithm 1), with population size $\mu = 2n + 2$ without reevaluation on LOTZ under one-bit prior noise of level p. Let T count the number of generations needed to obtain the entire Pareto front. For a sufficiently small absolute $\delta > 0$ and every constant crossover probability $p_c \in (0, 1)$, when $p \le \delta\, p_c(1 - p_c)/\log n$, we have $\mathbb{E}[T] = O(n^4)$ from an arbitrary initial population.*

Proof sketch. An individual storing $(k, n-k)$ is either π_k or a Hamming neighbour of π_k whose differing bit was flipped during evaluation. Thus an incorrect representative of position k differs from π_k in exactly one bit.

Stored front values are never lost: they are undominated, so individuals storing them have rank 1. Stable sorting gives at most two individuals with positive crowding distance per value, and all of them survive. Once every front value has two representatives, the population thereafter has exactly two per value, although their identities may change.

Before this happens, layer 1 may contain many off-front values, so a desired parent need not win a tournament often. Drawing it in both samples guarantees its selection. Skipping crossover, mutating at most two prescribed bits and evaluating without corruption then gives progress with probability $\Omega(1/(\mu n^2))$ per generation. At most $2n$ such steps find all front values, in $O(n^4)$ expected generations. Once a value has two representatives, it keeps at least two; the same argument with at most one mutation bit yields two per value in another $O(n^3)$ expected generations.

With exactly two individuals per front value, all have rank 1 and non-extreme individuals have equal crowding, so each wins a tournament with probability $\Theta(1/\mu)$. Let B_k indicate that neither individual storing $(k, n-k)$ equals π_k. Mutation creates a correctly evaluated individual π_k with probability $\Omega((1-p_c)/n)$ per generation, while a new incorrect representative of k appears with probability at most $\beta = O(p/n)$. If a correct offspring arrives and no incorrect representative does, the last member of the equal-value block is correct, because offspring follow parents. This individual survives selection. Thus, if a correct offspring arrives with probability r_k, the probability of repairing the block is at least $r_k - \beta \geq r_k/3$ for sufficiently small p. A block containing a correct representative cannot lose all correct representatives without a new incorrect one. An initial $O(n \log n)$ stage gives every block a correct representative at some generation boundary, not necessarily the same one, with high probability. This alone does not ensure simultaneous coverage, since correct representatives can be lost again. A conditional moment bound then shows that, at every generation of a second stage, most positions in each sufficiently long prefix and suffix of geometrically increasing lengths have correct representatives. Crossover uses these individuals to fill gaps: crossing π_i, π_j with $i < k < j$ and cut k produces π_k. At distance $d = \min\{k, n-k\} \geq C \log n$ from the endpoints, for sufficiently large C, this adds $\Omega(p_c d/n)$ to r_k; mutation handles the remaining positions.

We now consider all positions at the same final generation. For each k, we bound the probability that $B_k = 1$ then and that enough correct parents have remained available throughout the second stage. This probability is at most a term that decreases exponentially with the stage length plus $O(\beta/r_k)$. The interior terms are bounded using $\sum_{d=1}^{n} 1/d = O(\log n)$, while only $O(\log n)$ positions near the endpoints need the mutation bound. Together, these give $\sum_k \beta/r_k = O(p \log n/(p_c(1-p_c)))$. A union bound over positions, adding separately the probability of insufficient correct parents, shows that choosing δ small enough makes an $O(n \log n)$ epoch end with all $B_k = 0$ with constant probability. No independence between positions is required. This explains the sufficient noise range $p = O(1/\log n)$ for fixed p_c. These estimates hold whenever there are exactly two individuals per front value, so only constantly many epochs are needed in expectation. The full proof is in Appendix C. □

Preserving a front value does not guarantee preserving its true Pareto-optimal solution. Correctly evaluated offspring must also replace individuals storing incorrect fitness. Crossover provides more opportunities away from the endpoints, where many pairs of Pareto-optimal parents can produce a missing solution; mutation supplies progress near the endpoints. Comparing how often errors appear with how often correct solutions replace them shows that progress is frequent enough in the stated noise range.

3.2 REEVALUATION

Theorem 2 (reevaluation, upper bound)**.** *Consider NSGA-II (Algorithm 1), with population size $\mu = 2n + 2$ and reevaluation on LOTZ under one-bit prior noise of level p. Let T count the number of generations needed to obtain the entire Pareto front. For every constant crossover probability $p_c \in (0, 1)$, there is a constant $c' > 0$ such that, when $p \leq c'/n^2$, we have $\mathbb{E}[T] = O(n^2)$ from an arbitrary initial population.*

Proof sketch. Set $\Psi_t = \Phi_1(P_t) + 2Z_t + \beta_t \in [0, 4n+3]$, where Z_t counts missing Pareto strings, $\Phi_1 = 2n - \max \mathrm{LO} - \max \mathrm{TZ}$ measures distance to the two extremes, and β_t indicates whether the last survival step used a corrupted evaluation, with $\beta_0 = 1$. Each lost Pareto string requires a distinct corrupted evaluation, so the expected increase in Z_t is $O(\mu p)$. A decrease in the best LO or TZ value needs a separate bound. It requires either corruption of a lexicographically maximal individual attaining that value or a flip of one of at most two specific bits in another evaluation. These events have probability $O(p + \mu p/n)$; even a decrease of size n therefore contributes only $O(\mu p)$ in expectation. The expected increase in Ψ is thus $O(\mu p)$.

A survival step without corrupted evaluations restores the correct ranks and crowding distances used by subsequent tournaments. When $\beta_t = 0$, mutation can improve the best LO or TZ value while an extreme is missing; once both extremes are present, crossover with a suitable Pareto parent can fill a gap. These events decrease Ψ with probability $\Omega(\min\{1 - p_c, p_c\}/n)$, including a survival step without corruption, whose probability is at least $1/2$. When $\beta_t = 1$, that step instead resets β with constant probability. Choosing c' sufficiently small gives net drift $\Omega(1/n)$; additive drift yields $\mathbb{E}[T] = O(n^2)$. See Appendix D. □

Theorem 3 (reevaluation, lower bound)**.** *Consider NSGA-II (Algorithm 1), with population size $\mu = 2n+2$ and reevaluation on LOTZ under one-bit prior noise of level p, starting from a uniformly initialized population. Let T count the number of generations needed to obtain the entire Pareto front. There is an absolute constant $\kappa > 0$ such that, for every constant crossover probability $p_c \in [0, 1)$, there are constants $p_{\max}, C_2 > 0$ such that, for sufficiently large n, when $C_2/n \le p^2 \le p_{\max}^2$, we have*

$$\mathrm{P}\big[T > \exp(\kappa p^2 n)\big] \;\ge\; \tfrac{1}{2}, \qquad \textit{hence} \qquad \mathbb{E}[T] \;\ge\; \tfrac{1}{2}\big\lfloor \exp(\kappa p^2 n)\big\rfloor.$$

Proof sketch. Suppose $1 \le Z_t \le n/100$ and at most $n/100$ individuals are non-Pareto strings. An occupancy bound over the $\mu/2$ independent variation rounds shows that, with probability $1 - e^{-\Omega(n)}$, $\Theta(n)$ present Pareto strings have at most two copies and receive no new copy. Conditional on the offspring, all copies of each such string receive unfavorable fitness values with probability at least $p^2/625$, independently for different strings. With high probability, correctly evaluated front points dominate the affected copies and layer 1 fills the population, so none of those copies survives. Accounting for exceptions, the number Λ_t of losses satisfies $\mathbb{E}[e^{-2\Lambda_t} \mid \mathcal{F}_t] \le e^{-c_L n p^2}$ for a constant $c_L > 0$. Since the number Γ_t of gains is at most Z_t, this gives $\mathbb{E}[e^{\Gamma_t - \Lambda_t} \mid \mathcal{F}_t] \le e^{Z_t - c_L n p^2/2}$. Consequently, e^{-Z_t} is a supermartingale until Z_t leaves a sufficiently narrow band $1 \le Z_t \le W = \Theta(np^2)$ or the bound on non-Pareto individuals fails.

It remains to justify the bound on non-Pareto individuals when the process enters this band. Let s_t count Pareto strings with only one copy and J_t count non-Pareto individuals. Cloning such singletons reduces s_t, apart from new singletons caused by filling gaps and by noise, while $J_t \le 2Z_t + s_t$. After a constant number of generations this gives the required bound on J_t; occupancy bounds prevent Z_t from reaching W too soon. The bound then continues to hold while the population is close to coverage, except on rare concentration failures. The loss estimate also makes a jump from above W to below $b = \lfloor W/2 \rfloor$ exponentially unlikely. Every remaining entry starts with $Z_t \ge b$, so optional stopping bounds its probability of being zero at stopping by e^{-b}. Uniform initialization starts above the band with high probability. A union bound over $e^{\kappa n p^2}$ generations, including all exceptional events, proves the claim. See Appendix E. □

Multiple copies do not guarantee survival under reevaluation. At $\mu = 2n + 2$, many Pareto-optimal solutions still have at most two copies and receive no new copy. Near completion, noise causes $\Omega(np^2)$ losses in expectation. Crowding distance cannot save dominated individuals when the first layer fills the population. Since offspring are generated before reevaluation, newly generated solutions face the same risk. All front values must nevertheless have true representatives simultaneously.

3.3 Comparing Evaluation Strategies

The preceding bounds vary only the evaluation strategy, keeping population size, selection and variation operators unchanged.

Corollary 1. *Consider NSGA-II (Algorithm 1) on n-bit LOTZ with population size $\mu = 2n+2$ under one-bit prior noise of level p. Fix a constant crossover probability $p_c \in (0, 1)$, and let T count the*

number of generations needed to obtain the entire Pareto front. Let $\delta > 0$ be as in Theorem 1. Without reevaluation, we have $\mathbb{E}[T] = O(n^4)$ uniformly for $p \le \delta\, p_c(1-p_c)/\log n$. With reevaluation and uniform initialization, the expected runtime is superpolynomial when $p = \omega(\sqrt{\log n/n})$ and $p \le p_{\max}(p_c)$. For any fixed $K > 0$, when $\sqrt{K \log n/n} \le p \le p_{\max}(p_c)$, we have $\mathbb{E}[T] \ge n^{\Omega(K)}$ for sufficiently large n. Thus, from uniform initialization, switching from no reevaluation to reevaluation increases the expected runtime from polynomial to superpolynomial, with all other settings unchanged, when $p = \omega(\sqrt{\log n/n})$ and $p \le \delta\, p_c(1-p_c)/\log n$.

Proof. The first statement is Theorem 1. For the second, $K \log n \ge C_2$ for large n, and $p^2 n \ge K \log n$ turns Theorem 3 into $\mathbb{E}[T] \ge \frac{1}{2}\lfloor n^{\kappa K} \rfloor = n^{\Omega(K)}$. □

Corollary 1 shows that, under the stated noise conditions, reevaluation alone can increase NSGA-II's expected runtime on LOTZ from polynomial to superpolynomial. Multiple copies and crowding-distance selection do not prevent repeated noisy evaluations from removing Pareto-optimal solutions. This runtime separation disappears without noise. At $p = 0$, both strategies have the same runtime distribution in generations, so Theorem 2 gives $O(n^2)$ for both. The $O(n^4)$ bound in Theorem 1 is a looser bound that also applies throughout the stated noise range, not evidence of slower search at $p = 0$.

The bounds give sufficient noise conditions, not an exact point at which runtime becomes superpolynomial. They leave open the gap between the reevaluation bounds, the runtime without reevaluation at constant noise, and the effects of larger populations or different tie-breaking rules. Moreover, the small sufficient $p_{\max}$ in Appendix E means that the lower bound is an asymptotic result, not a quantitative prediction for the sizes used in our experiments.

4 When Reevaluation Wins: PairedCOCZ

The LOTZ comparison does not mean that avoiding reevaluation is always better. An incorrect fitness value may be retained indefinitely, and on some problems it can prevent true optima from entering the population. We introduce PairedCOCZ, a benchmark inspired by COCZ (Laumanns et al., 2004), to show this behavior. Within this comparison, population size and variation operators are again identical for the two evaluation policies.

PairedCOCZ function. Let $r \ge 1$ and $n > 2r$. Write $x = (w, u)$ with $w \in \{0,1\}^{2r}$, read as r pairs of bits, and $u \in \{0,1\}^{n-2r}$. Let $c_{11}(x)$ and $c_{00}(x)$ be the numbers of pairs equal to 11 and to 00, and let the *defect* $d(x)$ be the number of pairs equal to 01 or 10 plus the number of zeros of u. Then

$$\textsc{PairedCOCZ}(x) \;:=\; \big(c_{11}(x) - (r+1)\,d(x),\;\; c_{00}(x) - (r+1)\,d(x)\big),$$

to be maximized. The Pareto set $\{00,11\}^r \times \{1^{n-2r}\}$ contains 2^r strings, while the Pareto front $\{(j, r-j) : j = 0, \dots, r\}$ has $M = r+1$ values. We call the optimal strings with value $(j, r-j)$ class j. For fixed r, the front size is constant as n grows. Every one-bit neighbour of an optimum is nonoptimal, so one-bit noise always evaluates an optimum below the true front. Appendix F.1 gives examples and explains the function's properties.

Theorem 4. *Consider NSGA-II (Algorithm 1), with an even population size μ on* PairedCOCZ *under one-bit prior noise of level p. Fix the number of bit pairs $r \ge 1$, crossover probability $p_c \in [0,1)$ and $\delta \in (0, 2/3)$ as constants. Let T count the number of generations until the population contains a true optimum for every one of the $M = r+1$ front values. For sufficiently large $C = C(r, p_c)$ and sufficiently large n, when $C \ln n \le \mu \le n^{2/3-\delta}$ and $p = 1$, we have:*

- *(a) with reevaluation, $\mathbb{E}[T] = O(n^2 \mu^{4M-7})$ from every initial population; for $\mu = \Theta(\log n)$, this is $n^2 \operatorname{polylog}(n)$;*
- *(b) without reevaluation, $\mathrm{P}[T = \infty] \ge n^{-(n+1)\mu} > 0$ from every initial population with no true optimum for at least one front value, so $\mathbb{E}[T] = \infty$.*

For $p = 1 - q$ with $0 < q \le 2^{-2n}$, (a) still holds, whereas without reevaluation and from uniform initialization, $2^{n-o(n)} \le \mathbb{E}[T] < \infty$.

Proof sketch. At $p = 1$, optima are evaluated as one-defect strings, whereas noise can give a one-defect string a front value. Without reevaluation, these nonoptima retain their stored front values, strictly dominate stored optima and never decrease in number. Once all μ individuals store front values, no optimum can survive selection. A single generation producing μ such individuals has probability at least $n^{-(n+1)\mu}$ from any initial population with no true optimum for at least one front value. Under uniform initialization, starting with one such individual and cloning these individuals throughout the first generation reaches such a population with probability at least $\alpha_n = 2^{-n-o(n)}$. At $p = 1 - q$, independently requiring all evaluations through time B to be corrupted gives $\mathrm{P}[T > B] \geq \alpha_n(1-q)^{\mu(B+1)}$. Summing this tail yields $\mathbb{E}[T] \geq \alpha_n/(2\mu q)$ when $\mu q \leq 1/4$, hence the exponential lower bound.

With reevaluation, the observed defect differs from the true defect by one, so strings whose true defects differ by at least three remain correctly ordered. Selection and one-bit improvements first make at least $(1-\varepsilon)\mu$ individuals Pareto-optimal, for a small constant $\varepsilon > 0$. Crowding distance protects the endpoints of blocks with equal noisy fitness, and a class represented by only one individual gains a second surviving copy with probability $\Omega(1/\mu)$ per generation. Together with bounds on nonoptima that receive favorable fitness values, this bounds the probability of losing any class over a stage of at most $L_2 = \Theta(n^2/\mu^3)$ generations. A missing adjacent class is gained without another class being lost with probability $\Omega(1/(\mu n^2))$ per generation. Each stage therefore adds a class before any loss, while maintaining this proportion of optima, with probability $\Omega(\mu^{-4})$. An attempt of at most $M-1$ stages covers all classes with probability $\Omega(\mu^{-4(M-1)})$ and lasts $O(L_2)$ generations, since M is fixed. Repeating these attempts, including the time to prepare the population and to recover after rare failures, gives a constant success probability within $O(L_2\mu^{4(M-1)}) = O(n^2\mu^{4M-7})$ generations from every initial state. Restarting these blocks proves part (a). For $0 < q \leq 2^{-2n}$, coupling over such a polynomial-length block fails with probability at most its number of evaluations times q. The success probability remains bounded below by a positive constant, giving the same runtime bound. See Appendices F and G. □

The geometry above explains this reversal of the LOTZ comparison. At $p = 1$ every optimum has noisy fitness below the true front, while a one-defect string can have noisy fitness on it. Without reevaluation, improving stored values can therefore replace true optima with nonoptima. Once nonoptima storing front values fill the population, even a newly generated optimum is rejected: its noisy fitness is strictly dominated by their stored fitness values. The resulting trap comes from persistent evaluation errors, not from local optima of the true function.

Reevaluation breaks this mechanism without ever evaluating an optimum correctly at $p = 1$: it removes the permanent advantage of a favorable noisy fitness value. A one-defect string receives a front value on its next evaluation with probability at most $2/n$; otherwise its observed defect is two, whereas every optimum has observed defect one. An optimum of class j has noisy fitness $(j - M, r - j - M)$ with probability $1 - 2r/n$, when noise flips a bit of u. Subtracting the same constant from both objectives preserves the dominance relations, as also noted by Dang et al. (2023). The proof uses this fact to retain true optima and find missing classes through two-bit mutations of pairs. It must still control favorable noisy fitness values of nonoptima and losses of sparsely represented classes; preservation of true diversity is not deterministic.

Without reevaluation, $\mathbb{E}[T]$ is finite for $0 < q < 1$ (Proposition 2). The $2^{n-o(n)}$ lower bound is proved only for $0 < q \leq 2^{-2n}$ and does not imply that most runs take exponential time. Under the same population assumptions, the reevaluation upper bound extends to $0 \leq q \leq c_{\mathrm{R}}/(n^2\mu^{4M-6})$, with a sufficiently small constant $c_{\mathrm{R}} = c_{\mathrm{R}}(r, p_c) > 0$ (Theorem 12), allowing inverse-polynomial q when $\mu = \Theta(\log n)$.

5 Experiments

We examine whether the two effects identified by our analysis appear at moderate problem sizes: on LOTZ, repeated noisy evaluations remove Pareto-optimal solutions that retained fitness values would keep; on PairedCOCZ, retained incorrect fitness values exclude true optima. The constants in our theorems are far from these sizes, so the experiments concern the effects rather than the bounds. Every setting runs Algorithm 1 with $p_c = 0.5$ and uniform initialization 1000 times per strategy, and a run is complete once its population contains a true optimum of every front value. Reevaluation uses

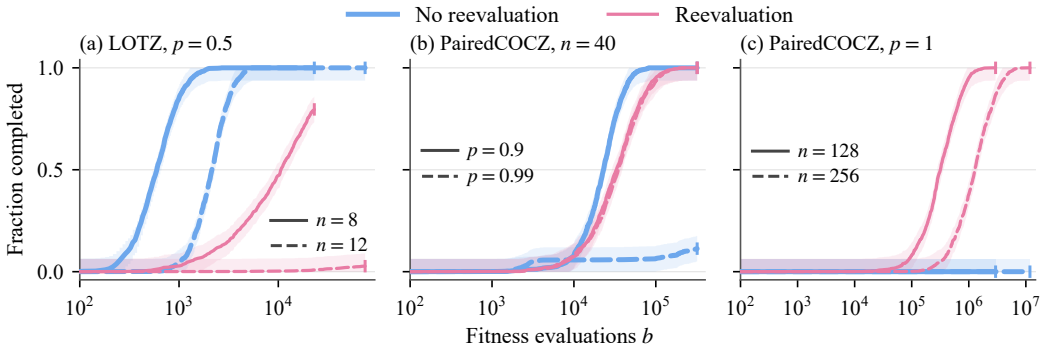

Figure 1: Fraction of runs that have covered the Pareto front after b fitness evaluations, with simultaneous 95% confidence bands; 1000 runs per strategy and setting. Each pair of curves ends at $b = \mu(1 + H)$, the cost of H generations without reevaluation. (a) LOTZ, $\mu = 2n + 2$, $H = 20n^2$. (b) PairedCOCZ, $r = 3$, $\mu = 32$, $H = 10^4$. (c) PairedCOCZ, $r = 3$, $\mu = 18$, $H = 10n^2$.

2μ evaluations per generation instead of μ, so Figure 1 compares the strategies at equal numbers of evaluations. Appendix H lists all settings and results and describes the simultaneous 95% confidence bands.

On LOTZ with $p = 0.5$ and $\mu = 2n + 2$, every run without reevaluation completes, on average after $0.54n^2$ to $0.60n^2$ generations for $n \in \{8, 12, 16, 24, 32, 64\}$ (Table 5). Within $H = 20n^2$ generations, more than 30 times this mean, reevaluation completes 95.6% of the runs at $n = 8$, 4.3% at $n = 12$ and none for $n \geq 16$; Figure 1(a) shows that the gap remains at equal evaluation budgets. Near completion, reevaluation loses $\Omega(np^2)$ Pareto-optimal solutions per generation in expectation (Section 3). Consistently, the gap in our settings appears only when p^2n is large: $p^2n \geq 2$ at $p = 0.5$, whereas for $p = 1/(8 \ln n)$ and $n \leq 128$, where $p^2n < 0.1$, both strategies complete every run and reevaluation needs only 4.8–8.3% more generations (Table 6).

On PairedCOCZ, the comparison reverses only when almost every evaluation is corrupted. At $n = 40$, $r = 3$ and $\mu = 32$, we compare the strategies after 32 032 evaluations, that is, 1000 generations without reevaluation and 500 with it (Figure 1(b) and Table 12). No reevaluation completes 40.5 and 27.0 percentage points more runs at $p = 0.5$ and $p = 0.9$, whereas reevaluation completes 40.0 and 41.6 points more at $p = 0.99$ and $p = 0.999$; all four simultaneous 95% intervals exclude zero. Counting generations instead would favor reevaluation already at $p = 0.9$, where it needs fewer generations on average (637 versus 784) but more evaluations (40 823 versus 25 131). At $p = 1$ with $n \in \{128, 256\}$ and $\mu = 18$, every run without reevaluation reaches a population of nonoptimal individuals that all store front values; no true optimum can survive selection from then on (Appendix G), so these runs never complete. Reevaluation completes every run, on average after 11 339 and 45 490 generations (Figure 1(c) and Table 11); the ratio of 4.0 between the two sizes is consistent with the factor n^2 in Theorem 4(a) for fixed μ. At $p = 1 - 2^{-2n}$, a noise-free evaluation is so unlikely at these sizes that the outcomes agree with those at $p = 1$ up to sampling variation (Appendix H).

6 Conclusion

For NSGA-II under one-bit prior noise, reevaluation can turn polynomial expected runtime into superpolynomial runtime on LOTZ with $\mu = 2n + 2$. On PairedCOCZ, it instead helps at and sufficiently near maximal noise. Neither strategy is always better: reevaluation can discard good solutions through new errors, while retaining incorrect fitness can exclude true optima. Future work could close the gaps between the proved polynomial and superpolynomial noise ranges on LOTZ.

AI Use Statement

In this work, we used generative AI tools to aid and polish the writing, including improving language clarity, grammar, and readability. We did not use generative AI tools to generate research results,

conduct data analysis, or develop substantive scientific claims; other disclosure categories are not applicable to this work. All AI-assisted text was carefully reviewed and revised by the authors to ensure accuracy, consistency with the underlying research, and appropriate scholarly presentation. We take responsibility for the final content of this work, including any text, claims, or artifacts produced with the aid of generative AI.

Reproducibility Statement

Complete proofs are provided in the appendices. Appendix H specifies the experimental protocol. The supplement includes source code, configurations, per-run seeds and outcomes, build metadata, and scripts reproducing all reported figures and tables.

References

Denis Antipov and Benjamin Doerr. Evolutionary algorithms are significantly more robust to noise when they ignore it. In *Proceedings of the 34th International Joint Conference on Artificial Intelligence (IJCAI'25)*, pp. 8842–8849, 2025.

Denis Antipov, Benjamin Doerr, and Alexandra Ivanova. Already moderate population sizes provably yield strong robustness to noise. In *Proceedings of the 26th ACM Conference on Genetic and Evolutionary Computation (GECCO'24)*, pp. 1524–1532, 2024.

Nicola Beume, Boris Naujoks, and Michael Emmerich. SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research*, 181:1653–1669, 2007.

Hans-Georg Beyer. Evolutionary algorithms in noisy environments: Theoretical issues and guidelines for practice. *Computer Methods in Applied Mechanics and Engineering*, 186(2-4):239–267, 2000.

Chao Bian and Chao Qian. Better running time of the non-dominated sorting genetic algorithm II (NSGA-II) by using stochastic tournament selection. In *Proceedings of the 17th International Conference on Parallel Problem Solving from Nature (PPSN'22)*, pp. 428–441, Dortmund, Germany, 2022.

Chao Bian, Shengjie Ren, Miqing Li, and Chao Qian. An archive can bring provable speed-ups in multi-objective evolutionary algorithms. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI'24)*, pp. 6905–6913, 2024.

Chao Bian, Yawen Zhou, Miqing Li, and Chao Qian. Stochastic population update can provably be helpful in multi-objective evolutionary algorithms. *Artificial Intelligence*, 341:104308, 2025.

Carlos A. Coello Coello and Gary B. Lamont. *Applications of Multi-Objective Evolutionary Algorithms*. World Scientific, Singapore, 2004.

Carlos A. Coello Coello, Gary B. Lamont, and David A. Van Veldhuizen. *Evolutionary Algorithms for Solving Multi-objective Problems*. Springer, 2007.

Duc-Cuong Dang, Andre Opris, Bahare Salehi, and Dirk Sudholt. Analysing the robustness of NSGA-II under noise. In *Proceedings of the 25th ACM Conference on Genetic and Evolutionary Computation (GECCO'23)*, pp. 642–651, Lisbon, Portugal, 2023.

Duc-Cuong Dang, Andre Opris, and Dirk Sudholt. Why dominance is not enough: Lessons from practical evolutionary multi-objective algorithms. *Algorithmica*, 88:60, 2026a.

Duc-Cuong Dang, Andre Opris, and Dirk Sudholt. SPEA2$^+$: Improved density estimation in SPEA2 with provable runtime guarantees. *arXiv preprint arXiv:2606.12382*, 2026b.

Kalyanmoy Deb. *Multi-objective Optimization using Evolutionary Algorithms*. Wiley, 2001.

Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.

Matthieu Dinot, Benjamin Doerr, Ulysse Hennebelle, and Sebastian Will. Runtime analyses of multi-objective evolutionary algorithms in the presence of noise. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI'23)*, pp. 5549–5557, Macao, SAR, China, 2023.

Benjamin Doerr. Probabilistic tools for the analysis of randomized optimization heuristics. In Benjamin Doerr and Frank Neumann (eds.), *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization*, Natural Computing Series, pp. 1–87. Springer, 2020.

Benjamin Doerr and Zhongdi Qu. Runtime analysis for the NSGA-II: Provable speed-ups from crossover. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI'23)*, pp. 12399–12407, Washington, DC, 2023a.

Benjamin Doerr and Zhongdi Qu. From understanding the population dynamics of the NSGA-II to the first proven lower bounds. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI'23)*, pp. 12408–12416, Washington, DC, 2023b.

Benjamin Doerr and Andrew M. Sutton. When resampling to cope with noise, use median, not mean. In *Proceedings of the 21st ACM Conference on Genetic and Evolutionary Computation (GECCO'19)*, pp. 242–248, 2019.

Benjamin Doerr, Tudor Ivan, and Martin S. Krejca. Speeding up the NSGA-II with a simple tie-breaking rule. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI'25)*, pp. 26964–26972, Philadelphia, PA, 2025a.

Benjamin Doerr, Martin S. Krejca, and Andre Opris. Tight runtime guarantees from understanding the population dynamics of the GSEMO multi-objective evolutionary algorithm. In *Proceedings of the 34th International Joint Conference on Artificial Intelligence (IJCAI'25)*, pp. 8876–8884, Montreal, Canada, 2025b.

Benjamin Doerr, Dimitri Korkotashvili, and Martin S. Krejca. Difficulties of the NSGA-II with the many-objective LeadingOnes problem. *IEEE Transactions on Evolutionary Computation*, 30(2): 773–782, 2026a.

Benjamin Doerr, Martin S. Krejca, and Milan Stanković. Improved runtime guarantees for the SPEA2 multi-objective optimizer. In *Proceedings of the 40th AAAI Conference on Artificial Intelligence (AAAI'26)*, pp. 36855–36863, Singapore, 2026b.

Benjamin Doerr, Martin S. Krejca, and Simon Wietheger. Speeding up the NSGA-II via dynamic population sizes. In *Proceedings of the 35th International Joint Conference on Artificial Intelligence (IJCAI'26)*, pp. 6273–6281, Bremen, Germany, 2026c.

Stefan Droste, Thomas Jansen, and Ingo Wegener. On the analysis of the (1+1) evolutionary algorithm. *Theoretical Computer Science*, 276(1-2):51–81, 2002.

Oliver Giel. Expected runtimes of a simple multi-objective evolutionary algorithm. In *Proceedings of the 2003 IEEE Congress on Evolutionary Computation (CEC'03)*, pp. 1918–1925, Canberra, Australia, 2003.

Chi Keong Goh and Kay Chen Tan. Noise handling in evolutionary multi-objective optimization. In *Proceedings of the IEEE Congress on Evolutionary Computation (CEC'06)*, pp. 1354–1361, Vancouver, Canada, 2006.

Jun He and Xin Yao. Drift analysis and average time complexity of evolutionary algorithms. *Artificial Intelligence*, 127(1):57–85, 2001.

Marco Laumanns, Lothar Thiele, and Eckart Zitzler. Running time analysis of multiobjective evolutionary algorithms on pseudo-Boolean functions. *IEEE Transactions on Evolutionary Computation*, 8(2):170–182, 2004.

Johannes Lengler. Drift analysis. In Benjamin Doerr and Frank Neumann (eds.), *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization*, Natural Computing Series, pp. 89–131. Springer, 2020.

Mian Li, Shapour Azarm, and Vikrant Aute. A multi-objective genetic algorithm for robust design optimization. In *Proceedings of the 7th ACM Conference on Genetic and Evolutionary Computation (GECCO'05)*, pp. 771–778, Washington, DC, 2005.

Pascal Massart. The tight constant in the Dvoretzky–Kiefer–Wolfowitz inequality. *The Annals of Probability*, 18(3):1269–1283, 1990.

Colin McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics, 1989*, volume 141 of *London Mathematical Society Lecture Note Series*, pp. 148–188. Cambridge University Press, 1989.

Chao Qian, Yang Yu, and Zhi-Hua Zhou. An analysis on recombination in multi-objective evolutionary optimization. *Artificial Intelligence*, 204:99–119, 2013.

Chao Qian, Yang Yu, Ke Tang, Yaochu Jin, Xin Yao, and Zhi-Hua Zhou. On the effectiveness of sampling for evolutionary optimization in noisy environments. *Evolutionary Computation*, 26(2): 237–267, 2018a.

Chao Qian, Yang Yu, and Zhi-Hua Zhou. Analyzing evolutionary optimization in noisy environments. *Evolutionary Computation*, 26(1):1–41, 2018b.

Chao Qian, Chao Bian, Yang Yu, Ke Tang, and Xin Yao. Analysis of noisy evolutionary optimization when sampling fails. *Algorithmica*, 83(4):940–975, 2021.

Shengjie Ren, Zimin Liang, Miqing Li, and Chao Qian. Not just for archiving: Provable benefits of reusing the archive in evolutionary multi-objective optimization. In *Proceedings of the 40th AAAI Conference on Artificial Intelligence (AAAI'26)*, pp. 37152–37160, Singapore, 2026.

Dirk Sudholt. Analysing the robustness of evolutionary algorithms to noise: Refined runtime bounds and an example where noise is beneficial. *Algorithmica*, 83:976–1011, 2021.

Guangyong Sun, Huile Zhang, Jianguang Fang, Guangyao Li, and Qing Li. A new multi-objective discrete robust optimization algorithm for engineering design. *Applied Mathematical Modelling*, 53:602–621, 2018.

Erfan Babaee Tirkolaee, Alireza Goli, Ali Faridnia, Mostafa Soltani, and Gerhard-Wilhelm Weber. Multi-objective optimization for the reliable pollution-routing problem with cross-dock selection using Pareto-based algorithms. *Journal of Cleaner Production*, 276:122927, 2020.

Qingfu Zhang and Hui Li. MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6):712–731, 2007.

Weijie Zheng and Benjamin Doerr. Better approximation guarantees for the NSGA-II by using the current crowding distance. In *Proceedings of the 24th ACM Conference on Genetic and Evolutionary Computation (GECCO'22)*, pp. 611–619, Boston, MA, 2022.

Weijie Zheng and Benjamin Doerr. Mathematical runtime analysis for the non-dominated sorting genetic algorithm II (NSGA-II). *Artificial Intelligence*, 325:104016, 2023.

Zhi-Hua Zhou, Yang Yu, and Chao Qian. *Evolutionary Learning: Advances in Theories and Algorithms*. Springer, Singapore, 2019.

The appendices give detailed statements and proofs of the main results in theorem order, including sufficient parameter choices. Theorem 1 is Theorem 6(ii), Theorem 2 is Theorem 7, and Theorem 3 is Theorem 8. Each is stated at the head of the appendix that proves it; Theorem 6 also bounds the time until every front value is stored. Appendices F and G prove the PairedCOCZ bounds, including the extension to noise levels below one.

A Probabilistic Tools

This appendix collects the probabilistic facts the other appendices cite; none refers to the NSGA-II. Throughout, $(\mathcal{F}_t)_{t\geq 0}$ is an arbitrary filtration. A process $(X_t)_{t\geq 0}$ is *adapted* if X_t is $\mathcal{F}_t$-measurable for every t. We write $a \wedge b := \min\{a, b\}$. For real random variables, $X \preceq Y$ denotes stochastic domination, that is, $\mathrm{P}[X \geq r] \leq \mathrm{P}[Y \geq r]$ for every real r. Between value vectors, the symbols $\preceq$ and $\succeq$ keep their meaning from Section 2, the coordinatewise order. The type of the two sides always shows which meaning applies.

A.1 Concentration Bounds

Theorem 5 (Chernoff bounds; see e.g. Doerr, 2020)**.** *Let $X = \sum_{i=1}^{m} X_i$, where $X_1, \ldots, X_m$ are independent random variables with values in $[0, 1]$.*

*(a) (*lower tail*) For every $\delta \in (0, 1)$,*

$$\mathrm{P}\big[X \leq (1-\delta)\mathbb{E}[X]\big] \;\leq\; \exp\!\big(-\delta^2 \mathbb{E}[X]/2\big).$$

*(b) (*upper tail*) For every $\delta \in (0, 1)$,*

$$\mathrm{P}\big[X \geq (1+\delta)\mathbb{E}[X]\big] \;\leq\; \exp\!\big(-\delta^2 \mathbb{E}[X]/3\big).$$

A.2 Additive Drift

The following drift result bounds an expected hitting time from above. It needs a potential that decreases in expectation by at least a fixed amount per step.

Lemma 1 (additive drift, hitting-time form; cf. He & Yao, 2001; Lengler, 2020)**.** *Let $(\Psi_t)_{t\geq 0}$ be adapted to $(\mathcal{F}_t)$ with $0 \leq \Psi_t \leq B$ for every t, and let τ be a stopping time. Suppose that there is a $\delta > 0$ such that, for every t,*

$$\mathbb{E}[\Psi_{t+1} - \Psi_t \mid \mathcal{F}_t] \;\leq\; -\delta \qquad \text{on the event } \{t < \tau\}.$$

Then

$$\mathbb{E}[\tau] \;\leq\; \frac{B}{\delta}.$$

Proof. Step 1: a supermartingale. Define the stopped and compensated process $M_t := \Psi_{t\wedge\tau} + \delta\,(t \wedge \tau)$. On the event $\{t < \tau\}$ we have $M_{t+1} - M_t = \Psi_{t+1} - \Psi_t + \delta$, and the hypothesis gives $\mathbb{E}[M_{t+1} - M_t \mid \mathcal{F}_t] \leq 0$. On the event $\{t \geq \tau\}$ the process is constant, $M_{t+1} = M_t$. Hence (M_t) is a supermartingale.

Step 2: a bound for the stopped time. The supermartingale property gives $\mathbb{E}[M_t] \leq \mathbb{E}[M_0]$. We use the definition of M_t, then $\Psi \geq 0$, then this inequality, then $\Psi_0 \leq B$:

$$\delta\,\mathbb{E}[t \wedge \tau] \;=\; \mathbb{E}[M_t] - \mathbb{E}[\Psi_{t\wedge\tau}] \;\leq\; \mathbb{E}[M_t] \;\leq\; \mathbb{E}[M_0] \;=\; \mathbb{E}[\Psi_0] \;\leq\; B.$$

Step 3: the limit. Let $t \to \infty$. By monotone convergence, $\mathbb{E}[t \wedge \tau]$ increases to $\mathbb{E}[\tau]$. Hence $\delta\,\mathbb{E}[\tau] \leq B$. □

B Auxiliary Lemmas for LOTZ

This appendix collects deterministic facts about the values of strings under LeadingOnesTrailingZeroes and how the survival selection (Algorithm 1) treats them; the only probability involved is the

elementary one in Lemma 5(b). The statements concern the *observed fitness values* used in a reevaluating survival step. Every such value is the true fitness value of some string, as is every stored fitness value without reevaluation, so the statements also apply under that strategy.

We recall from Section 2 the value $g(x) = (\mathrm{LO}(x), \mathrm{TZ}(x))$, the Pareto strings $\pi_k = 1^k 0^{n-k}$, the Pareto set F, and the word *slot* for the index k.

- *Front and off-front values.* A *front value* is a vector $(k, n-k)$ with $0 \le k \le n$. A value (a, b) with $a + b < n$ is called *off-front*. By Lemma 2 every value is of one of these two types.
- *Combined population and fitness values.* We write R for the combined parent and offspring population (2μ individuals) that enters a survival-selection step of Algorithm 1. The *observed fitness value* $\tilde{v}(z)$ of an individual $z \in R$ is the objective vector used in this step; below, *fitness value* refers to this vector unless explicitly qualified as true. An individual $z \in R$ with $\tilde{v}(z) = (k, n-k)$ is a *representative* of the front value $(k, n-k)$, or briefly of k.
- *Noise.* An evaluation of y returns $\tilde{v}(y) = g(y)$ with probability $1-p$ and $\tilde{v}(y) = g(y \oplus e_I)$ otherwise (Section 2). For the proofs, let $B \sim \mathrm{Bernoulli}(p)$, independent of I, indicate which case occurs: $B = 0$ for the first and $B = 1$ for the second. All evaluations use fresh independent pairs (B, I). The evaluation is *corrupted* if $B = 1$. A survival step is *corruption-free* if none of its 2μ evaluations is corrupted.
- *Layer 1 and blocks. Layer 1* is the set L_1 of Algorithm 1. It consists of the individuals whose fitness value is not strictly dominated by any fitness value in R. A *layer-1 fitness value* is the fitness value of a layer-1 individual. Whether an individual lies in layer 1 depends only on its fitness value. The *block* of a fitness value v is the set of individuals in R with fitness value v.
- *Hamming neighbours.* $\mathrm{Ham}_1(x)$ is the set of strings that differ from x in exactly one bit.
- *Substrings.* $x[i..j]$ is the substring of x from position i to position j. The symbol $\delta_{i,j}$ is 1 if $i = j$ and 0 otherwise.

B.1 Front Values and Dominance

Lemma 2. *The following two statements hold.*

(a) Every string $y \in \{0,1\}^n$ satisfies

$$\mathrm{LO}(y) + \mathrm{TZ}(y) \;\le\; n,$$

and equality holds if and only if $y \in F$.

(b) Hence every fitness value (a, b) satisfies $a + b \le n$. A fitness value with $a + b = n$ has the unique preimage π_a: the only string with value $(a, n-a)$ is π_a.

Proof. Step 1: the inequality. Let $a = \mathrm{LO}(y)$ and $b = \mathrm{TZ}(y)$. The leading block of ones occupies the positions $1, \dots, a$, and the trailing block of zeros occupies the positions $n-b+1, \dots, n$. If $a + b > n$, then position a belongs to both blocks, which is a contradiction.

Step 2: the equality case. If $a + b = n$, then the first a bits are 1 and the last $b = n - a$ bits are 0. This determines $y = \pi_a$. Conversely, $g(\pi_a) = (a, n-a)$.

Step 3: fitness values. An observed fitness value is obtained by applying g to the evaluated string, which is the individual's genotype with possibly one flipped bit. Hence it is the true fitness value of some string, and (a) applies to it. □

Lemma 3. *Let y be an individual of R whose fitness value equals a front value $(k, n-k)$. Then no fitness value in R strictly dominates $(k, n-k)$. Consequently, y lies in layer 1 of the non-dominated sort of R.*

Proof. A strict dominator (a, b) would need $a \ge k$ and $b \ge n - k$ with one of the two inequalities strict. This gives $a + b > n$, which contradicts Lemma 2. □

B.2 Single-Bit Flips of Pareto Strings

One-bit noise evaluates a string after flipping at most one bit. For a Pareto string the possible outcomes are listed next.

Lemma 4. *The single-bit flips of the Pareto strings have the following values.*

*(a) (*interior slots*) Let $0 < k < n$. Flipping bit i of π_k yields the following string and value.*

bit	*string*	*value*	*type*
$i < k$	$1^{i-1}\,0\,1^{k-i}\,0^{n-k}$	$(i-1,\ n-k)$	*off-front*
$i = k$	$1^{k-1}\,0^{n-k+1} = \pi_{k-1}$	$(k-1,\ n-k+1)$	*front value of slot* $k-1$
$i = k+1$	$1^{k+1}\,0^{n-k-1} = \pi_{k+1}$	$(k+1,\ n-k-1)$	*front value of slot* $k+1$
$i > k+1$	$1^{k}\,0^{i-k-1}\,1\,0^{n-i}$	$(k,\ n-i)$	*off-front*

Thus exactly 2 of the n flips produce front values, namely those of the two neighbouring slots. The remaining $n-2$ flips produce off-front values.

*(b) (*extreme slots*) Flipping bit i of π_n yields the value $(i-1, \delta_{i,n})$. This is the front value of π_{n-1} if $i = n$, and an off-front value otherwise. Symmetrically, flipping bit i of π_0 yields the value $(\delta_{i,1}, n-i)$. This is the front value of π_1 if $i = 1$, and an off-front value otherwise.*

*(c) (*dominance*) Every off-front fitness value (a, b), that is, every fitness value with $a + b < n$, is* strictly *dominated by the front value $(k', n-k')$ of every slot $k' \in [a,\ n-b]$.*

Proof. Parts (a) and (b) follow by direct computation of LO and TZ of the displayed strings. In (b) the flipped strings are $1^{i-1}\,0\,1^{n-i}$ and $0^{i-1}\,1\,0^{n-i}$. For (c), let $k' \in [a, n-b]$. Then $k' \geq a$ gives the first coordinate, and $n - k' \geq b$ gives the second. Moreover $a + b < n = k' + (n-k')$ forces strictness in at least one coordinate. □

The next lemma characterises the representatives of a front value $(k, n-k)$. Besides the copies of π_k, only Hamming neighbours of π_k with a corrupted evaluation are possible.

Lemma 5. *Consider one evaluation of a string y with noise pair (B, I), and let $0 \leq k \leq n$.*

(a) The fitness value satisfies the following equivalences:

$$\begin{aligned} B = 0: &\quad \tilde{v}(y) = (k, n-k) \iff y = \pi_k, \\ B = 1: &\quad \tilde{v}(y) = (k, n-k) \iff y \oplus e_I = \pi_k. \end{aligned}$$

(b) Hence a representative of $(k, n-k)$ is of exactly one of two kinds.

(i) It is the string π_k itself, with an evaluation that did not change its value; we call it a correct *representative. For π_k this forces $B = 0$, because every single-bit flip changes $g(\pi_k)$ (Lemma 4).*

(ii) It is an incorrect *representative: a string $y \in \mathrm{Ham}_1(\pi_k)$ whose evaluation flipped exactly the bit in which y and π_k differ. For a given $y \in \mathrm{Ham}_1(\pi_k)$ this event has probability exactly p/n per evaluation.*

(c) In a corruption-free survival step, the representatives of $(k, n-k)$ are exactly the copies of π_k.

Proof. Step 1: part (a). This is the equality case of Lemma 2. The evaluated string is y when $B = 0$ and $y \oplus e_I$ when $B = 1$. The value $(k, n-k)$ has the unique preimage π_k. So the evaluated string must equal π_k.

Step 2: part (b). If $B = 0$, part (a) gives $y = \pi_k$. If $B = 1$, it gives $y \oplus e_I = \pi_k$, so $y \in \mathrm{Ham}_1(\pi_k)$ and I equals the position in which y and π_k differ. For a given $y \in \mathrm{Ham}_1(\pi_k)$ this event has probability $p \cdot \frac{1}{n}$, because B and I are independent.

Step 3: part (c). In a corruption-free survival step every evaluation has $B = 0$. By Step 2 only kind (i) remains. □

B.3 Non-Dominated Sorting and Selection

Survival selection never removes the last individual with a given layer-1 fitness value. This preserves the fitness value, not necessarily the genotype of the individuals having it.

Lemma 6. *Among the distinct layer-1 fitness values of any combined population there is at most one with any given first coordinate, and at most one with any given second coordinate. Consequently, the distinct layer-1 fitness values are strictly increasing in one coordinate as the other decreases: if (a, b) and (a', b') are distinct layer-1 fitness values with $a < a'$, then $b > b'$.*

Proof. Let (k, b) and (k, b') be two distinct fitness values with the same first coordinate, say with $b > b'$. Then $(k, b) \succ (k, b')$, so they cannot both be layer-1 fitness values. The argument for the second coordinate is symmetric. Monotonicity follows because layer-1 fitness values are pairwise incomparable. □

Lemma 7. *Consider any survival step (Algorithm 1) with population size $\mu \geq 2(n+1)$, and let P_{t+1} be the population that it returns.*

- *(a) The distinct layer-1 fitness values number at most $n + 1$.*
- *(b) The following three statements hold.*
  - *(b1) For each distinct layer-1 fitness value, the members of its block with positive crowding distance are exactly the first and last members of the block, called its* block endpoints*. They coincide for a singleton block; otherwise there are exactly two.*
  - *(b2) The layer-1 individuals with positive crowding distance number at most $2(n+1) \leq \mu$, and* all of them survive*.*
  - *(b3) Each distinct layer-1 fitness value is the fitness value of some member of P_{t+1}.*

Proof. Step 1: part (a). By Lemmas 2 and 6, the distinct layer-1 fitness values have pairwise distinct first coordinates in $\{0, \ldots, n\}$, so there are at most $n + 1$ of them. Their second coordinates are also pairwise distinct, so individuals with equal values in either objective have equal fitness values.

Step 2: at most two members of each block have positive crowding (b1). In each per-objective stable sort of layer 1 (Algorithm 1), equal fitness values occupy contiguous positions. Consider an element that is interior to its block. Both of its sort neighbours have the same value in that objective. Its normalized gap contribution is therefore 0 when the objective range is positive; when the range is zero, its contribution is 0 by the zero-range convention. Only the first and the last element of the block can receive a positive or infinite contribution. A block is contiguous in *both* sorts, and its pair of extremal elements is the same in both sorts. Indeed, both stable sorts preserve the input order within the block, so its first and last members coincide across sorts. Hence at most 2 members of each block have positive crowding distance.

Step 3: the count in (b2). By (a) and Step 2, layer 1 contains at most $2(n + 1) \leq \mu$ individuals with positive crowding distance.

Step 4: they all survive (b2). The truncation step of Algorithm 1 takes whole layers first and cuts the last admitted layer by crowding distance in descending order. If layer 1 has size at least μ, then the μ survivors are chosen from layer 1. The members with positive crowding distance come first, and all of them are taken because they number at most μ; the remaining survivors have crowding distance zero. If layer 1 has size less than μ, then it survives entirely. In both cases every layer-1 individual with positive crowding distance survives.

Step 5: each layer-1 fitness value is kept (b3). Every block endpoint has positive crowding distance: if it is an endpoint of the whole sort, it receives $+\infty$; otherwise its neighbour outside the block has a different coordinate by Step 1. The corresponding objective range is then strictly positive, and dividing the positive gap by this range preserves positivity. This also applies to a singleton block. If the layer contains just one distinct fitness value, both objective ranges are zero, but the block endpoints are whole-sort boundaries and still receive $+\infty$; only its interior members receive zero. Thus Step 4 preserves the block endpoints and hence at least one individual with each distinct layer-1 fitness value. □

The last lemma of this subsection concerns the binary tournaments rather than the survival step. It involves no noise and applies to both evaluation strategies.

Lemma 8. *Consider one binary tournament in any generation. It uses the stored ranks and the stored crowding distances that the previous survival step produced.*

(a) A position with stored rank 1 *and stored crowding* $+\infty$ *wins the tournament with probability at least* $\frac{1}{2\mu}$.

(b) Let T *be any set of positions that all have stored rank* 1*. The tournament returns a member of* T *with probability at least* $(|T|/\mu)^2$.

Proof. Part (a). The position is among the two samples with probability

$$1 - \left(1 - \frac{1}{\mu}\right)^2 \;\geq\; \frac{1}{\mu}.$$

Condition on the position being sampled. It wins outright against any opponent of stored rank larger than 1 or of finite stored crowding. Against an opponent that also has stored rank 1 and stored crowding $+\infty$, it wins the fair-coin tie with probability $\frac{1}{2}$. So its conditional winning probability is at least $\frac{1}{2}$, and the claim follows.

Part (b). With probability $(|T|/\mu)^2$ both samples lie in T. Then the winner lies in T, whichever of the two it is. □

B.4 Mutation and One-Point Crossover

Lemma 9. *Let* y *be a string. Define its set of* productive bits *by*

$$\mathcal{P}(y) := \{\mathrm{LO}(y) + 1,\ n - \mathrm{TZ}(y)\} \cap \{1, \ldots, n\}.$$

For every $i \in \{1, \ldots, n\} \setminus \mathcal{P}(y)$, $g(y \oplus e_i) \preceq g(y)$ *coordinatewise.*

Proof. Put $a = \mathrm{LO}(y)$ and $b = \mathrm{TZ}(y)$. If $a < n$, increasing LO by a single bit flip requires changing the first zero, at position $a + 1$; if $a = n$, an increase is impossible. Similarly, if $b < n$, increasing TZ requires changing the last one, at position $n - b$; if $b = n$, an increase is impossible. Thus flipping a valid position outside $\mathcal{P}(y)$ increases neither coordinate, including at the boundary strings 0^n and 1^n. □

Lemma 10. *Let* $x = \pi_a$ *and* $y = \pi_b$ *be Pareto strings with* $b \leq a$*, and let* $c \in \{1, \ldots, n-1\}$ *be a cut. Let* $z := x[1..c]\, y[c{+}1..n]$ *be the child that takes its prefix from* x *and its suffix from* y*. Then*

$$z \;=\; \begin{cases} \pi_{\max(c,b)} & \text{if } c \leq a, \\ \pi_a & \text{if } c > a. \end{cases}$$

In particular, the following two statements hold.

*(i) (*closure*) The child* z *with its prefix from the larger-index parent is a Pareto string.*

*(ii) (*the two extremes*) For the parents* (π_n, π_0)*, every cut* c *yields* π_c *as the first child. The second child is* $\pi_0[1..c]\,\pi_n[c{+}1..n] = 0^c 1^{n-c}$*, which is* not *a Pareto string for* $0 < c < n$.

Proof. Step 1: the case $c \leq a$*.* The prefix of z is 1^c. The suffix $y[c{+}1..n]$ equals 0^{n-c} if and only if $c \geq b$, and then $z = \pi_c$. For $c < b$ the suffix begins with 1^{b-c}, and the concatenation is $1^b 0^{n-b} = \pi_b$. In both cases $z = \pi_{\max(c,b)}$.

Step 2: the case $c > a$*.* Since $c > a \geq b$, the suffix of y is 0^{n-c}, so $z = 1^a 0^{c-a} 0^{n-c} = \pi_a$.

Step 3: items (i) and (ii). The two formulas prove (i). For (ii), take $a = n$ and $b = 0$. Then $c \leq a$ for every cut, and the first child is $\pi_{\max(c,0)} = \pi_c$. The second child is computed directly. □

C Proof of Theorem 1: Noise Upper Bound for the No-Reevaluation Strategy

This appendix analyses the NSGA-II *without* reevaluation: every individual is evaluated exactly once, at the moment it is created, and the value returned by that evaluation is its *stored value*,

used by all later sorts and carried forward unchanged by a survivor. One fact from Appendix B is used throughout: a stored value is always the true value of *some* string, namely of the string that was actually evaluated, so the lemmas of Appendix B, which are phrased for fitness values, hold verbatim for stored values. Part (ii) of the following theorem is Theorem 1 with $\delta = c_\star$; part (i) bounds the time until all front values are stored, which is used in its proof.

Theorem 6. *Let $p, p_c \in [0, 1)$, put $A := (1-p)(1-p_c)$, and let $\mu = 2n+2$ and $n \geq 3$. For the no-reevaluating NSGA-II of Algorithm 1, from an arbitrary initial population:*

(i) all front values are eventually stored and remain stored forever; the first generation T_{val} in which every front value is the stored value of some member of the population satisfies

$$\mathbb{E}[T_{\mathrm{val}}] \leq \frac{8e\mu n^3}{A};$$

(ii) there is an absolute $c_\star > 0$ such that, for every fixed $p_c \in (0, 1)$, all sufficiently large n, and

$$p \leq c_\star \frac{p_c(1-p_c)}{\log n},$$

$$\mathbb{E}[T] \leq \frac{8e\mu n^3}{A} + \frac{4e\mu n(n+1)}{A} + O\Big(\frac{n \log n}{1-p_c}\Big) = O\Big(\frac{n^4}{1-p_c}\Big).$$

Let T_{eval} count fitness evaluations until the Pareto front is covered, including initialization. Then $\mathbb{E}[T_{\mathrm{eval}}] = \mu + \mu\mathbb{E}[T]$; in particular, part (ii) gives $O(n^5/(1-p_c))$ evaluations.

As in Appendix B, an individual with stored value $(k, n-k)$ is a *representative* of k. It is either the Pareto string π_k (a correct representative) or a Hamming neighbour of π_k whose evaluation was corrupted (an incorrect representative). The analysis has four phases. Phase 0 finds the first stored front value. Phase 1 finds all front values; once present, a stored front value is never lost (Lemma 11). In Phase 2 every front value obtains two representatives; from then on the population contains exactly two representatives of each front value, since $2(n+1) = \mu$ (Lemma 13), although these individuals may change in every generation. In Phase 3 the incorrect representatives are replaced: both representatives of a value k may still be incorrect ($B_k = 1$), and the front is covered once $B_k = 0$ for all k simultaneously. Phases 0 to 2 (Section C.2) only use offspring of a parent that is sampled twice in a tournament, and they work for every $p_c \in [0, 1)$. Phase 3 (Section C.3) consists of epochs whose one-generation estimates hold conditionally on *every* history after Phase 2, so a failed epoch is simply followed by a new one. Only Phase 3 uses that parents precede offspring in R_t: if an offspring is a correct representative of k and no offspring is an incorrect one, the last member of block k is a correct offspring and survives. Preserving front values alone would not ensure that a correct string survives.

C.1 Notation and Two Elementary Estimates

Standing conventions. Throughout, $\mu = 2n+2$ and $n \geq 3$, so that $\mu \geq 8$. Here P_t is the population at the start of generation t, Q_t is its offspring multiset, $R_t = P_t \cup Q_t$ is the combined population, and $\mathcal{F}_t$ is the history up to the start of generation t. The survival step of generation t selects P_{t+1} from R_t. For the upper bounds we put $A := (1-p)(1-p_c)$, the probability that a round skips crossover and that one evaluation is not corrupted.

Stable sorting. The combined population R_t lists all members of P_t before all members of Q_t. Each layer inherits this order, and its two ascending objective sorts are stable. Within a block, both sorts therefore put the parents before the offspring and preserve their respective input orders. The order of the survivors in P_{t+1} need not match their order in R_t: the next combined population again puts all parents before all offspring. Stored crowding distances, computed at the preceding survival step, remain the distances used by the next generation's tournaments.

Representatives and blocks.

- A *representative of* k is an individual whose stored value is $(k, n-k)$, and *block* k of a population or of R_t is the set of its representatives of k.

- By Lemma 5, a representative of k is either the string π_k, whose evaluation was then not corrupted (a *correct* representative), or a string $y \in \mathrm{Ham}_1(\pi_k)$ whose evaluation flipped exactly the bit in which y and π_k differ (an *incorrect* representative). For a given $y \in \mathrm{Ham}_1(\pi_k)$ the latter has probability exactly p/n. In both cases a representative of k differs from π_k in at most one bit.
- The *endpoints* of a block are its first and last members in the input order retained by both objective sorts. They coincide when the block is a singleton and otherwise are two distinct individuals. From T_{lock} on (defined below), the two members of each block of P_t are the endpoints retained from the sort of R_{t-1}, or from the initial sort if $t = 0$; they need not be endpoints in the sort of R_t.
- A *new representative of k in generation t* is an offspring created in generation t whose stored value is $(k, n-k)$; it is correct or incorrect as above.

Hitting times.

- T_{val} is as in Theorem 6(i), and T_{lock} is the first generation $t \geq T_{\mathrm{val}}$ such that every block of P_t has at least two members. Since $|P_t| = \mu = 2(n+1)$, every block of $P_{T_{\mathrm{lock}}}$ then has exactly two members, and this remains so afterwards (Lemma 13).
- A statement holds *from every history after* T_{lock} if it holds conditionally on $\mathcal{F}_t$ for every $t \geq T_{\mathrm{lock}}$ and every realisation of $\mathcal{F}_t$.
- For $t \geq T_{\mathrm{lock}}$,

$$B_k(t) := \mathbf{1}\{\text{neither member of block } k \text{ in } P_t \text{ equals } \pi_k\}$$

indicates that block k is *bad*, that is, both of its members are incorrect representatives; otherwise block k is *good*. We also write $d(k) := \min\{k, n-k\}$ for the distance of slot k from the nearer end of the front.

Two elementary estimates. The first estimate concerns bit-wise mutation. Let $S \subseteq \{1, \dots, n\}$ be a prescribed set of positions with $|S| \leq 2$. Then

$$\mathrm{P}[\text{mutation flips exactly the positions in } S] = \frac{1}{n^{|S|}}\Big(1-\frac{1}{n}\Big)^{n-|S|} \;\geq\; \begin{cases} 1/(en) & \text{if } |S| \leq 1, \\ 1/(en^2) & \text{if } |S| \leq 2. \end{cases} \tag{M}$$

To see this, recall that $(1-1/n)^{n-1} \geq 1/e$. For $|S| \in \{1,2\}$ we bound $(1-1/n)^{n-|S|} \geq (1-1/n)^{n-1} \geq 1/e$. For $S = \emptyset$ we use $(1-1/n)^n \geq (1-1/n)/e \geq 1/(en)$, which holds because $n \geq 2$. The second line of (M) follows from the first because $1/(en) \geq 1/(en^2)$.

The second estimate passes from one reproduction round to one generation. Let $m \geq 1$ and $q \in [0,1]$ with $mq \leq 1$. Then

$$1-(1-q)^m \;\geq\; 1-e^{-mq} \;\geq\; mq - \frac{(mq)^2}{2} \;\geq\; \frac{mq}{2}. \tag{R}$$

We use (R) as follows. Given $\mathcal{F}_t$, the $\mu/2$ reproduction rounds of generation t are independent and identically distributed. If one round succeeds with probability at least q, and $\mu q/2 \leq 1$, then generation t contains at least one successful round with probability at least $1-(1-q)^{\mu/2} \geq \mu q/4$.

C.2 Analysis of Phases 0 to 2

This section covers Phases 0 to 2 of the upper bound. From here on we assume the hypotheses of Theorem 6: $p, p_c \in [0,1)$, $A = (1-p)(1-p_c)$, $\mu = 2n+2$, the stable sorting rule of Section 2.3, and an arbitrary initial population.

Phases 0 and 1: all front values are found and never lost. The first lemma says that stored front values are permanent, and that the largest coordinate sum of a stored value never decreases.

Lemma 11. *Consider the no-reevaluating NSGA-II with $\mu = 2n+2$.*

(a) If some member of P_t has the stored value $(k, n-k)$, then some member of $P_{t'}$ has it for every $t' > t$.

(b) The quantity

$$V_t^{\max} := \max\{a+b : \ (a,b) \text{ is a stored value in } P_t\}$$

is non-decreasing in t.

Proof. (a) Let $(k, n-k)$ be stored in P_t, hence in R_t. By Lemma 3 a stored front value is never strictly dominated, so its representatives have rank 1 in every sort. By Lemma 7(b) at least one representative, a block endpoint, survives the survival step. The stored value of a survivor never changes. So $(k, n-k)$ is stored in P_{t+1}, and the claim follows by induction on t'.

(b) Let s be the maximal coordinate sum of a stored value in R_t, and let (a, b) be a stored value with $a + b = s$. No stored value in R_t strictly dominates (a, b), because a dominator would have a larger coordinate sum. Hence (a, b) has rank 1 and is preserved by the same argument as in (a). Therefore $V_{t+1}^{\max} = s$. Since $P_t \subseteq R_t$, we have $s \geq V_t^{\max}$: new values can only raise the maximum. □

Proof of Theorem 6(i). Permanence is Lemma 11(a). It remains to bound $\mathbb{E}[T_{\text{val}}]$.

Step 1: creating a prescribed offspring. Fix an individual $x \in P_t$ and a prescribed set S of at most two positions. A fixed reproduction round creates the offspring $x \oplus \bigoplus_{i \in S} e_i$ with an uncorrupted evaluation if the following four independent events occur.

- *Selection.* Both samples of the first tournament are x. Then x wins the tournament, whatever the stored ranks are. The probability is $1/\mu^2$.
- *No crossover.* The round skips crossover, so the first child equals x before mutation. The probability is $1 - p_c$.
- *Mutation.* Mutation flips exactly the positions in S. By (M) the probability is at least $1/(en^2)$.
- *Uncorrupted evaluation.* The evaluation of the child is not corrupted. The probability is $1 - p$.

So one round creates this offspring with probability at least

$$q \ := \ \frac{1}{\mu^2} \cdot (1 - p_c) \cdot \frac{1}{en^2} \cdot (1 - p) \ = \ \frac{A}{e\mu^2 n^2}.$$

By (R), applied to the $\mu/2$ rounds of a generation, a prescribed progress event of this kind occurs in one generation with probability at least

$$\sigma := \frac{\mu q}{4} = \frac{A}{4e\mu n^2}.$$

This bound holds conditionally on every history $\mathcal{F}_t$.

Step 2: reaching the first front value (Phase 0). Suppose that $V_t^{\max} < n$. Choose the first individual $x \in P_t$ whose stored value (a, b) satisfies $a + b = V_t^{\max}$. Let $\widehat{x}$ be the lexicographically first string in $\{x\} \cup \mathrm{Ham}_1(x)$ with $g(\widehat{x}) = (a, b)$; one exists because the stored value came from a one-bit prior-noise evaluation of x. Both choices are $\mathcal{F}_t$-measurable and need not reveal the creation-noise draw. We flip one productive bit of $\widehat{x}$ (Lemma 9), namely position $a + 1$. In detail: position $a + 1$ of $\widehat{x}$ is a zero because $\mathrm{LO}(\widehat{x}) = a < n$, and position $n - b$ is a one because $\mathrm{TZ}(\widehat{x}) = b < n$. We have $a + 1 \leq n - b$ because $a + b < n$, and the two positions differ because they hold different bits. Hence $a + 1 < n - b$. The string $z := \widehat{x} \oplus e_{a+1}$ has at least $a + 1$ leading ones and still exactly b trailing zeros, so the coordinate sum of $g(z)$ is at least $a + b + 1$. Starting from x, the required mutation flips the symmetric difference of $\{a + 1\}$ and the set of positions where x and $\widehat{x}$ differ. This is an $\mathcal{F}_t$-measurable set of at most two positions, with coincident flips cancelling. So the event of Step 1 with this x and this set creates z and stores $g(z)$. By the proof of Lemma 11(b), $V_{t+1}^{\max}$ is the maximal coordinate sum in R_t, so $V_{t+1}^{\max} \geq a + b + 1$. The maximal stored coordinate sum thus increases, and the increase is permanent by Lemma 11(b). Since $0 \leq V_t^{\max} \leq n$, at most n such successes reach $V_t^{\max} = n$, that is, a stored front value.

Step 3: spreading along the front. Thereafter let $Y \subseteq \{0, \ldots, n\}$ be the set of values k such that $(k, n-k)$ is stored in P_t. By Lemma 11(a) this set never loses an element. If $Y \neq \{0, \ldots, n\}$, choose some $k \in Y$ and $k' \in \{k-1, k+1\}$ with $k' \notin Y$. Let x be a representative of k. It is within

Hamming distance one of π_k, and $\pi_{k'}$ differs from π_k in exactly one bit (position $k+1$ if $k' = k+1$, position k if $k' = k-1$). So producing $\pi_{k'}$ from x again needs at most two prescribed flips: the bit in which x differs from π_k, if any, and the bit in which π_k and $\pi_{k'}$ differ. The event of Step 1 creates $\pi_{k'}$ with an uncorrupted evaluation, which stores the value $(k', n-k')$. By Lemma 11(a) the new value is present in P_{t+1}. At most another n successes acquire all front values.

Step 4: expected time. In every generation before T_{val}, conditionally on the history, a success in the sense of Step 2 or Step 3 occurs with probability at least σ. At most $2n$ successes are needed. So T_{val} is dominated by a sum of $2n$ geometric waiting times with parameter σ, and

$$\mathbb{E}[T_{\mathrm{val}}] \;\le\; \frac{2n}{\sigma} \;=\; \frac{8e\mu n^3}{A}.$$

This is the bound of part (i). □

Phase 2: two representatives of every front value. The next lemma keeps up to two representatives of each front value, although not necessarily the same individuals across generations.

Lemma 12. *Let $0 \le k \le n$, and consider a survival step whose combined population contains $m \ge 1$ representatives of k. Their block has exactly $\min\{m, 2\}$ endpoints, all with positive crowding distance, and all of them survive. If at least one offspring is a representative of k, the last endpoint is an offspring.*

Proof. The representatives have rank 1 by Lemma 3. Their common fitness value forms one block in each objective sort, with the same two endpoints. Every interior member has zero crowding distance, whereas both endpoints have positive crowding distance and survive by Lemma 7(b). For a singleton, the unique member is an endpoint and survives. Parents precede offspring in the input to both stable sorts, so the last endpoint is an offspring whenever the block contains an offspring. □

Lemma 13. *Assume the hypotheses of Theorem 6. Then*

$$\mathbb{E}[T_{\mathrm{lock}} - T_{\mathrm{val}}] \;\le\; \frac{4e\mu n(n+1)}{A}.$$

Moreover, for every $t \ge T_{\mathrm{lock}}$:

- *(a) the population P_t consists of exactly two representatives of each of the $n+1$ front values; each survival step selects the two endpoints of every block of R_t, so this pattern remains unchanged;*
- *(b) an offspring whose stored value is not a front value is removed in its first survival step;*
- *(c) in every tournament, every member of P_t is the winner with probability in $[1/(2\mu),\, 2/\mu]$.*

Proof. Step 1: a block gains a second member. Let $t \ge T_{\mathrm{val}}$ and let $x \in P_t$ be a representative of k. It differs from π_k in at most one bit. We use Step 1 of the proof of Theorem 6(i) with S the set of positions in which x and π_k differ ($|S| \le 1$): both samples of the first tournament are x, there is no crossover, mutation flips exactly S, and the evaluation is not corrupted. By (M) this mutation has probability at least $1/(en)$. So one round creates a copy of π_k with an uncorrupted evaluation with probability at least

$$q \;:=\; \frac{1}{\mu^2} \cdot (1-p_c) \cdot \frac{1}{en} \cdot (1-p) \;=\; \frac{A}{e\mu^2 n},$$

and by (R) one generation does so with probability at least $\mu q/4 = A/(4e\mu n)$. The copy is a second member of block k of R_t.

Step 2: two members stay two members. Once block k of R_t has at least two members, its two endpoints survive by Lemma 12. So block k has at least two members in the next population, and by induction in all later ones.

Step 3: the time bound. As long as some block of P_t has a single member, fix such a block. By Step 1 it gains a second member in the current generation with probability at least $A/(4e\mu n)$, conditionally on every history. By Step 2 the number of blocks with at least two members never decreases. Summing the geometric bounds over the $n+1$ blocks gives

$$\mathbb{E}[T_{\mathrm{lock}} - T_{\mathrm{val}}] \;\le\; (n+1) \cdot \frac{4e\mu n}{A}.$$

Step 4: the pattern is permanent, (a) and (b). Consider a survival step at which every one of the $n+1$ blocks of R_t has at least two members. By Lemma 12 each block contributes two rank-one endpoints with positive crowding distance. These are $2(n+1) = \mu$ individuals. By Lemma 7(b) all rank-one individuals with positive crowding distance survive. Hence the next population consists of exactly these μ endpoints, two per block. The next combined population again has at least two members in every block, so the argument repeats. This is (a). An offspring with an off-front stored value is not among the μ endpoints, so it is removed in its first survival step. This is (b).

Step 5: tournaments, (c). All members of P_t have stored rank 1, so a tournament is decided by the stored crowding distances and, on a tie, by a fair coin. We first determine the crowding distances. They were computed at the survival step that produced P_t (for $t = 0$: at the initial sorting call). At that step rank 1 consisted of the $n+1$ front-value blocks, each with at least two members; by Lemma 7(a) there is no further rank-one value. In the sort by the first objective the blocks appear in the order $k = 0, \dots, n$, in the sort by the second objective in the order $k = n, \dots, 0$. An endpoint that is the first or the last element of a sorted list receives crowding distance $+\infty$; we call it an *extreme* endpoint. Both objective ranges in this complete-front layer equal n. Any other endpoint has, in each sort, one neighbour in its own block and one in an adjacent block, so its gap is 1 and its normalized contribution is $1/n$. Its stored crowding distance is therefore $2/n$; we call it a *typical* endpoint. There are exactly four extreme endpoints, the two retained members of each of blocks 0 and n: the two ascending objective sorts reverse the block order but preserve the order within each block.

Now let x be a typical endpoint. It wins the tournament if both samples are x, or if exactly one sample is x, the other sample is one of the other $\mu - 5$ typical endpoints, and the coin favours x:

$$\mathrm{P}[x \text{ wins}] = \frac{1}{\mu^2} + 2 \cdot \frac{1}{\mu} \cdot \frac{\mu - 5}{\mu} \cdot \frac{1}{2} = \frac{\mu - 4}{\mu^2}.$$

Let x be an extreme endpoint. It wins if both samples are x, or if exactly one sample is x and the other sample is a typical endpoint, or if the other sample is one of the other three extreme endpoints and the coin favours x:

$$\mathrm{P}[x \text{ wins}] = \frac{1}{\mu^2} + 2 \cdot \frac{1}{\mu} \cdot \frac{\mu - 4}{\mu} + 2 \cdot \frac{1}{\mu} \cdot \frac{3}{\mu} \cdot \frac{1}{2} = \frac{2\mu - 4}{\mu^2}.$$

Since $\mu \geq 8$, we have $(\mu - 4)/\mu^2 \geq 1/(2\mu)$ and $(2\mu - 4)/\mu^2 \leq 2/\mu$. So the probabilities lie in the claimed interval. □

C.3 Analysis of Phase 3

We keep the hypotheses of Theorem 6 and consider generations $t \geq T_{\mathrm{lock}}$.

How a block changes. Block k of R_t consists of its two members in P_t, followed by all new representatives of k in generation t. Its first and last members survive (Lemmas 12 and 13). A block without new representatives keeps both members. If a new correct representative appears and no new incorrect one, the last member is a correct offspring and survives; if new incorrect representatives also appear, this need not hold. Recall that $B_k(t) = 1$ exactly when neither of the two representatives of k in P_t equals π_k. If $B_k(t) = 0$ for every k, then every π_k is a member of P_t, so $T \leq t$. The event "$B_k = 0$ for every k" implies coverage and is therefore sufficient for an upper bound.

Let $\mathcal{A}_{k,t}$ be the event that generation t creates a new correct representative of k, let $\mathcal{L}_{k,t}$ be the event that it creates a new incorrect representative of k, and define the *repair event*

$$\mathcal{R}_{k,t} := \mathcal{A}_{k,t} \cap \mathcal{L}_{k,t}^{c} \cap \{B_k(t+1) = 0\}.$$

If $B_k(t) = 0$ and $\mathcal{L}_{k,t}$ does not occur, block k of R_t contains at most one incorrect representative. Its two survivors therefore include a correct one. This proves the pathwise inclusion

$$\{B_k(t+1) = 1\} \subseteq \big(\{B_k(t) = 1\} \cap \mathcal{R}_{k,t}^{c}\big) \cup \mathcal{L}_{k,t}. \tag{D}$$

On $\mathcal{A}_{k,t} \cap \mathcal{L}_{k,t}^{c}$, every new representative of k is correct, and there is at least one. Since all of them follow the parents, the last endpoint is correct. Thus $\mathcal{R}_{k,t} = \mathcal{A}_{k,t} \cap \mathcal{L}_{k,t}^{c}$. Consequently, whenever $\mathrm{P}[\mathcal{A}_{k,t} \mid \mathcal{F}_t] \geq r$ and $\mathrm{P}[\mathcal{L}_{k,t} \mid \mathcal{F}_t] \leq \beta$, every history with $B_k(t) = 1$ satisfies

$$\mathrm{P}[\mathcal{R}_{k,t} \mid \mathcal{F}_t] = \mathrm{P}[\mathcal{A}_{k,t} \cap \mathcal{L}_{k,t}^{c} \mid \mathcal{F}_t] \geq r - \beta. \tag{RR}$$

In particular, $\beta \le r/2$ gives a repair probability of at least $r/2$. We use the weaker bound $r/3$ in the estimates below. This estimate uses no independence between $\mathcal{A}_{k,t}$ and $\mathcal{L}_{k,t}$.

One-generation rates. The next lemma collects the three estimates on new representatives that Phase 3 needs. All three hold from every history after T_{lock}, so they apply again at the start of every epoch.

Lemma 14. *Assume the hypotheses of Theorem 6. For every $t \ge T_{\mathrm{lock}}$ and conditionally on every $\mathcal{F}_t$:*

*(a) (*correct representatives*) generation t creates a new correct representative of k with probability at least*

$$r_0 := \frac{A}{8en};$$

*(b) (*incorrect representatives of one value*) for an absolute constant C_B (one may take $C_B = 300$), the probability that generation t creates a new incorrect representative of a fixed k is at most*

$$\beta := \frac{C_B p}{n};$$

*(c) (*incorrect representatives of an interval, one round*) if I is an interval of s values of k and X_I counts the offspring of one reproduction round that are incorrect representatives of some $k \in I$, then for an absolute constant C_I,*

$$\mathbb{E}[e^{\lambda X_I} \mid \mathcal{F}_t] \le \exp\Big\{C_I(e^{2\lambda}-1)\frac{ps}{n^2}\Big\} \qquad (\lambda \ge 0).$$

Proof. Proof of (a). Choose a member of block k of P_t by an $\mathcal{F}_t$-measurable rule. It is either π_k or differs from it in one bit. A fixed round creates a correct representative of k if the following four independent events occur.

- The chosen member wins the first tournament. By Lemma 13(c) the probability is at least $1/(2\mu)$.
- The round skips crossover. The probability is $1 - p_c$.
- Mutation flips exactly the bits in which the chosen member differs from π_k (none or one). By (M) the probability is at least $1/(en)$.
- The evaluation is not corrupted. The probability is $1 - p$.

So one round succeeds with probability at least

$$q \;:=\; \frac{1}{2\mu}\cdot(1-p_c)\cdot\frac{1}{en}\cdot(1-p) \;=\; \frac{A}{2e\mu n}.$$

By (R) the $\mu/2$ rounds of a generation create a correct representative of k with probability at least $\mu q/4 = A/(8en) = r_0$.

Proof of (b) and (c). Fix k, a round, and one of the two children of this round. We first bound the probability that this child is an incorrect representative of k.

Step 1: ideal parents. Both selected parents are members of P_t: the first is a representative of some value i, the second of some value j. Replace them temporarily by the *ideal* parents π_i and π_j. Each actual parent differs from its ideal string in at most one bit, and each position of a child is copied from one parent. So each actual parent changes the crossover child in at most one bit, and the actual pre-mutation child differs from its ideal counterpart in at most two bits. Thus, if the actual pre-mutation child is within distance two of π_k, its ideal counterpart is within distance four.

Step 2: where the ideal child can lie. On the skip branch the ideal child is a parent, π_i or π_j. On the crossover branch with cut c, consider the first child $\pi_i[1..c]\,\pi_j[c{+}1..n]$; the second child is the same with i and j exchanged. Inspect whether each of i, j lies below or above c.

- If $i \ge c$ and $j \le c$, the child is $1^c0^{n-c} = \pi_c$.

- If $i \geq c$ and $j > c$, the child is $1^c 1^{j-c} 0^{n-j} = \pi_j$.
- If $i < c$ and $j \leq c$, the child is $1^i 0^{c-i} 0^{n-c} = \pi_i$.
- In the remaining case $i < c < j$, the child is $w = 1^i 0^{c-i} 1^{j-c} 0^{n-j}$.

In the first three cases the ideal child is one of π_i, π_j, π_c. Since π_a and π_k differ in $|a-k|$ bits, distance at most four from π_k forces $|i-k| \leq 4$, $|j-k| \leq 4$ or $|c-k| \leq 4$. In the remaining case, compare w and π_k on the positions $1..c$ and on the positions $c{+}1..n$:

$$\text{the distance of } w \text{ and } \pi_k \text{ is } \begin{cases} |i-k| + (j-c) & \text{if } k \leq c, \\ (c-i) + |j-k| & \text{if } k > c. \end{cases}$$

So distance at most four from π_k forces $|i-k| \leq 4$ or $|j-k| \leq 4$. Hence in every case the value i or j of a selected parent, or the crossover cut, lies in the window $\{k-4, \ldots, k+4\}$.

Step 3: this has probability $O(1/n)$. The window contains at most 9 values, hence at most 18 members of P_t. By Lemma 13(c) each of them wins a tournament with probability at most $2/\mu$. So a given parent belongs to the window with probability at most $36/\mu \leq 18/n$. The cut c is uniform on $\{1, \ldots, n-1\}$, so it lies in the window with probability at most $9/(n-1) \leq 18/n$. By Steps 1 and 2,

$$\mathrm{P}[\text{the pre-mutation child is within distance two of } \pi_k \mid \mathcal{F}_t] \;\leq\; \frac{18}{n} + \frac{18}{n} + \frac{18}{n} = \frac{54}{n}.$$

Step 4: mutation from distance at least three. Otherwise the pre-mutation child has distance $h \geq 3$ from π_k. Mutation reaches $\mathrm{Ham}_1(\pi_k)$ in one of two ways. It flips exactly $h-1$ of the h differing bits and nothing else, or it flips all h differing bits and exactly one of the $n-h$ other bits. The probability is at most

$$h n^{-(h-1)} (1-1/n)^{n-h+1} + (n-h) n^{-(h+1)} (1-1/n)^{n-h-1} \leq \frac{4}{n^2}.$$

For the inequality, $h n^{-(h-1)}$ is non-increasing in $h \geq 3$ and equals $3/n^2$ at $h = 3$, and $(n-h) n^{-(h+1)} \leq n^{-h} \leq 1/n^2$.

Step 5: one child. By Steps 3 and 4, one child lies in $\mathrm{Ham}_1(\pi_k)$ after mutation with probability at most $54/n + 4/n^2 \leq 58/n$, which is $O(1/n)$. By Lemma 5, such a child becomes a representative of k only if the noise flips the unique bit in which it differs from π_k. This costs another factor p/n, independently of everything before. Hence

$$\mathrm{P}[\text{a given child is an incorrect representative of } k \mid \mathcal{F}_t] \;\leq\; \frac{58p}{n^2}.$$

Step 6: proof of (c). Sum the bound of Step 5 over the s values in I and the two children of the round. The expected round count is

$$\mathbb{E}[X_I \mid \mathcal{F}_t] \;\leq\; \frac{116\,ps}{n^2} = O\Big(\frac{ps}{n^2}\Big).$$

The two siblings are kept as one variable $0 \leq X_I \leq 2$. By convexity of $x \mapsto e^{\lambda x}$ on $[0,2]$,

$$e^{\lambda X_I} \leq 1 + (e^{2\lambda} - 1)\frac{X_I}{2}.$$

Take conditional expectations and use $1 + y \leq e^y$:

$$\mathbb{E}[e^{\lambda X_I} \mid \mathcal{F}_t] \;\leq\; 1 + (e^{2\lambda} - 1)\frac{58\,ps}{n^2} \;\leq\; \exp\Big\{58\,(e^{2\lambda} - 1)\frac{ps}{n^2}\Big\}.$$

This proves (c), for instance with $C_I = 58$.

Step 7: proof of (b). Sum the bound of Step 5 over the $\mu \leq 3n$ children of one generation:

$$\mathrm{P}[\text{generation } t \text{ creates an incorrect representative of } k \mid \mathcal{F}_t] \;\leq\; 3n \cdot \frac{58p}{n^2} = \frac{174\,p}{n} \;\leq\; \frac{300\,p}{n}.$$

This proves (b). □

Epochs. The next lemma is the core of Theorem 6(ii). Its success probability holds from every history after T_{lock}, so a failed epoch can be followed by a new one.

Lemma 15. *There is an absolute constant $c_\star > 0$ with the following property. Fix $p_c \in (0,1)$, let n be sufficiently large, and let*

$$p \;\le\; c_\star \,\frac{p_c(1-p_c)}{\log n}.$$

Then there is an epoch length $T_{\mathrm{ep}} = O(n \log n/(1-p_c))$ such that an epoch of T_{ep} generations, started at any time $t_0 \ge T_{\mathrm{lock}}$, ends with $B_k = 0$ for all k with probability at least $3/4$:

$$\mathrm{P}\big[B_k(t_0+T_{\mathrm{ep}}) = 0 \text{ for all } 0 \le k \le n \;\big|\; \mathcal{F}_{t_0}\big] \;\ge\; \frac{3}{4} \qquad \text{for every } t_0 \ge T_{\mathrm{lock}} \text{ and every } \mathcal{F}_{t_0}.$$

Proof. All probabilities in this proof are conditional on the history $\mathcal{F}_{t_0}$ at the start of the epoch, and all bounds are uniform in this history.

Step 0: parameters. Let

$$D = 2^{\lceil \log_2(K_D \log n)\rceil}, \quad T_1 = \left\lceil \frac{192en\log n}{A} \right\rceil,$$
$$T_2 = \left\lceil \frac{24n\log n}{\rho A} \right\rceil, \qquad \rho = \frac{1}{10368},$$

where K_D is a sufficiently large absolute constant. So D is a power of two with $K_D \log n \le D \le 2K_D \log n$. Let $\mathcal{I}_D$ contain all dyadic prefixes $\{0,\dots,2^a-1\}$ and all dyadic suffixes $\{n-2^a+1,\dots,n\}$ of the index set $\{0,\dots,n\}$ whose power-of-two length 2^a lies in $[D, n+1]$. This family has $O(\log n)$ members. The epoch consists of Stage 1, of length T_1, followed by Stage 2, of length T_2; so $T_{\mathrm{ep}} = T_1 + T_2$. After reducing $c_\star$ if necessary, the assumption gives $p \le 1/2$, so $A \ge (1-p_c)/2$ and $T_{\mathrm{ep}} = O(n\log n/A) = O(n\log n/(1-p_c))$. Also require $c_\star \le \rho/(4C_B)$. Since $\log n \ge 1$ for the values of n under consideration, the noise assumption then gives

$$\beta = \frac{C_B p}{n} \le \frac{2C_B c_\star p_c}{\log n}\frac{A}{n} \le \frac{\rho A}{2n}.$$

This is at most $r_0/2$, and will also be at most half of every rate r_k used in Stage 2. Thus (RR) loses only a factor of three when converting these rates into repair probabilities.

The plan is as follows. In Stage 1 every block becomes good at least once. A bound on the number of new incorrect representatives then guarantees that, throughout Stage 2, most blocks of every interval in $\mathcal{I}_D$ are good (the event $G(t)$ of Step 3). On this event, crossover between correct representatives repairs block k at a rate proportional to $d(k)$. A recursion for $\mathrm{P}[B_k = 1]$ and a harmonic sum finish the proof.

Step 1: in Stage 1 every block becomes good. By Lemma 14(a), (RR), and Step 0, every history with $B_k = 1$ has conditional probability at least $r_0/3$ of making $B_k = 0$ at the next generation boundary. Let σ_k be the first boundary at or after t_0 at which $B_k = 0$; in particular, $\sigma_k = t_0$ if block k starts good. Iterating the conditional bound while the block has not yet been good gives

$$\mathrm{P}[\sigma_k > t_0 + T_1] \le (1 - r_0/3)^{T_1} \le \exp\left(-\frac{A}{24en}\cdot\frac{192en\log n}{A}\right) = n^{-8}.$$

Let $\mathcal{E}_1 := \{\sigma_k \le t_0 + T_1 \text{ for all } k\}$. A union bound over the $n+1$ blocks gives

$$\mathrm{P}[\mathcal{E}_1^c] \le (n+1)n^{-8} \le 2n^{-7} = O(n^{-7}).$$

Step 2: few new incorrect representatives. For an interval I of values let X_I^{ep} be the number of offspring created during both stages that are incorrect representatives of some $k \in I$. The epoch has $(\mu/2)(T_1+T_2) \le Cn^2\log n/A$ reproduction rounds, for an absolute constant C. Apply Lemma 14(c) to each round, conditionally on everything before that round. This is allowed because the rounds of one generation are i.i.d. given $\mathcal{F}_t$, and because the lemma holds from every history after T_{lock}. Iterating over complete reproduction rounds and then over generations (tower property) gives, with an absolute constant C_E,

$$\mathbb{E}e^{\lambda X_I^{\mathrm{ep}}} \le \exp\left\{C_I(e^{2\lambda}-1)\frac{p|I|}{n^2}\cdot\frac{\mu}{2}(T_1+T_2)\right\} \le \exp\left\{C_E(e^{2\lambda}-1)\frac{p\log n}{A}|I|\right\}.$$

Now use the assumed noise bound and Chernoff's inequality, that is, Markov's inequality for $e^{\lambda X_I^{\mathrm{ep}}}$ with $\lambda = 1$:

$$\mathrm{P}[X_I^{\mathrm{ep}} \geq |I|/8] \;\leq\; \exp\!\left\{-\frac{|I|}{8} + C_E(e^2-1)\frac{p\log n}{A}|I|\right\}.$$

By Step 0 and the assumption on p,

$$\frac{p\log n}{A} \;\leq\; \frac{2p\log n}{1-p_c} \;\leq\; 2c_\star p_c \;\leq\; 2c_\star.$$

If $c_\star \leq 1/(32C_E(e^2-1))$, the exponent is at most $-|I|/16$. So

$$\mathrm{P}[X_I^{\mathrm{ep}} \geq |I|/8] \leq e^{-c_E|I|} \qquad \text{with } c_E = \tfrac{1}{16}.$$

Let $\mathcal{E}_2$ be the event that $X_I^{\mathrm{ep}} < |I|/8$ for every $I \in \mathcal{I}_D$. Every $I \in \mathcal{I}_D$ has $|I| \geq D \geq K_D \log n$, so $e^{-c_E|I|} \leq n^{-c_E K_D}$. A union bound over the $O(\log n)$ intervals gives $\mathrm{P}[\mathcal{E}_2^c] = o(n^{-2})$ after increasing K_D.

The event $\mathcal{E}_2$ concerns the whole epoch. We never condition on it, because conditioning would change the one-generation probabilities of Lemma 14; it is only used to bound the probability that G fails (Step 7).

Step 3: most blocks stay good. If block i is good at time t, one of its two members in P_t is π_i and can serve as a correct crossover parent. Let

$$G(t) := \big\{\text{every } I \in \mathcal{I}_D \text{ contains at least } 7|I|/8 \text{ good blocks at time } t\big\}.$$

We claim that on $\mathcal{E}_1 \cap \mathcal{E}_2$ the event $G(t)$ holds at every generation boundary of Stage 2. Indeed, let block k be bad at such a boundary t. On $\mathcal{E}_1$, block k was good at some boundary σ_k during Stage 1, including its initial boundary. Since it is bad at t, it made a good-to-bad transition at a boundary after σ_k. By (D), that transition requires a new incorrect representative of k in the corresponding generation. Map each currently bad block to one such offspring. The map is injective, because an offspring represents one value only. So the number of bad blocks in I at time t is at most $X_I^{\mathrm{ep}} < |I|/8$ on $\mathcal{E}_2$. Consequently every $I \in \mathcal{I}_D$ contains at least $7|I|/8$ good blocks throughout Stage 2, except on the event $\mathcal{E}_1^c \cup \mathcal{E}_2^c$ of probability $o(n^{-2})$.

Step 4: repair rates on $G(t)$. Let $d(k) = \min\{k, n-k\}$ and define

$$r_k := \frac{\rho(1-p)}{n}\left((1-p_c) + p_c d(k)\mathbf{1}_{\{d(k)\geq D\}}\right).$$

We claim: for every generation t of Stage 2 and every history $\mathcal{F}_t$ on which $G(t)$ holds, generation t creates a new correct representative of k with conditional probability at least r_k.

Step 4a: good blocks on both sides. Let $D \leq k \leq n/2$. The largest dyadic prefix of length at most k has length $2^a \in (k/2, k]$. It belongs to $\mathcal{I}_D$ because D is a power of two with $D \leq k$. All its indices are at most $2^a - 1 < k$. On $G(t)$ it contains at least $\frac{7}{8}\cdot\frac{k}{2} \geq k/4$ good blocks below k. The largest dyadic suffix whose length is at most $n-k$ has length $2^b \in ((n-k)/2,\, n-k]$. It lies strictly above k, because its indices are at least $n - 2^b + 1 \geq k+1$. It belongs to $\mathcal{I}_D$ because $k \leq n/2$ gives $2^b > n/4 \geq D$ for large n. On $G(t)$ it contains at least $\frac{7}{8}\cdot\frac{n}{4} \geq n/16$ good blocks there.

Step 4b: parents on both sides of k. Consider a pair of correct representatives π_j with $j > k$ and π_i with $i < k$ in P_t. By Lemma 13(c) a given member of P_t wins a given tournament with probability at least $1/(2\mu)$, and the two tournaments of a round are independent. The pair can be returned in two orders, and these two events are disjoint. For $n \geq 8$ we have $\mu \leq 9n/4$. So the two tournaments return such a pair with probability at least

$$2\cdot\frac{k}{4}\cdot\frac{n}{16}\cdot\frac{1}{(2\mu)^2} = \frac{kn}{128\mu^2} \;\geq\; \frac{kn}{128\cdot(81/16)\,n^2} = \frac{k}{648n}$$

per round.

Step 4c: creating π_k by crossover. Given such a pair, the round creates a correct representative of k if four further independent events occur.

- Crossover is applied. The probability is p_c.

- The cut is at k. The probability is $1/(n-1)$; note that $1 \le D \le k \le n/2 \le n-1$. The child that takes its prefix from π_j and its suffix from π_i is then $1^k 0^{n-k} = \pi_k$. This is the first child if π_j is the first parent, and the second child otherwise.
- Mutation is the identity. The probability is $(1-1/n)^n \ge 1/4$.
- The evaluation is not corrupted. The probability is $1-p$.

So one round succeeds with probability at least

$$q \;:=\; \frac{k}{648n} \cdot p_c \cdot \frac{1}{n-1} \cdot \frac{1}{4} \cdot (1-p).$$

By (R), and since $\mu/(4(n-1)) \ge 1/2$, some round of the generation creates a correct representative of k with probability at least

$$\frac{\mu q}{4} \;=\; \frac{\mu}{4(n-1)} \cdot \frac{p_c(1-p)k}{2592n} \;\ge\; \frac{p_c(1-p)k}{5184n}.$$

The right half $n/2 < k \le n-D$ is symmetric, with $n-k$ in place of k. So for $d(k) \ge D$, crossover creates a correct representative of k with probability at least $p_c(1-p)d(k)/(5184n)$ on $G(t)$.

Step 4d: combining with mutation. By Lemma 14(a), a correct representative of k is created with probability at least $r_0 = (1-p)(1-p_c)/(8en)$, whether or not $G(t)$ holds. Both numbers are lower bounds for the same probability, so this probability is at least their maximum, and hence at least their average. For $d(k) \ge D$ the average is

$$\frac{1-p}{n}\left(\frac{1-p_c}{16e} + \frac{p_c d(k)}{10368}\right) \;\ge\; r_k,$$

because $16e \le 10368 = 1/\rho$. For $d(k) < D$ we have $r_0 \ge \rho(1-p)(1-p_c)/n = r_k$. This proves the claim of Step 4. Note for later that in all cases

$$r_k \;\ge\; \frac{\rho A}{n}.$$

Step 0 gives $\beta \le r_k/2$. By (RR), whenever both $G(t)$ and $B_k(t) = 1$ hold, the repair event $\mathcal{R}_{k,t}$ has conditional probability at least $r_k/3$.

Step 5: a recursion. Measure the following times from the beginning of Stage 2, so t runs from 0 to T_2. Let τ be the first t at which $G(t)$ fails, and put $x_{k,t} := \mathrm{P}[B_k(t) = 1,\ \tau > t]$. The event $\{\tau > t\}$ is determined by $\mathcal{F}_t$, and on it $G(t)$ holds. By (D),

$$\{B_k(t+1) = 1,\ \tau > t+1\} \;\subseteq\; \big(\{B_k(t) = 1,\ \tau > t\} \cap \mathcal{R}_{k,t}^c\big) \cup \mathcal{L}_{k,t}.$$

Condition on $\mathcal{F}_t$. The first event has conditional probability at most $\mathbf{1}_{\{B_k(t)=1,\,\tau>t\}}(1-r_k/3)$ by Step 4. The second event has conditional probability at most β by Lemma 14(b); this uses neither $G(t)$ nor any conditioning on $\mathcal{E}_2$. Taking expectations gives a recursion, and unrolling it gives the bound at time T_2:

$$x_{k,t+1} \le (1-r_k/3)x_{k,t} + \beta,$$
$$x_{k,T_2} \le (1-r_k/3)^{T_2} + \beta \sum_{s \ge 0} (1-r_k/3)^s \le e^{-r_k T_2/3} + \frac{3\beta}{r_k}.$$

Step 6: summing over the blocks. Since $r_k \ge \rho A/n$, we have $r_k T_2/3 \ge 8 \log n$ and $e^{-r_k T_2/3} \le n^{-8}$. So the exponential terms sum to at most $(n+1)n^{-8} \le 2n^{-7}$. For the second terms, at most $2D$ indices k have $d(k) < D$, and for them $r_k \ge \rho(1-p)(1-p_c)/n$. Every value $d \ge D$ equals $d(k)$ for at most two indices k, and for them $r_k \ge \rho(1-p)p_c d/n$. With $\beta = C_B p/n$ this gives

$$\sum_{k=0}^{n} \frac{3\beta}{r_k} \;\le\; \frac{3C_B p}{\rho(1-p)}\left(\frac{2D}{1-p_c} + \frac{2}{p_c}\sum_{d=D}^{n}\frac{1}{d}\right).$$

Integral comparison gives $\sum_{d=D}^{n} 1/d \le 1 + \ln n$. Moreover $D \le 2K_D \log n$, $1 - p \ge 1/2$, and both $1/(1-p_c)$ and $1/p_c$ are at most $1/(p_c(1-p_c))$. Hence, for an absolute constant C,

$$\sum_{k=0}^{n} \frac{3\beta}{r_k} \le C \, \frac{p \log n}{p_c(1-p_c)} \;\le\; C\, c_\star.$$

Choose $c_\star$ so that this is at most $1/8$, the preceding Chernoff bound (Step 2) holds, and the condition $\beta \le \rho A/(2n)$ of Step 0 holds.

Step 7: conclusion. If some block is bad at the end of the epoch, then either G failed during Stage 2 ($\tau \le T_2$), or $B_k(T_2) = 1$ and $\tau > T_2$ for some k. By Step 3, $\mathrm{P}[\tau \le T_2] \le \mathrm{P}[\mathcal{E}_1^c] + \mathrm{P}[\mathcal{E}_2^c]$. Adding this probability and applying a union bound over k gives

$$\mathrm{P}[\exists k :\; B_k = 1 \text{ at the end of the epoch}] \;\le\; \underbrace{2n^{-7} + o(n^{-2})}_{\mathrm{P}[\tau \le T_2]} + \underbrace{2n^{-7} + \tfrac{1}{8}}_{\sum_k x_{k,T_2}} \;\le\; \frac{1}{4}$$

for all sufficiently large n. This proves an epoch success probability of at least $3/4$. □

Proof of Theorem 6(ii). Let $\tau_3 := \min\{s \ge 0 : B_k(T_{\mathrm{lock}} + s) = 0 \text{ for all } k\}$. Pointwise $T \le T_{\mathrm{lock}} + \tau_3$, since the front is covered as soon as all B_k vanish. The epoch of Lemma 15 has length $T_{\mathrm{ep}} = O(n \log n/(1-p_c))$ and succeeds with probability at least $3/4$ from every history after T_{lock}. The time T_{lock} is a stopping time and is finite almost surely by Theorem 6(i) and Lemma 13. On each event $\{T_{\mathrm{lock}} = s\}$, apply this epoch bound successively at $s, s + T_{\mathrm{ep}}, s + 2T_{\mathrm{ep}}, \ldots$. The tower property, followed by summation over s, gives $\mathrm{P}[\tau_3 > jT_{\mathrm{ep}} \mid \mathcal{F}_{T_{\mathrm{lock}}}] \le 4^{-j}$ for every integer $j \ge 0$. Summing this tail bound gives

$$\mathbb{E}[\tau_3 \mid \mathcal{F}_{T_{\mathrm{lock}}}] \le \frac{4}{3} T_{\mathrm{ep}} = O\Big(\frac{n \log n}{1-p_c}\Big).$$

Theorem 6(i) and Lemma 13 give

$$\mathbb{E}[T_{\mathrm{lock}}] \le \frac{8e\mu n^3}{A} + \frac{4e\mu n(n+1)}{A}.$$

Adding the bounds by the tower property proves the stated estimate. After reducing $c_\star$ we have $p \le 1/2$, so $A \ge (1-p_c)/2$ and the total is $O(n^4/(1-p_c))$. Without reevaluation the algorithm performs μ evaluations at initialization and μ in each generation, so $T_{\mathrm{eval}} = \mu + \mu T$. □

D Proof of Theorem 2: Noise Upper Bound for the Reevaluation Strategy

Theorem 7. *Fix a constant $p_c \in (0,1)$ and let $\mu = 2n + 2$. With $c_A = \min\{\frac{1}{480e}, \frac{p_c}{1105920}\}$ and $C_A = 20 \max\{\frac{32e}{1-p_c}, \frac{4608e}{p_c}\}$, for every $n \ge 3$ and $p \le c_A(1-p_c)/(\mu n)$, the reevaluating NSGA-II satisfies, uniformly over every initial state, $\mathbb{E}[T] \le C_A n^2$ generations.*

This implies Theorem 2 with $c' = c_A(1-p_c)/3$, since $\mu \le 3n$ makes $p \le c'/n^2$ sufficient for the displayed noise condition. The strategy is reevaluation throughout, $p_c \in (0,1)$ is constant, $\mu = 2n + 2$ (so $2n \le \mu \le 3n$ for $n \ge 2$) and $p \le c_A(1-p_c)/(\mu n)$, whence

$$2\mu p \;\le\; \frac{2c_A}{n} \;\le\; \frac{1}{2}.$$

In each generation, the *variation step* creates the offspring Q_t, and the *survival step* evaluates the combined population R_t and selects P_{t+1}. On a corruption-free survival step, ranking and normalized crowding are computed from the true fitness values of R_t. The resulting ranks and crowding distances are used by the next generation's tournaments; the current variation step may still use ranks and crowding distances computed from corrupted evaluations. Sections D.1 and D.3 use the noise-free progress arguments of Bian & Qian (2022), with the crossover probability kept as a parameter. What noise adds is the bookkeeping of Section D.2: a corrupted evaluation can remove a Pareto string, reduce a best objective value or invalidate the data the next generation's tournaments

rely on, and each of these costs at most $O(\mu p) = O(1/n)$ per generation, against a progress rate $\Omega(1/n)$ with a larger constant. An additive drift argument on $\Psi_t = \Phi_1(P_t) + 2Z_t + \beta_t$ concludes (Section D.4).

We use the terminology of Appendix B (fitness value, representative, layer 1, block). Let $Y_t := \{k : \pi_k \in P_t\}$ be the set of present slots and $Z_t := n + 1 - |Y_t|$ the deficit; slots outside Y_t are missing. The *true fitness value* of an individual z is $g(z)$; its observed fitness value $\tilde{v}(z)$ may differ from it, and for $y \in R_t$ we write (B_y, I_y) for the noise pair of the evaluation of y in the survival step. The number of corrupted evaluations in the survival step of generation t is $N_t \sim \mathrm{Bin}(2\mu, p)$, and $\mathcal{C}_t := \{N_t = 0\}$ is the event that this survival step is corruption-free; on $\mathcal{C}_t$ every fitness value equals the true value of the individual. Every survivor stores its layer index (*stored rank*) and its crowding distance (*stored crowding*), which are the data the tournaments of the next generation use. To record whether these data came from a corrupted survival step, set $\beta_0 := 1$ and, for $t \geq 1$, $\beta_t := \mathbf{1}\{\text{the survival step of generation } t-1 \text{ had at least one corruption}\}$, so that $\beta_{t+1} = 0$ exactly on $\mathcal{C}_t$. Finally $\Phi_1(P) := 2n - \max_{x\in P} \mathrm{LO}(x) - \max_{x\in P} \mathrm{TZ}(x) \in [0, 2n]$, which vanishes if and only if $\pi_n \in P$ and $\pi_0 \in P$, because $\mathrm{LO}(x) = n$ forces $x = 1^n$ and $\mathrm{TZ}(x) = n$ forces $x = 0^n$. We write $x^+ := \max\{x, 0\}$ and $\Delta X := X_{t+1} - X_t$.

D.1 Ranks and Crowding Distances after a Corruption-Free Generation

In this subsection we fix a generation t and work on the event $\mathcal{C}_t$, on which every fitness value equals the true value of the individual.

Lemma 16. *Consider generation t of the reevaluating NSGA-II. On the event $\mathcal{C}_t$,*

$$\max_{x\in P_{t+1}} \mathrm{LO}(x) \;\geq\; \max_{y\in R_t} \mathrm{LO}(y) \;\geq\; \max_{x\in P_t} \mathrm{LO}(x).$$

The same holds with TZ *in place of* LO.

Proof. Step 1: a candidate value. Let $L' := \max_{y\in R_t} \mathrm{LO}(y)$, and let T' be the maximal TZ value among the LO-maximisers of R_t. Put $\hat{v} := (L', T')$.

Step 2: no true value strictly dominates $\hat{v}$. A strict dominator among the true values in R_t would either exceed L' in the first coordinate, which is impossible by the maximality of L', or it would equal L' in the first coordinate and exceed T' in the second, which is impossible by the choice of T'.

Step 3: an individual with fitness value $\hat{v}$ survives. On $\mathcal{C}_t$ all fitness values are true values. So $\hat{v}$ is a layer-1 fitness value. By Lemma 7(b) some individual with fitness value $\hat{v}$ survives. Its true value is $\hat{v}$, so $\max_{x\in P_{t+1}} \mathrm{LO}(x) \geq L'$. This is the first inequality. The second inequality holds because $P_t \subseteq R_t$. The argument for TZ is symmetric. □

The next lemma describes the ranks and crowding distances stored after a corruption-free survival step and used by the next generation's tournaments.

Lemma 17. *Consider generation t of the reevaluating NSGA-II. On the event $\mathcal{C}_t$, the stored state at time $t+1$ has the following properties.*

- *(a) If $\pi_n \in P_{t+1}$, then some position holding π_n has stored rank 1 and stored crowding $+\infty$. The same holds for π_0.*
- *(b) Every position holding any Pareto string has stored rank 1.*
- *(c) Some position holding a string of maximal* LO *value in P_{t+1} has stored rank 1 and stored crowding $+\infty$. The same holds for* TZ*.*

Proof. On $\mathcal{C}_t$ all fitness values are true values. We prove (b) first.

Step 1: part (b). A front value is never strictly dominated by any true value (Lemma 3). So every copy of a Pareto string is in layer 1. The stored rank of a survivor is its layer index.

Step 2: part (a). Assume that $\pi_n \in P_{t+1}$. Then R_t contains a copy of π_n. On $\mathcal{C}_t$ this copy has fitness value $(n, 0)$, and it lies in layer 1 by Step 1. The value $(n, 0)$ is the unique layer-1 fitness value with first coordinate n. Indeed, any other fitness value (n, b) would have $b \geq 1$ and coordinate

sum larger than n. Hence, in the LO-sort of layer 1, the block of $(n, 0)$ is last. The last element of the sort is an endpoint of that block, and it receives the contribution $+\infty$. On $\mathcal{C}_t$ all representatives of $(n, 0)$ are copies of π_n (Lemma 5). The endpoint has positive crowding distance, so it survives (Lemma 7(b)). It has stored rank 1 and stored crowding $+\infty$. The argument for π_0 is symmetric.

Step 3: part (c). Let $\hat{v} = (L', T')$ be as in the proof of Lemma 16. Every fitness value with first coordinate L' other than $\hat{v}$ has a second coordinate smaller than T', and is strictly dominated by $\hat{v}$. Hence $\hat{v}$ is the unique layer-1 fitness value with first coordinate L', and no layer-1 fitness value has a larger first coordinate. As in Step 2, the LO-sort of layer 1 ends inside the block of $\hat{v}$. Its end element receives $+\infty$ and survives. Fitness values are true values, so this element has true LO value L'. By Lemma 16, $L' = \max_{x \in P_{t+1}} \mathrm{LO}(x)$. The argument for TZ is symmetric. □

D.2 Effects of Corrupted Evaluations

Now the noise is arbitrary. We fix a generation t and condition on $\mathcal{F}_t$ and on the offspring Q_t. The remaining randomness is the fresh evaluation noise of the survival step and any tie-breaking in crowding-based truncation. All pointwise bounds below hold for every realization of these random choices. Noise can remove Pareto strings, reduce the best LO and TZ values, and change the stored ranks and crowding distances. The next three lemmas bound these three costs in turn, each by a constant multiple of μp, which up to a factor 2 is the expected number of corruptions in a generation.

Lemma 18. *Consider generation t of the reevaluating NSGA-II. Pointwise, that is, for every outcome, $|Y_t \setminus Y_{t+1}| \le N_t$. In particular,*

$$\mathbb{E}\big[\, |Y_t \setminus Y_{t+1}| \,\big|\, \mathcal{F}_t \big] \;\le\; 2\mu p.$$

Proof. We exhibit an injection from the lost slots to the corrupted evaluations of the survival step. Let $k \in Y_t \setminus Y_{t+1}$ be a lost slot. Then every copy of π_k in R_t failed to survive.

Case (a): some copy of π_k was corrupted. We map k to the first such copy in position order. This image did not survive. If it had survived, the slot would be present.

Case (b): every copy of π_k has its true fitness value $(k, n-k)$. By Lemma 3 each such copy is in layer 1. By Lemma 7(b) at least one representative of k survives. That survivor is not a copy of π_k, because the slot was lost. So by Lemma 5 it is an incorrect representative, and its evaluation was corrupted. We map k to the first surviving incorrect representative of k.

Injectivity. Within case (a), distinct slots have disjoint sets of copies. Within case (b), a position carries one fitness value, and the front values $(k, n-k)$ are distinct for distinct k. Across the cases, the images of case (a) did not survive and the images of case (b) did.

Conclusion. Every image is a corrupted evaluation, and there are N_t corrupted evaluations. This proves the pointwise bound. The bound on the expectation follows from $\mathbb{E}[N_t \mid \mathcal{F}_t] = 2\mu p$. □

Lemma 19. *Consider generation t of the reevaluating NSGA-II, and let $L := \max_{y \in R_t} \mathrm{LO}(y)$ be the largest true* LO *value in R_t.*

(a) The best LO *value is lost with small probability:*

$$\mathrm{P}\Big[\max_{x \in P_{t+1}} \mathrm{LO}(x) < L \;\Big|\; \mathcal{F}_t,\, Q_t\Big] \;\le\; p + \frac{4\mu p}{n}.$$

The symmetric statement holds for TZ.

(b) Consequently,

$$\mathbb{E}\big[(\Phi_1(P_{t+1}) - \Phi_1(P_t))^+ \,\big|\, \mathcal{F}_t\big] \;\le\; 2n\Big(p + \frac{4\mu p}{n}\Big) \;\le\; \frac{9c_A(1-p_c)}{n}.$$

Proof. Step 1: a lexicographically maximal individual and two bad events. Let $\hat{x}$ be a member of R_t that is maximal for the lexicographic order (LO first, then TZ) on true values. So $g(\hat{x}) = (L, \hat{T})$, where $\hat{T}$ is maximal among the LO-maximisers. We define two bad events.

The first bad event is $E_1 := \{\hat{x}\text{'s evaluation is corrupted}\}$. We have $\mathrm{P}[E_1] \le p$.

The second bad event E_2 says that the evaluation of some individual flips a productive bit and the flipped string has a value of at least $(L, \hat{T})$:

$$E_2 := \big\{\exists y \in R_t \; \exists i \in \mathcal{P}(y) : \; (B_y, I_y) = (1, i) \text{ and } g(y \oplus e_i) \succeq (L, \hat{T})\big\}.$$

Each of the 2μ individuals contributes at most 2 candidate flips, and each candidate flip is realised with probability p/n. So

$$\mathrm{P}[E_2] \;\le\; 2 \cdot 2\mu \cdot \frac{p}{n} \;=\; \frac{4\mu p}{n}.$$

Step 2: claim. Off $E_1 \cup E_2$, some survivor has true LO *value L.* Off E_1, the individual $\hat{x}$ has fitness value $(L, \hat{T})$. Consider any fitness value $r = \tilde{v}(y)$ with $r \succeq (L, \hat{T})$. There are three cases.

- *The evaluation of y did not change its value.* Then $g(y) = r \succeq (L, \hat{T})$, and the maximality of $(L, \hat{T})$ forces $g(y) = (L, \hat{T})$. In particular y has true LO value L.
- *The evaluation of y flipped a non-productive bit.* Then $r \preceq g(y)$ by Lemma 9. So again $g(y) \succeq (L, \hat{T})$, which forces $g(y) = (L, \hat{T})$. Now $(L, \hat{T}) = g(y) \succeq r \succeq (L, \hat{T})$ gives $r = g(y)$. The value did not change, and we are back in the previous case.
- *The evaluation of y flipped a productive bit*, and $r \succeq (L, \hat{T})$. This is exactly an E_2-event, which is excluded.

Hence, off $E_1 \cup E_2$, every fitness value $\succeq (L, \hat{T})$ equals $(L, \hat{T})$, and every individual with this fitness value has true value $(L, \hat{T})$. In particular no fitness value strictly dominates $(L, \hat{T})$. So $(L, \hat{T})$ is a layer-1 fitness value, and $\hat{x}$ has it. Some individual with this fitness value survives (Lemma 7(b)), and that survivor has true LO value L.

Step 3: part (a). By Step 2, the event $\{\max_{x \in P_{t+1}} \mathrm{LO}(x) < L\}$ is contained in $E_1 \cup E_2$. By Step 1 its conditional probability is at most $p + \frac{4\mu p}{n}$. The argument for TZ is symmetric.

Step 4: part (b). The potential Φ_1 increases only if the maximal LO value or the maximal TZ value drops. A drop of either maximum is at most n. Since $L \ge \max_{x \in P_t} \mathrm{LO}(x)$, a drop of the maximal LO value implies $\max_{x \in P_{t+1}} \mathrm{LO}(x) < L$, and similarly for TZ. With part (a) this gives the first inequality of (b). For the second inequality we use the hypothesis $p \le c_A(1 - p_c)/(\mu n)$ and then $\mu \ge 2n$:

$$2np + 8\mu p \;\le\; \frac{2c_A(1-p_c)}{\mu} + \frac{8c_A(1-p_c)}{n} \;\le\; \frac{9c_A(1-p_c)}{n}.$$

This proves the lemma. □

Lemma 20. *Consider generation t of the reevaluating NSGA-II.*

(a) The event $\mathcal{C}_t$ has probability

$$\mathrm{P}[\mathcal{C}_t \mid \mathcal{F}_t, Q_t] \;=\; (1-p)^{2\mu} \;\ge\; 1 - 2\mu p \;\ge\; \frac{1}{2}.$$

(b) The event $\mathcal{C}_t$ is independent of Q_t given $\mathcal{F}_t$.

(c) On $\mathcal{C}_t$, the stored state at time $t+1$ satisfies Lemma 17. In particular $\beta_{t+1} = 0$.

Proof. The evaluation noise of the survival step consists of 2μ independent $\mathrm{Bernoulli}(p)$ corruption indicators, drawn independently of $(\mathcal{F}_t, Q_t)$. The event $\mathcal{C}_t$ says that all of them are 0. This gives the equality in (a) and the independence in (b). Bernoulli's inequality gives $(1-p)^{2\mu} \ge 1 - 2\mu p$, and $2\mu p \le \frac{1}{2}$ holds by the hypothesis on p. Part (c) is Lemma 17 together with the definition of β. □

D.3 Progress after a Corruption-Free Generation

We call the state at time t *good* if $\beta_t = 0$; then the survival step of generation $t-1$ was corruption-free, so the ranks and crowding distances used by the tournaments of generation t satisfy Lemma 17. The next two lemmas describe progress: while an extreme solution is missing, mutation can improve the best objective value; once both extremes are present, crossover can fill a gap in the front.

Lemma 21. *Consider generation t of the reevaluating NSGA-II, and let $c_2 := \frac{1}{8e}$. If the state at time t is good and $\Phi_1(P_t) > 0$, then*

$$\mathrm{P}\big[\Phi_1(P_{t+1}) \le \Phi_1(P_t) - 1 \,\big|\, \mathcal{F}_t\big] \;\ge\; \frac{c_2(1-p_c)}{2n}.$$

Proof. Assume without loss of generality that $L := \max_{x\in P_t} \mathrm{LO}(x) < n$. If only the TZ side is deficient, the argument is symmetric.

Step 1: an individual with maximal LO *and infinite crowding distance.* By Lemma 17(c) some position holds a string x_L with $\mathrm{LO}(x_L) = L$, stored rank 1 and stored crowding $+\infty$.

Step 2: one round. Consider any fixed round of the variation step. We call it a success if the following three events occur.

- The first tournament returns the position of x_L. This has probability at least $\frac{1}{2\mu}$ (Lemma 8(a)).
- The round skips crossover. This has probability $1 - p_c$.
- Mutation flips exactly bit $L+1$ of x_L and no other bit. Bit $L+1$ is 0 because $\mathrm{LO}(x_L) = L$. This has probability $\frac{1}{n}(1-\frac{1}{n})^{n-1} \ge \frac{1}{en}$.

The resulting offspring y' has $\mathrm{LO}(y') \ge L+1$. Hence a deterministic lower bound on the per-round success probability is

$$\alpha := \frac{1-p_c}{2e\mu n} \le \frac{2}{\mu}.$$

Step 3: all rounds. The $\mu/2$ rounds are independent conditional on $\mathcal{F}_t$. For any deterministic per-round lower bound $\alpha \in [0, 2/\mu]$, we have $(\mu/2)\alpha \le 1$, and hence

$$\mathrm{P}[\text{some round succeeds}] \;\ge\; 1-(1-\alpha)^{\mu/2} \;\ge\; \frac{(\mu/2)\alpha}{1+(\mu/2)\alpha} \;\ge\; \frac{\mu\alpha}{4}, \tag{1}$$

and with Step 2 the right-hand side is at least $\frac{1-p_c}{8en}$.

Step 4: a corruption-free survival step. The event $\mathcal{C}_t$ is independent of Q_t given $\mathcal{F}_t$ and has probability at least $\frac{1}{2}$ (Lemma 20). On $\mathcal{C}_t$, Lemma 16 applies and gives

$$\max_{x\in P_{t+1}} \mathrm{LO}(x) \;\ge\; \max_{y\in R_t} \mathrm{LO}(y) \;\ge\; L+1 \qquad \text{and} \qquad \max_{x\in P_{t+1}} \mathrm{TZ}(x) \;\ge\; \max_{x\in P_t} \mathrm{TZ}(x).$$

So Φ_1 decreases by at least 1. The product of the two probabilities is $\frac{1-p_c}{8en} \cdot \frac{1}{2} = \frac{c_2(1-p_c)}{2n}$, which is the claim. □

Lemma 22. *Consider generation t of the reevaluating NSGA-II. Let $G_t := |Y_{t+1} \setminus Y_t|$ be the number of slots gained, and let $\gamma_1 := \frac{1}{36864}$. If the state at time t is good, $\Phi_1(P_t) = 0$ and $Z_t \ge 1$, then*

$$\mathrm{P}[G_t \ge 1 \mid \mathcal{F}_t] \;\ge\; \frac{\gamma_1 p_c}{n}.$$

Proof. Step 1: setup. Since $\Phi_1(P_t) = 0$, we have $\pi_n, \pi_0 \in P_t$. Each of the two strings has a position with stored rank 1 and stored crowding $+\infty$ (Lemma 17(a)). Both extremes are present and $Z_t \ge 1$, so some interior slot is missing:

$$Z^{\mathrm{int}} \;:=\; \big|\{1,\dots,n-1\} \setminus Y_t\big| \;\ge\; 1.$$

Let $i := (n-1) - Z^{\mathrm{int}}$ be the number of interior slots that are present. We distinguish two cases, depending on i. In both cases we bound the probability that a fixed round of the variation step creates a missing Pareto string, and then we apply (1).

Step 2: the case $i < (n-1)/2$, crossover of the two extremes. In a fixed round, consider the following events.

- The two tournaments return these positions of π_n and π_0, in this order. By Lemma 8(a) and independence of the two tournaments, this has probability at least $(\frac{1}{2\mu})^2$.

- Crossover is applied. This has probability p_c.
- The cut hits a missing interior slot. By Lemma 10(ii) the first child is π_c, where the cut c is uniform. So this has probability $\frac{Z^{\text{int}}}{n-1} \geq \frac{1}{2}$.
- Mutation leaves that child unchanged. This has probability $(1-\frac{1}{n})^n \geq \frac{1}{4}$ for $n \geq 2$.

The per-round probability is therefore at least $\alpha_E := p_c/(32\mu^2) \leq 2/\mu$. Over the $\mu/2$ rounds, (1) gives a probability of at least

$$\frac{\mu\alpha_E}{4} \;=\; \frac{p_c}{128\mu} \;\geq\; \frac{p_c}{384n}.$$

Step 3: the case $i \geq (n-1)/2$, crossover of an extreme and an interior Pareto string. Fix positions e_n holding π_n and e_0 holding π_0, each with stored rank 1 and crowding distance $+\infty$. Write w_n, w_0 for their respective probabilities of winning the first tournament; each is at least $1/(2\mu)$ by Lemma 8(a). Let T be the set of positions holding interior Pareto strings, and let q_x be the probability that its position $x \in T$ wins the second tournament. These positions all have stored rank 1, so Lemma 8(b) gives $\sum_{x\in T} q_x \geq (|T|/\mu)^2 \geq i^2/\mu^2$.

Let $M := \{1,\dots,n-1\} \setminus Y_t$ be the missing interior slots. For $x \in T$, write $k = k(x)$ for the slot it holds and put $a_k := |M \cap \{k+1,\dots,n-1\}|$ and $b_k := |M \cap \{1,\dots,k-1\}|$. Since k is present, $a_k + b_k = |M| = Z^{\text{int}}$. With ordered parents (e_n, x), each of the a_k cuts produces the missing string as the first child. With ordered parents (e_0, x), each of the b_k cuts produces it as the second child. In each case the designated child is unchanged by mutation with probability at least $1/4$. The parent choices and cuts in this sum are disjoint events. Independence of the two tournaments, the crossover coin, the cut and mutation therefore gives the following bound on the per-round success probability s_D, without assuming uniform parent selection:

$$\begin{aligned} s_D &\geq \frac{p_c}{4(n-1)} \sum_{x\in T} q_x\big(w_n a_{k(x)} + w_0 b_{k(x)}\big) \\ &\geq \frac{p_c Z^{\text{int}}}{8\mu(n-1)} \sum_{x\in T} q_x \\ &\geq \frac{p_c i^2 Z^{\text{int}}}{8\mu^3(n-1)}. \end{aligned}$$

For convenience we weaken this bound by a factor of four and use the valid per-round lower bound

$$\alpha_D := \frac{p_c i^2 Z^{\text{int}}}{32\mu^3(n-1)} \leq \frac{1}{32\mu} \leq \frac{2}{\mu},$$

where $i \leq \mu$ and $Z^{\text{int}} \leq n-1$ justify the first upper bound. We insert $i \geq \frac{n-1}{2}$, $Z^{\text{int}} \geq 1$ and $\mu \leq 3n$. Over the $\mu/2$ rounds, (1) gives a probability of at least

$$\frac{\mu\alpha_D}{4} \;\geq\; \frac{\mu}{4} \cdot \frac{\frac{(n-1)^2}{4}\, p_c}{32\mu^3(n-1)} \;=\; \frac{p_c(n-1)}{512\mu^2} \;\geq\; \frac{p_c}{18432n}.$$

This is $2\gamma_1 p_c/n$.

Step 4: a corruption-free survival step. In either case, the variation step creates some missing Pareto string π_j with probability at least $2\gamma_1 p_c/n$. The event $\mathcal{C}_t$ is independent of Q_t given $\mathcal{F}_t$ and has probability at least $\frac{1}{2}$ (Lemma 20). On $\mathcal{C}_t$ the following holds. The value $(j, n-j)$ is a layer-1 fitness value (Lemma 3). Some representative of j survives (Lemma 7(b)). On $\mathcal{C}_t$ all representatives of j are copies of π_j (Lemma 5). Hence $\pi_j \in P_{t+1}$ and $G_t \geq 1$. Multiplying the two probabilities gives the claim. □

D.4 Proof of Theorem 7

Proof of Theorem 7. Step 1: the potential. Define

$$\Psi_t \;:=\; \Phi_1(P_t) + 2Z_t + \beta_t \;\in [0, 4n+3] \qquad \text{and} \qquad \tau \;:=\; \inf\{t : Z_t = 0\} = T.$$

Fix $t < \tau$ and condition on $\mathcal{F}_t$. All expectations in Steps 2 and 3 are conditional on $\mathcal{F}_t$.

Step 2: the expected increase, in every state. In *every* state, the expected increase of Ψ is bounded by three terms:

$$\mathbb{E}\big[(\Delta\Phi_1)^+\big] \;\le\; \frac{9c_A(1-p_c)}{n} \qquad \text{(Lemma 19)},$$

$$\mathbb{E}\big[2(\Delta Z)^+\big] \;\le\; 2\,\mathbb{E}[N_t] \;=\; 4\mu p \;\le\; \frac{4c_A(1-p_c)}{n} \qquad \text{(Lemma 18)},$$

$$\mathbb{E}\big[(\Delta\beta)^+\big] \;\le\; \mathrm{P}[N_t \ge 1] \;\le\; 2\mu p \;\le\; \frac{2c_A(1-p_c)}{n} \qquad \text{(Lemma 20)}.$$

For the second line, note that $\Delta Z = |Y_t \setminus Y_{t+1}| - |Y_{t+1} \setminus Y_t|$, so $(\Delta Z)^+ \le |Y_t \setminus Y_{t+1}| \le N_t$. For the third line, note that $\beta_{t+1} = 1$ requires $N_t \ge 1$. The last inequality of each of the two lines uses $p \le c_A(1-p_c)/(\mu n)$. The three bounds add up to $\frac{15c_A(1-p_c)}{n}$. We now distinguish three cases.

Step 3, case (i): $\beta_t = 0$ and $\Phi_1(P_t) > 0$. The success event of Lemma 21 decreases Φ_1 by at least 1, and all increases are already accounted for by the bounds of Step 2. So

$$\mathbb{E}[\Delta\Psi_t] \;\le\; -\frac{c_2(1-p_c)}{2n} + \frac{15c_A(1-p_c)}{n} \;\le\; -\frac{c_2(1-p_c)}{4n} \qquad \text{once } c_A \le \frac{c_2}{60}.$$

Step 3, case (ii): $\beta_t = 0$ and $\Phi_1(P_t) = 0$. Then $Z_t \ge 1$ because $t < \tau$. Since $\Delta Z \le N_t - \mathbf{1}\{G_t \ge 1\}$, we have $\mathbb{E}[2\Delta Z] \le 4\mu p - 2\mathrm{P}[G_t \ge 1]$. We bound the increases of Φ_1 and of β by Step 2, and we bound $\mathrm{P}[G_t \ge 1]$ from below by Lemma 22. This gives

$$\begin{aligned}\mathbb{E}[\Delta\Psi_t] \;&\le\; \frac{9c_A(1-p_c)}{n} + \frac{2c_A(1-p_c)}{n} + \frac{4c_A(1-p_c)}{n} - \frac{2\gamma_1 p_c}{n} \\ &\le\; -\frac{\gamma_1 p_c}{n} \qquad \text{once } c_A \le \frac{\gamma_1 p_c}{30}.\end{aligned}$$

Step 3, case (iii): $\beta_t = 1$. By Lemma 20, $\mathrm{P}[\beta_{t+1} = 0] \ge \frac{1}{2}$. So $\mathbb{E}[\Delta\beta] \le -\frac{1}{2}$, and

$$\mathbb{E}[\Delta\Psi_t] \;\le\; \frac{13c_A(1-p_c)}{n} - \frac{1}{2} \;\le\; -\frac{1}{4} \qquad \text{for } n \ge 3,$$

since $c_A \le 1/(480e)$ gives $13c_A(1-p_c)/n < 1/4$.

Step 4: the drift theorem. Let

$$c_A \;:=\; \min\Big\{\frac{c_2}{60},\, \frac{\gamma_1 p_c}{30}\Big\} \qquad \text{and} \qquad \delta \;:=\; \frac{\min\{c_2(1-p_c),\, 4\gamma_1 p_c\}}{4n} \wedge \frac{1}{4}.$$

With $c_2 = \frac{1}{8e}$ and $\gamma_1 = \frac{1}{36864}$ this is the constant c_A of the theorem. By Step 3 the hypothesis of Lemma 1 holds on $\{t < \tau\}$ with this δ and with $B = 4n + 3 \le 5n$. We obtain

$$\begin{aligned}\mathbb{E}[T] \le \frac{5n}{\delta} &= 20n^2 \max\Big\{\frac{8e}{1-p_c}, \frac{9216}{p_c}\Big\} \\ &\le 20n^2 \max\Big\{\frac{32e}{1-p_c}, \frac{4608e}{p_c}\Big\} = C_A n^2.\end{aligned}$$

The cap $1/4$ in δ is inactive for $n \ge 3$. The final inequality merely rounds up the constants, using $8e \le 32e$ and $9216 \le 4608e$. All three drift cases hold uniformly over the stored state, including $\beta_0 = 1$, so the same bound already covers arbitrary initialization. □

E Proof of Theorem 3: Noise Lower Bound for the Reevaluation Strategy

This appendix proves the reevaluation lower bound of Theorem 3.

Theorem 8. *There is an absolute constant $c_L > 0$, and for every constant $p_c \in [0,1)$ there are constants $p_{\max} = p_{\max}(p_c) > 0$, C_2 and n_0 such that for all $n \ge n_0$ and every p with $C_2/n \le p^2 \le p_{\max}^2$, the reevaluating NSGA-II with $\mu = 2n+2$, from uniform initialization, satisfies $\mathrm{P}[T > \exp(\frac{c_L}{64} n p^2)] \ge \frac{1}{2}$.*

Theorem 8 is Theorem 3 with $\kappa = c_L/64$. The rest of this appendix proves it. The proof constants are collected in Table 2; the asymptotic cutoff $n_0 = n_0(p_c)$ remains implicit.

Write $L_n := np^2$ and $H := \lfloor \exp(\frac{c_L}{64} L_n) \rfloor$. Since T is an integer, Theorem 8 asserts that $\mathrm{P}[T \leq H] \leq \frac{1}{2}$. The proof follows the deficit Z_t, the number of missing Pareto slots, and shows that it does not reach 0 within H generations.

The main step is Lemma 25. If the deficit and the number of non-Pareto individuals are both at most $n/100$, then $\Theta(n)$ slots have at most two copies and receive no new copy in the variation step, and each of them is lost with probability at least $p^2/625$, because noise can reduce the fitness values of both copies in the same survival step. Hence the number Λ_t of lost slots satisfies $\mathbb{E}[e^{-2\Lambda_t} \mid \mathcal{F}_t] \leq e^{-c_L L_n}$. Gains are bounded trivially, by Z_t. So, while $Z_t \leq W := \lfloor \frac{c_L}{4} L_n \rfloor$, the process e^{-Z_t} is a supermartingale, and by optional stopping a passage from $b := \lfloor W/2 \rfloor$ down to 0 has probability at most e^{-b}. The negative-drift theorem is not used.

Two conditions are needed. First, the number J_t of non-Pareto individuals must satisfy $J_t \leq n/100$. Pointwise $J_t \leq 2Z_t + s_t$, where s_t counts the singleton slots, and a typical singleton is cloned in the variation step with probability at least $2\delta = (1 - p_c)/32$; this is the only place where p_c enters. Hence, while $Z_t \leq \varepsilon n$, s_t contracts by the factor $1 - \delta$ per generation, falls below $n/400$ within $d_0 = \lceil 9/\delta \rceil$ generations, and stays there (Lemma 32). Second, a linear deficit shrinks by at most a constant factor per generation, except on an exponentially unlikely occupancy failure (Corollary 2). On these regularity events, d_0 generations after the deficit falls below εn it is still above $c_{d_0} n > W$. The band $[0, W]$ is therefore reached only after the bound on J_t holds; another consequence of Lemma 25 makes entry below b exponentially unlikely. Each generation fails with probability $e^{-\Omega(n)}$ or $e^{-\Omega(L_n)}$, and H is small enough for a union bound. The singleton lemmas assume at most $r_n = 8Np + 8 \ln n$ corrupted evaluations per generation, which holds for all H generations with probability $1 - O(n^{-3})$ (Lemma 33). No case analysis of crossover is needed: it enters only through the facts that a round has two offspring and that it is skipped with probability $1 - p_c$.

E.1 Notation, Constants and Two Tail Bounds

Setting. We use Algorithm 1 with the reevaluation strategy, uniform initialization and $\mu = 2n + 2$. Let $m_k(t)$ count copies of π_k in P_t; slot k is present if $m_k(t) \geq 1$ and missing otherwise. The deficit is $Z_t := n + 1 - \#\{k : m_k(t) \geq 1\}$, the number of non-Pareto individuals is $J_t := \#\{x \in P_t : x \notin F\}$ (counting multiplicities), irrespective of their observed fitness values, and $s_t := \#\{k : m_k(t) = 1\}$ counts singleton slots. Let $\mathcal{F}_t$ be the history up to the start of generation t; the variation and survival steps are as in Appendix D. In addition:

- $N := n + 1$ is the number of slots. So $\mu = 2N$, and the variation step consists of $M := \mu/2 = N$ reproduction rounds. Given $\mathcal{F}_t$, the rounds are independent and identically distributed. Each round produces two offspring.
- $L_n := np^2$. The symbol $\ln$ is the natural logarithm.
- C_t is the number of corrupted evaluations in the survival step of generation t, which produces P_{t+1}; thus $C_t \sim \mathrm{Bin}(4N, p)$. Further, $C_{-1} \sim \mathrm{Bin}(2N, p)$ is the number of corrupted evaluations in the initial sort of P_0. The evaluation noise of this survival step is independent of $\mathcal{F}_t$ and of the variation step. Note that C_{t-1} is $\mathcal{F}_t$-measurable and C_t is not.
- *The sort that produced* P_t is the survival step of generation $t - 1$ if $t \geq 1$, and the initial sort if $t = 0$. The stored ranks and stored crowding distances used by the tournaments of generation t come from this sort.
- $\Lambda_t := \#\{k : \pi_k \in P_t,\ \pi_k \notin P_{t+1}\}$ is the number of *lost* slots, and $\Gamma_t := \#\{k : \pi_k \notin P_t,\ \pi_k \in P_{t+1}\}$ is the number of *gained* slots. Hence

$$Z_{t+1} - Z_t = \Lambda_t - \Gamma_t, \qquad 0 \leq \Gamma_t \leq Z_t. \tag{2}$$

- The *block* of a fitness value v in a sort is the set of sorted individuals whose fitness value equals v. A block of a layer-1 fitness value is contiguous in both objective orders, and its first and last element (its *endpoints*) are the same individuals in both orders (Algorithm 1, line 15, and Lemma 7).

Constants. Table 2 lists the working constants used below. We record three consequences of the definitions. Since $0 < \delta \le \frac{1}{64}$,

$$\varepsilon < \frac{\delta\sigma}{8} < \frac{1}{400}, \qquad 2\varepsilon + \sigma < \frac{1}{100}, \qquad p_F < \frac{\delta\sigma}{384} < 10^{-6}. \tag{3}$$

Since $\theta(c) < \frac{1}{2}$ for every $c > 0$,

$$c_{d_0} < c_{d_0-1} < \cdots < c_0 < \varepsilon/2. \tag{4}$$

Table 2: Constants used in the proof of Theorem 8. Here $N = n+1$ and $L_n = np^2$.

Symbol	Value	Role
Absolute constants		
$\theta(c)$	$\frac{1}{2}e^{-8/c}$	occupancy: surviving fraction (Lemma 24)
$\gamma(c)$	$c^2 e^{-16/c}/64$	occupancy: failure exponent (Lemma 24)
β	$\theta(\frac{1}{5})/5 = e^{-40}/10$	βn slots get no new copy (Lemma 25)
c_F	$\gamma(\frac{1}{5})/2$	failure exponent in Lemma 25
c_L	$\beta(1-e^{-2})/1250$	loss constant (Lemma 25)
p_1	$\min\{\frac{1}{200}, \sqrt{c_F/(2c_L)}\}$	largest p in Lemma 25
σ	$1/400$	upper bound σn on s_t
C_{typ}	20	Lemma 29
Constants that depend on p_c only		
δ	$(1-p_c)/64$	contraction rate of s_t
d_0	$\lceil 9/\delta \rceil = \lceil 576/(1-p_c) \rceil$	generations to reach $s_t \le \sigma n$
ε	$\delta\sigma/\big(8(20\delta+1)\big)$	near-coverage periods: $0 < Z_t \le \varepsilon n$
p_F	$\delta\sigma/\big(128(20\delta+3)\big)$	largest p in Lemma 32
c_s	$\delta^2\sigma^2/128$	failure exponent of cloning (Lemma 30)
$c_0, \ldots, c_{d_0}$	$c_0 := \varepsilon\theta(\varepsilon),\ c_{j+1} := c_j\theta(c_j)$	$Z \ge c_j n$ after j generations of a period
γ_*	$\min\{\gamma(\varepsilon), \gamma(c_0), \ldots, \gamma(c_{d_0})\}$	common occupancy exponent
Quantities that depend on n and p		
r_n	$8Np + 8\ln n$	bound on corrupted evaluations
a_n	$(20\delta+1)\varepsilon n + (20\delta+3)r_n + 20\delta$	additive term in (16)
W	$\lfloor \frac{c_L}{4} L_n \rfloor$	upper edge of the band
b	$\lfloor W/2 \rfloor$	lower bound on Z at entry into $[1, W]$
κ	$c_L/64$	exponent of the horizon
H	$\lfloor e^{\kappa L_n} \rfloor$	horizon
Constants of Theorem 8		
$p_{\max}$	$\min\big\{p_1,\ p_F,\ 1/\kappa,\ \sqrt{2c_{d_0}/c_L},\ \sqrt{\min\{\gamma_*, c_s\}/(2\kappa)}\big\}$	
C_2	$\max\big\{1/c_L,\ 16/c_L,\ \frac{64}{7c_L}\ln(64e)\big\}$	

These sufficient constants are deliberately conservative: $c_L \approx 2.94 \times 10^{-22}$, $\kappa \approx 4.59 \times 10^{-24}$ and $C_2 \approx 1.61 \times 10^{23}$. The reciprocal sequence satisfies $c_{j+1}^{-1} = 2c_j^{-1}\exp(8/c_j)$; its iteration makes $1/c_{d_0}$ and the sufficient size threshold enormous, while $p_{\max}$ becomes extremely small. This is an asymptotic existence proof, not a calibrated finite-size prediction. At $np^2 = C_2$ the horizon is only $H = 2$; the horizon grows when $np^2 \to \infty$ and becomes superpolynomial when $np^2 = \omega(\log n)$.

Two inequalities. Besides Theorem 5 we use two standard tail bounds.

Lemma 23 (bounded differences; Bernstein)**.** *Let $\xi_1, \ldots, \xi_m$ be independent random variables.*

(a) (McDiarmid, 1989). *Let f be a real function such that $|f(x) - f(x')| \le c_i$ whenever x and x' differ only in coordinate i. Then, for every $\lambda > 0$,*

$$\mathrm{P}\big[f(\xi_1, \ldots, \xi_m) \le \mathbb{E}f(\xi_1, \ldots, \xi_m) - \lambda\big] \ \le\ \exp\Big(-\frac{2\lambda^2}{\sum_{i=1}^m c_i^2}\Big).$$

(b) (Doerr, 2020). *Suppose* $\xi_i \in [0,1]$*, and let* $\xi := \sum_i \xi_i$ *and* $v \geq \mathrm{Var}[\xi]$*. Then, for every* $\lambda > 0$*,*

$$\mathrm{P}\big[\xi \geq \mathbb{E}[\xi] + \lambda\big] \;\leq\; \exp\Big(-\frac{\lambda^2}{2(v + \lambda/3)}\Big).$$

E.2 Occupancy Bounds

The first lemma says: of z given Pareto strings, a constant fraction is *not* produced by the variation step, simply because it creates only μ offspring. It is used twice. With the missing slots as targets, it bounds how fast the deficit can fall (Corollary 2). With low-multiplicity present slots as targets, it shows that many of them receive no fresh copy (Lemma 25).

Lemma 24. *Fix a generation* t*. Let* $\mathcal{S}$ *be an* $\mathcal{F}_t$*-measurable set of* $z \geq 2$ *slots. Let* K *be the number of slots* $k \in \mathcal{S}$ *such that no offspring in* Q_t *equals* π_k*. Then, for every* $p_c \in [0,1]$*, every* $p \in [0,1]$ *and every population* P_t*:*

(a) $\mathbb{E}[K \mid \mathcal{F}_t] \;\geq\; z\Big(1 - \dfrac{2}{z}\Big)^M$*;*

(b) $\mathrm{P}\Big[K < \frac{1}{2}\mathbb{E}[K \mid \mathcal{F}_t] \;\Big|\; \mathcal{F}_t\Big] \;\leq\; \exp\Big(-\dfrac{\mathbb{E}[K \mid \mathcal{F}_t]^2}{32M}\Big)$*;*

(c) if $z \geq cn$ *for some* $c > 0$*, and* $n \geq 4/c$*, then*

$$\mathrm{P}\big[K < \theta(c)\, z \;\big|\; \mathcal{F}_t\big] \;\leq\; e^{-\gamma(c)n}, \qquad \theta(c) = \tfrac{1}{2}e^{-8/c}, \quad \gamma(c) = \frac{c^2 e^{-16/c}}{64}.$$

Proof. All probabilities are conditional on $\mathcal{F}_t$. For a round $r \in \{1, \ldots, M\}$ let $A_r \subseteq \mathcal{S}$ be the set of slots $k \in \mathcal{S}$ such that one of the two offspring of round r (after mutation) equals π_k. Then $|A_r| \leq 2$, and K is the number of $k \in \mathcal{S}$ outside $A_1 \cup \cdots \cup A_M$.

Step 1: the expectation. Put $q_k := \mathrm{P}[k \in A_r]$; it does not depend on r because the rounds are identically distributed. Then

$$\sum_{k \in \mathcal{S}} q_k = \mathbb{E}|A_r| \leq 2.$$

The rounds are independent, so $\mathbb{E}[K] = \sum_{k \in \mathcal{S}} (1 - q_k)^M$. The map $q \mapsto (1-q)^M$ is convex and decreasing on $[0,1]$. Hence

$$\begin{aligned} \frac{1}{z}\sum_{k \in \mathcal{S}} (1 - q_k)^M &\geq \Big(1 - \frac{1}{z}\sum_{k \in \mathcal{S}} q_k\Big)^M && \text{(Jensen's inequality)} \\ &\geq \Big(1 - \frac{2}{z}\Big)^M && \text{(monotonicity, } \textstyle\sum_k q_k \leq 2 \text{, and } z \geq 2\text{)}. \end{aligned}$$

This is (a).

Step 2: concentration. K is a function of the M independent rounds. Replacing the outcome of one round changes A_r only, and $|A_r| \leq 2$ before and after. Hence K changes by at most 4. Lemma 23(a) with $c_i = 4$ and $\lambda = \frac{1}{2}\mathbb{E}[K]$ gives

$$\mathrm{P}\big[K < \tfrac{1}{2}\mathbb{E}[K]\big] \;\leq\; \exp\Big(-\frac{2\,(\mathbb{E}[K]/2)^2}{16M}\Big) = \exp\Big(-\frac{\mathbb{E}[K]^2}{32M}\Big).$$

This is (b).

Step 3: linearly many targets. Let $z \geq cn$ and $n \geq 4/c$. Then $2/z \leq \frac{1}{2}$, and $1 - x \geq e^{-2x}$ for $x \in [0, \frac{1}{2}]$. With $M = n + 1 \leq 2n$ we get

$$\Big(1 - \frac{2}{z}\Big)^M \;\geq\; e^{-4M/z} \;\geq\; e^{-8n/(cn)} = e^{-8/c}, \qquad \text{so} \qquad \mathbb{E}[K] \;\geq\; z\, e^{-8/c} = 2\theta(c) z.$$

Hence $\{K < \theta(c)z\} \subseteq \{K < \frac{1}{2}\mathbb{E}[K]\}$. By (b), and by $z \geq cn$ and $32M \leq 64n$,

$$\mathrm{P}[K < \theta(c)z] \;\leq\; \exp\Big(-\frac{z^2 e^{-16/c}}{32M}\Big) \;\leq\; \exp\Big(-\frac{c^2 n^2 e^{-16/c}}{64n}\Big) = e^{-\gamma(c)n}.$$

This is (c). □

Corollary 2. *Let $c > 0$ and $n \geq 4/c$. On the event $\{Z_t \geq cn\}$,*

$$\mathrm{P}\big[Z_{t+1} < \theta(c)\, Z_t \;\big|\; \mathcal{F}_t\big] \;\leq\; e^{-\gamma(c)n}.$$

This holds for every $p_c \in [0,1]$ and every $p \in [0,1]$.

Proof. Apply Lemma 24(c) to the set $\mathcal{S}$ of slots that are missing in P_t; then $z = Z_t \geq cn \geq 4$. A slot counted by K is absent from P_t and from Q_t. Since $P_{t+1} \subseteq P_t \cup Q_t$, it is absent from P_{t+1} as well. Hence $Z_{t+1} \geq K$, and $\{Z_{t+1} < \theta(c)Z_t\} \subseteq \{K < \theta(c)z\}$. □

E.3 Loss of Pareto-Optimal Solutions

Recall that Λ_t is the number of slots that are present in P_t and missing in P_{t+1}, that $\Gamma_t \leq Z_t$ is the number of gained slots, and that $L_n = np^2$.

Lemma 25. *Let $c_L = \beta(1-e^{-2})/1250$ and p_1 be the absolute constants of Table 2. There is an absolute n_0 such that the following holds for every $n \geq n_0$, every $p_c \in [0,1]$, and every $p \leq p_1$ with $c_L L_n \geq 1$. On the $\mathcal{F}_t$-measurable event*

$$1 \leq Z_t \leq n/100 \qquad \textit{and} \qquad J_t \leq n/100, \tag{5}$$

we have

$$\mathbb{E}\big[e^{-2\Lambda_t} \;\big|\; \mathcal{F}_t\big] \leq \exp(-c_L L_n), \tag{6}$$

$$\mathbb{E}\big[e^{\Gamma_t - \Lambda_t} \;\big|\; \mathcal{F}_t\big] \leq \exp\big(Z_t - \tfrac{c_L}{2} L_n\big), \tag{7}$$

$$\mathrm{P}\big[Z_{t+1} < x \;\big|\; \mathcal{F}_t\big] \leq \exp\big(x - \tfrac{c_L}{2} L_n\big) \qquad \textit{for every real } x. \tag{8}$$

Proof. Fix a population P_t that satisfies (5). All probabilities are conditional on $\mathcal{F}_t$. Steps 1–3 concern the variation step, Step 4 concerns the noise of the survival step, Step 5 combines them, and Steps 6–8 are calculations.

Step 1: at least $n/5$ central slots have one or two copies. There are $N - Z_t$ present slots. Let h of them have at least three copies. The population has $2N$ positions, so

$$(N - Z_t - h) + 3h \leq 2N, \qquad \text{that is,} \qquad h \leq \frac{N + Z_t}{2}.$$

Hence the number of slots with $m_k(t) \in \{1,2\}$ is

$$N - Z_t - h \;\geq\; \frac{N - 3Z_t}{2}.$$

At most $n/4 + 2$ slots lie outside the interval $[n/8, 7n/8]$. Let

$$\mathcal{S} := \big\{k \in [n/8,\, 7n/8] :\; m_k(t) \in \{1,2\}\big\}.$$

By $Z_t \leq n/100$ and $N = n+1$,

$$|\mathcal{S}| \;\geq\; \frac{n + 1 - 3n/100}{2} - \frac{n}{4} - 2 \;=\; 0.235\, n - 1.5 \;\geq\; \frac{n}{5} \qquad \text{for } n \geq 43.$$

Step 2: at least βn slots of $\mathcal{S}$ receive no fresh copy. Let

$$\mathcal{B} := \{k \in \mathcal{S} :\; \text{no offspring in } Q_t \text{ equals } \pi_k\}, \qquad \mathcal{A}_1 := \{|\mathcal{B}| \geq \beta n\}.$$

Apply Lemma 24(c) to $\mathcal{S}$ with $c = \frac{1}{5}$; this needs $n \geq 20$. Then $K = |\mathcal{B}|$ and $z = |\mathcal{S}| \geq n/5$, so $\theta(\frac{1}{5})z \geq \theta(\frac{1}{5})n/5 = \beta n$ and $\mathrm{P}[\mathcal{A}_1^c] \;\leq\; e^{-\gamma(1/5)\, n}$. This bound does not depend on the selection rule or on p_c.

Step 3: at least $n/9$ Pareto offspring. Let $\mathcal{A}_2$ be the event that Q_t contains at least $n/9$ Pareto strings. We show $\mathrm{P}[\mathcal{A}_2^c] \leq e^{-0.03n}$. Consider one round.

- A tournament returns a non-Pareto individual only if one of its two samples is non-Pareto. This has probability at most $2J_t/\mu \leq \frac{2n/100}{2n+2} < 0.01$. The two tournaments are independent. So both parents are Pareto strings with probability at least 0.99^2.

- If both parents are Pareto strings, at least one child before mutation is a Pareto string. If crossover is skipped, the children are the parents. Otherwise, the child that takes its prefix from the parent with the larger index is a Pareto string, for every cut (Lemma 10).
- Mutation leaves a given child unchanged with probability $(1-1/n)^n \geq \frac{1}{4}$.

Hence a round produces at least one Pareto offspring with probability at least $0.99^2/4 \geq 0.24$. The N rounds are independent. So the number F of rounds with a Pareto offspring has $\mathbb{E}[F] \geq 0.24N$. By Theorem 5 (lower tail, $\delta = \frac{1}{2}$),

$$\mathrm{P}[F < n/9] \;\leq\; \mathrm{P}\big[F \leq \tfrac{1}{2}\mathbb{E}F\big] \;\leq\; e^{-\mathbb{E}[F]/8} \;\leq\; e^{-0.03N},$$

where the first inequality uses $n/9 \leq 0.12N \leq \frac{1}{2}\mathbb{E}F$.

Step 4: few corrupted evaluations. Let $\mathcal{D} := \{C_t < n/20\}$. We have $C_t \sim \mathrm{Bin}(4N, p)$ and $p \leq p_1 \leq \frac{1}{200}$, so C_t is stochastically dominated by $\xi \sim \mathrm{Bin}(4N, \frac{1}{200})$ with $\mathbb{E}\xi = N/50$. For $n \geq 4$ we have $1.9N/50 \leq n/20$. By Theorem 5 (upper tail, $\delta = \frac{9}{10}$),

$$\mathrm{P}[\mathcal{D}^c] \;\leq\; \mathrm{P}\big[\xi \geq 1.9\,\mathbb{E}\xi\big] \;\leq\; e^{-0.81N/150} \;\leq\; e^{-n/200}.$$

On $\mathcal{A}_2 \cap \mathcal{D}$ we have two facts.

(i) *Layer 1 fills the population.* The combined population $P_t \cup Q_t$ contains at least $(\mu - J_t) + n/9$ Pareto strings. Fewer than $n/20$ of them have a corrupted evaluation. So the number of Pareto strings evaluated without corruption is larger than

$$\mu - \frac{n}{100} + \frac{n}{9} - \frac{n}{20} \;>\; \mu.$$

Such a string π_j has its true fitness value $(j, n-j)$ and lies in layer 1 (Lemma 3). Hence $|L_1| \geq \mu$, so truncation uses $i = 1$ in Algorithm 1, and every survivor belongs to L_1.

(ii) *Uncorrupted evaluations.* Let $\mathcal{K}$ be a set of at least $n/6$ present slots. Choose one copy in P_t for each $k \in \mathcal{K}$. These are at least $n/6 > n/20$ distinct individuals, so one of them is evaluated without corruption. Its fitness value is $(k', n-k')$ with $k' \in \mathcal{K}$.

Step 5: noise that removes all copies of a slot. For $k \in \mathcal{S}$ define the set of bit positions

$$\rho_k := \begin{cases} \big[k + \lceil n/5 \rceil,\; k + \lfloor n/4 \rfloor\big] & \text{if } k \leq n/2, \\ \big[k - \lfloor n/4 \rfloor,\; k - \lceil n/5 \rceil\big] & \text{if } k > n/2. \end{cases}$$

In the first case $k + \lfloor n/4 \rfloor \leq 3n/4$, and in the second case $k - \lfloor n/4 \rfloor > n/4$; so $\rho_k \subseteq \{1, \ldots, n\}$. Moreover $|\rho_k| \geq n/20 - 1 \geq n/25$ for $n \geq 100$, and $0 < k < n$ because $n/8 \leq k \leq 7n/8$. Let I_k be the indicator of the event that, in the survival step of generation t, every copy of π_k in P_t is corrupted by a flip of a bit in ρ_k. *Claim.* On $\mathcal{A}_2 \cap \mathcal{D}$, if $k \in \mathcal{B}$ and $I_k = 1$, then slot k is lost. Consequently

$$\Lambda_t \;\geq\; \sum_{k \in \mathcal{B}} I_k \qquad \text{on } \mathcal{A}_2 \cap \mathcal{D}. \tag{9}$$

Proof of the claim. Let $k \leq n/2$ and let $i \in \rho_k$ be the flipped bit of some copy of π_k. Then $i \geq k+2$, and by Lemma 4 the copy has the off-front value $(k, n-i)$. This value is strictly dominated by the front value of every slot in $[k, i]$ (Lemma 4(c)). The slots $k+1, \ldots, i$ number $i - k \geq n/5$. At most $Z_t \leq n/100$ of them are missing, so at least $n/6$ are present. By fact (ii), one of them has a copy whose observed fitness value equals its true front value. So the copy of π_k is not in layer 1, and by fact (i) it does not survive. This holds for every copy of π_k in P_t. As $k \in \mathcal{B}$, no offspring equals π_k. Hence $\pi_k \notin P_{t+1}$. For $k > n/2$ the argument is symmetric: a flip of bit $i \in \rho_k$ gives the fitness value $(i-1, n-k)$, which is strictly dominated by the front value of every slot in $[i-1, k]$, and the slots $i, \ldots, k-1$ number $k - i \geq n/5$.

Step 6: the moment generating function on the good event. Condition on $\mathcal{F}_t$ and on Q_t. This fixes $\mathcal{B}$, $\mathcal{A}_1$ and $\mathcal{A}_2$. On $\mathcal{A}_1$ choose, by a fixed rule, a set $\mathcal{B}_0 \subseteq \mathcal{B}$ with $|\mathcal{B}_0| = \lceil \beta n \rceil$. The indicators I_k, $k \in \mathcal{B}_0$, depend on the evaluation noise of pairwise disjoint sets of individuals, and the evaluation noise of the survival step is independent of $\mathcal{F}_t$ and Q_t. Hence they are independent.

A slot $k \in \mathcal{B}_0$ has $m_k(t) \in \{1, 2\}$ copies, and each evaluation is corrupted by a bit in ρ_k with probability $p|\rho_k|/n \geq p/25$. Therefore

$$\mathrm{P}[I_k = 1 \mid \mathcal{F}_t, Q_t] \;\geq\; (p/25)^{m_k(t)} \;\geq\; \frac{p^2}{625} =: q_D.$$

Using $1 - x \leq e^{-x}$,

$$\begin{aligned}
\mathbb{E}\Big[e^{-2\sum_{k\in\mathcal{B}_0} I_k} \;\Big|\; \mathcal{F}_t, Q_t\Big] &= \prod_{k\in\mathcal{B}_0} \Big(1 - (1-e^{-2})\mathrm{P}[I_k = 1 \mid \mathcal{F}_t, Q_t]\Big) && \text{(independence)} \\
&\leq \exp\big(-(1-e^{-2})\, q_D\, \lceil \beta n \rceil\big) \\
&\leq \exp\Big(-\frac{\beta(1-e^{-2})}{625}\, np^2\Big) = e^{-2c_L L_n} && \text{(definition of } c_L\text{)}.
\end{aligned}$$

Step 7: proof of (6). Put $\mathcal{A} := \mathcal{A}_1 \cap \mathcal{A}_2$. Split according to $\mathcal{A} \cap \mathcal{D}$:

$$\begin{aligned}
\mathbb{E}\big[e^{-2\Lambda_t}\big] &\leq \mathrm{P}[\mathcal{A}^c \cup \mathcal{D}^c] + \mathbb{E}\big[e^{-2\Lambda_t}\mathbf{1}_{\mathcal{A}\cap\mathcal{D}}\big] && (e^{-2\Lambda_t} \leq 1) \\
&\leq \mathrm{P}[\mathcal{A}^c \cup \mathcal{D}^c] + \mathbb{E}\Big[\mathbf{1}_{\mathcal{A}}\, e^{-2\sum_{k\in\mathcal{B}_0} I_k}\Big] && \text{(by (9), } \mathcal{B}_0 \subseteq \mathcal{B}\text{; drop } \mathbf{1}_{\mathcal{D}}\text{)} \\
&\leq \big(e^{-\gamma(1/5)n} + e^{-0.03N} + e^{-n/200}\big) + e^{-2c_L L_n} && \text{(Steps 2, 3, 4 and 6)} \\
&\leq e^{-c_F n} + e^{-2c_L L_n} && (c_F = \gamma(\tfrac{1}{5})/2,\ n \text{ large}).
\end{aligned}$$

In the last line we used that $\gamma(\frac{1}{5})$ is smaller than 0.03 and $\frac{1}{200}$, so the bracket is at most $3e^{-\gamma(1/5)n} \leq e^{-c_F n}$ for large n. Now use $p \leq p_1$. By the definition of p_1 we have $2c_L p_1^2 \leq c_F$, hence $2c_L L_n = 2c_L np^2 \leq c_F n$ and $e^{-c_F n} \leq e^{-2c_L L_n}$. With $c_L L_n \geq 1$,

$$\mathbb{E}\big[e^{-2\Lambda_t}\big] \;\leq\; 2e^{-2c_L L_n} \;=\; 2e^{-c_L L_n} \cdot e^{-c_L L_n} \;\leq\; \frac{2}{e}\, e^{-c_L L_n} \;\leq\; e^{-c_L L_n}.$$

Step 8: proof of (7) *and* (8). By the Cauchy–Schwarz inequality, by $\Gamma_t \leq Z_t$ and by (6),

$$\mathbb{E}\big[e^{\Gamma_t - \Lambda_t}\big] \;\leq\; \sqrt{\mathbb{E}\big[e^{2\Gamma_t}\big]}\,\sqrt{\mathbb{E}\big[e^{-2\Lambda_t}\big]} \;\leq\; e^{Z_t}\, e^{-c_L L_n/2}.$$

This is (7). For (8), recall $Z_t - Z_{t+1} = \Gamma_t - \Lambda_t$ from (2). By Markov's inequality and (7),

$$\mathrm{P}[Z_{t+1} < x] \;=\; \mathrm{P}\big[e^{\Gamma_t - \Lambda_t} > e^{Z_t - x}\big] \;\leq\; e^{-(Z_t - x)}\, \mathbb{E}\big[e^{\Gamma_t - \Lambda_t}\big] \;\leq\; e^{x - c_L L_n/2}.$$

□

E.4 Bounding the Number of Non-Pareto Individuals

Lemma 25 needs $J_t \leq n/100$. This subsection provides it. The plan is as follows. The number of non-Pareto individuals is controlled by the number s_t of singleton slots (Lemma 26). A singleton slot stops being a singleton when its string is cloned in the variation step and no corrupted evaluation interferes in the survival step (Lemma 31). A singleton is cloned with constant probability if it wins tournaments often enough. This holds for *good* singletons (Definition 1, Lemma 30), and all but few singletons are good (Lemma 29). Lemma 32 combines these facts.

Lemma 26. *For every population P_t,*

$$J_t \;\leq\; 2Z_t + s_t. \tag{10}$$

Proof. The population has $\mu - J_t$ Pareto strings. There are $N - Z_t$ present slots. A singleton slot contributes one copy, and each of the other $N - Z_t - s_t$ present slots contributes at least two. Hence $\mu - J_t \;\geq\; s_t + 2(N - Z_t - s_t) = 2N - 2Z_t - s_t$. With $\mu = 2N$ this is (10). □

Unobserved front values. Fix $t \geq 0$ and consider the sort that produced P_t. Let

$$\mathcal{U}_t := \{k \in \{0, \ldots, n\} : \text{ no individual in that sort has fitness } (k, n-k)\}, \qquad D_t := |\mathcal{U}_t|.$$

Both are $\mathcal{F}_t$-measurable.

Lemma 27. *For every $t \geq 0$:*

(a) $D_t \leq Z_t + C_{t-1}$;

(b) in the sort that produced P_t, the first coordinate of every off-front layer-1 fitness value lies in $\mathcal{U}_t$. Hence there are at most D_t distinct off-front layer-1 fitness values.

Proof. (a) At most Z_t slots of $\mathcal{U}_t$ are missing in P_t. Let $k \in \mathcal{U}_t$ be present, and choose a copy of π_k in P_t. This copy took part in the sort. Its evaluation in that sort was corrupted; otherwise its fitness value would have been $g(\pi_k) = (k, n-k)$. Copies chosen for distinct slots are distinct individuals. So at most C_{t-1} slots of $\mathcal{U}_t$ are present.

(b) Let (a, b') be an off-front layer-1 fitness value; then $a + b' < n$ (Lemma 2). If some individual had fitness value $(a, n-a)$, that value would strictly dominate (a, b'), and (a, b') would not be in layer 1. Hence $a \in \mathcal{U}_t$. Distinct layer-1 fitness values have distinct first coordinates (Lemma 6). So $(a, b') \mapsto a$ is an injection into $\mathcal{U}_t$. □

The normalization ranges and comparison threshold. For a sorting call let L_1 be its complete first layer, and let $R_j := \max_{x \in L_1} \tilde{v}_j(x) - \min_{x \in L_1} \tilde{v}_j(x)$ for $j = 1, 2$. Distinct layer-1 values differ in both coordinates (Lemma 6), so either both ranges are positive or the layer has one distinct value and both are zero. Put

$$c^*(L_1) := \begin{cases} 2/R_1 + 2/R_2, & R_1, R_2 > 0, \\ +\infty, & R_1 = R_2 = 0. \end{cases} \tag{11}$$

For the sort that produced P_t, write $c_t^* := c^*(L_1)$ and denote its ranges by $R_{1,t}, R_{2,t}$. These quantities are $\mathcal{F}_t$-measurable and refer to the whole layer before truncation, as do the stored crowding distances. The ranges need not be equal or equal to n.

Lemma 28. *Consider any sort by Algorithm 1. Let v be a layer-1 fitness value whose block has $m \geq 1$ members, and use the ranges and threshold of* (11) *for this sort.*

(a) The block has $\min\{m, 2\}$ endpoints. Each endpoint has positive normalized crowding distance. Every other member of the block has crowding distance 0.

(b) If the sort is a survival step, both endpoints survive.

(c) Let $v = (k, n-k)$ with $0 < k < n$, and suppose that both $(k-1, n-k+1)$ and $(k+1, n-k-1)$ are fitness values in the same sort. Then $R_1, R_2 > 0$. If $m = 1$, the only member of the block has crowding distance exactly $c^*(L_1)$; if $m \geq 2$, each endpoint has distance $c^*(L_1)/2$. In particular, every endpoint has distance at most $c^*(L_1)$.

Proof. (a) Blocks are contiguous in both objective orders and have the same endpoint individuals. With a single distinct value, these endpoints are whole-layer boundaries and receive $+\infty$; interior members receive zero by the zero-range convention. Otherwise both ranges are positive. Each endpoint is a whole-layer boundary or has a neighbour outside its block, giving an infinite or strictly positive normalized contribution. Interior members have zero gaps in both orders.

(b) Apply Lemma 7(b).

(c) The three specified front values lie in layer 1 by Lemma 3, so both ranges are positive. The immediate neighbouring blocks of $(k, n-k)$ have values $(k-1, n-k+1)$ and $(k+1, n-k-1)$: coordinates are integral and layer-1 values have distinct first coordinates (Lemma 6). The block contains no whole-layer boundary. For $m = 1$, each objective contributes a gap of 2 divided by its own range, giving $c^*(L_1)$. For $m \geq 2$, each endpoint has one neighbour in its block and one in an adjacent block in each objective sort. Both numerators are 1, giving $1/R_1 + 1/R_2 = c^*(L_1)/2$. No equality of the ranges is used. □

Definition 1. Fix $t \geq 0$ and consider the sort that produced P_t. A singleton slot k of P_t is *good* if

(G1) $0 < k < n$;

(G2) the evaluation of the only copy of π_k in P_t was not corrupted in that sort;

(G3) no individual other than a copy of π_k had the fitness value $(k, n-k)$ in that sort;

(G4) $k-1 \notin \mathcal{U}_t$ and $k+1 \notin \mathcal{U}_t$.

Otherwise it is *bad*. B_t is the number of bad singleton slots. V_t is the number of positions of P_t whose stored rank is 1 and whose stored crowding distance is larger than c_t^*. If both ranges are zero then $c_t^* = +\infty$ and $V_t = 0$. Both B_t and V_t are $\mathcal{F}_t$-measurable.

Lemma 29. *For every $t \geq 0$:*

(a) if a good singleton slot exists, both normalization ranges are positive. Its copy has stored rank 1 and stored normalized crowding distance c_t^*. Hence the positions that strictly beat it in a tournament are exactly the V_t positions of Definition 1;

(b) $B_t \leq 2C_{t-1} + 4D_t + 2$ and $V_t \leq 10D_t + 4$;

(c) consequently, with $C_{\mathrm{typ}} = 20$,

$$B_t \;\leq\; C_{\mathrm{typ}}\,(Z_t + C_{t-1} + 1), \qquad V_t \;\leq\; C_{\mathrm{typ}}\,(Z_t + C_{t-1} + 1). \tag{12}$$

Proof. (a) Let k be a good singleton slot and y the copy of π_k in P_t. By (G2), y had fitness value $(k, n-k)$, so y is in layer 1 (Lemma 3) and has stored rank 1. By (G3), every member of the block of $(k, n-k)$ is a copy of π_k. Suppose the block had two or more members. Its two endpoints are copies of π_k and belong to P_t: by Lemma 28(b) if $t \geq 1$, and because the initial sort discards nobody if $t = 0$. This contradicts $m_k(t) = 1$. So the block is $\{y\}$. By (G1), (G4) and Lemma 28(c), both ranges are positive and the crowding distance of y is c_t^*. A position strictly beats y in a tournament if its stored rank is smaller than 1, which is impossible, or if its stored rank is 1 and its stored crowding distance is larger than c_t^*.

(b) *Bound on B_t.* A bad singleton slot violates one of (G1)–(G4).

- (G1) fails for at most 2 slots.
- If (G2) fails for k, the copy of π_k is a corrupted individual. Distinct slots give distinct individuals. So (G2) fails for at most C_{t-1} slots.
- If (G3) fails for k, some $y' \neq \pi_k$ had fitness value $(k, n-k)$. By Lemma 5 the evaluation of y' was corrupted. An individual has one fitness value. So (G3) fails for at most C_{t-1} slots.
- (G4) fails for at most $2D_t$ slots.

Hence $B_t \leq 2C_{t-1} + 2D_t + 2 \leq 2C_{t-1} + 4D_t + 2$.

Bound on V_t. If both ranges are zero, $V_t = 0$ and the bound is immediate. Otherwise $c_t^* > 0$. Let x be a position counted by V_t, and let v be its fitness value. Since x has positive crowding distance, it is an endpoint of the block of v (Lemma 28(a)), and a block has at most two endpoints.

- v is off-front. There are at most D_t such blocks (Lemma 27(b)), hence at most $2D_t$ positions.
- $v = (k, n-k)$ with $0 < k < n$ and $k-1, k+1 \notin \mathcal{U}_t$. Then the crowding distance is at most c_t^* by Lemma 28(c), using the same two ranges as for the good singleton. So x is not counted by V_t.
- $v = (k, n-k)$ with $0 < k < n$, and $k-1 \in \mathcal{U}_t$ or $k+1 \in \mathcal{U}_t$. There are at most $2D_t$ such slots, hence at most $4D_t$ positions.
- $v = (k, n-k)$ with $k \in \{0, n\}$: at most 4 positions.

Hence $V_t \le 6D_t + 4 \le 10D_t + 4$.

(c) Insert $D_t \le Z_t + C_{t-1}$ (Lemma 27(a)):

$$B_t \le 2C_{t-1} + 4(Z_t + C_{t-1}) + 2 = 4Z_t + 6C_{t-1} + 2 \le 20(Z_t + C_{t-1} + 1),$$
$$V_t \le 10(Z_t + C_{t-1}) + 4 \le 20(Z_t + C_{t-1} + 1).$$

□

Cloning. Let X_t be the number of good singleton slots k of P_t such that some offspring in Q_t equals π_k. The next lemma is the only place where p_c enters the proof of Theorem 8.

Lemma 30. *Let $p_c \in [0, 1)$ be a constant, $p \le p_F$, and $n \ge n_0(p_c)$. On the $\mathcal{F}_t$-measurable event $\{Z_t \le \varepsilon n,\ C_{t-1} \le r_n\}$:*

(a) $B_t \le \sigma n/4$ and $V_t \le \mu/4$;

(b) for every good singleton slot k,

$$\mathrm{P}\big[\textit{some offspring in } Q_t \textit{ equals } \pi_k \bigm| \mathcal{F}_t\big] \ \ge\ 2\delta = \frac{1 - p_c}{32};$$

(c) if $s_t - B_t \ge \sigma n/4$, then

$$\mathrm{P}\big[X_t < \delta\,(s_t - B_t) \bigm| \mathcal{F}_t\big] \ \le\ e^{-c_s n}, \qquad c_s = \frac{\delta^2\sigma^2}{128}. \tag{13}$$

Proof. (a) By (12), both B_t and V_t are at most $20(\varepsilon n + r_n + 1)$. We bound the two main terms. By (3) and $\delta \le \frac{1}{64}$,

$$20\varepsilon n \ <\ \tfrac{5}{2}\,\delta\sigma n \ \le\ \tfrac{5}{128}\,\sigma n \ <\ \frac{\sigma n}{8}.$$

By $p \le p_F < \delta\sigma/384$,

$$20 r_n \ \le\ 160 N p_F + 160 \ln n \ \le\ \tfrac{5}{12}\,\delta\sigma N + 160 \ln n \ \le\ \frac{\sigma n}{16} \qquad \text{for large } n.$$

Hence $20(\varepsilon n + r_n + 1) \le \sigma n/4$ for large n, and $\sigma n/4 < \mu/4$.

(b) *One tournament.* Let g be the position of the copy of π_k. The tournament samples an ordered pair of positions (i, j) uniformly from μ^2 pairs. It returns g if $i = j = g$. If exactly one of i, j equals g and the other position does not strictly beat g, it returns g with probability at least $\frac{1}{2}$ (fair coin on a full tie). By Lemma 29(a), exactly V_t positions strictly beat g. Hence

$$\mathrm{P}[\text{the tournament returns } g] \ \ge\ \frac{1}{\mu^2} + \frac{2(\mu - 1 - V_t)}{\mu^2} \cdot \frac{1}{2} = \frac{\mu - V_t}{\mu^2} \ \ge\ \frac{3}{4\mu},$$

where the last step uses $V_t \le \mu/4$ from (a).

One round. Consider the event that the first tournament returns g, crossover is skipped, and mutation does not change the first child. Then the first offspring equals π_k. The three parts are independent. Crossover is skipped with probability $1 - p_c$, and $(1 - 1/n)^n \ge \frac{1}{4}$. So

$$\mathrm{P}[\text{a given round produces a copy of } \pi_k] \ \ge\ \frac{3}{4\mu} \cdot (1 - p_c) \cdot \frac{1}{4} \ \ge\ \frac{1 - p_c}{8\mu}.$$

All rounds. The $\mu/2$ rounds are independent given $\mathcal{F}_t$. With $1 - x \le e^{-x}$, and $1 - e^{-x} \ge x/2$ for $x \in [0, 1]$,

$$\mathrm{P}[\text{some round produces a copy of } \pi_k] \ \ge\ 1 - \Big(1 - \frac{1 - p_c}{8\mu}\Big)^{\mu/2} \ \ge\ 1 - e^{-(1-p_c)/16} \ \ge\ \frac{1 - p_c}{32} = 2\delta.$$

(c) There are $s_t - B_t$ good singleton slots. By (b), $\mathbb{E}[X_t \mid \mathcal{F}_t] \ge 2\delta(s_t - B_t)$. Given $\mathcal{F}_t$, the variable X_t is a function of the M independent rounds. A round has two offspring, so replacing one round changes X_t by at most 2. Lemma 23(a) with $c_i = 2$ and $\lambda = \delta(s_t - B_t) \ge \delta\sigma n/4$ gives

$$\mathrm{P}\big[X_t < \delta(s_t - B_t) \bigm| \mathcal{F}_t\big] \ \le\ \exp\Big(-\frac{2\lambda^2}{4M}\Big) \ \le\ \exp\Big(-\frac{\delta^2\sigma^2 n^2}{32(n+1)}\Big) \ \le\ \exp\Big(-\frac{\delta^2\sigma^2 n}{128}\Big).$$

□

How the number of singletons changes. The next lemma is deterministic. It says that the number of singletons can grow only through gained slots and corrupted evaluations, and that it shrinks by the number of cloned good singletons.

Lemma 31. *For every $t \geq 0$, pointwise,*

$$s_{t+1} \;\leq\; s_t - X_t + Z_t + 3C_t. \tag{14}$$

Proof. Let k be a singleton slot of P_{t+1}. At least one of the following three cases holds.

Case 1: slot k is missing in P_t. There are at most Z_t such slots.

Case 2: slot k is a singleton in P_t and no offspring in Q_t equals π_k. The X_t cloned good singleton slots are not of this kind. So there are at most $s_t - X_t$ such slots.

Case 3: $P_t \cup Q_t$ contains at least two copies of π_k. We charge k to a corrupted evaluation in the survival step of generation t.

- *Case 3a: at most one copy of π_k is evaluated without corruption.* Then some copy of π_k has a corrupted evaluation. Charge k to this copy.
- *Case 3b: at least two copies of π_k are evaluated without corruption.* Their fitness value is $(k, n-k)$. So the block of $(k, n-k)$ has at least two members, and its two endpoints survive (Lemma 28(a),(b)). If both endpoints were copies of π_k, we would have $m_k(t+1) \geq 2$. So one endpoint is an individual $y' \neq \pi_k$ with fitness value $(k, n-k)$. By Lemma 5 the evaluation of y' was corrupted. Charge k to y'.

A corrupted individual is charged at most twice: at most once in Case 3a (by its own slot, if it is a Pareto string) and at most once in Case 3b (by the slot corresponding to its observed fitness value). So Case 3 occurs for at most $2C_t \leq 3C_t$ slots. Adding the three cases gives (14). □

The contraction lemma. We say that *cloning fails* in generation t if the event

$$\mathcal{Y}_t := \big\{ Z_t \leq \varepsilon n,\;\; C_{t-1} \leq r_n,\;\; s_t \geq \sigma n/2,\;\; X_t < \delta(s_t - B_t) \big\}$$

occurs. We call the transition from t to $t+1$ *regular* if

$$Z_t \leq \varepsilon n, \qquad C_{t-1} \leq r_n, \qquad C_t \leq r_n, \qquad \text{and } \mathcal{Y}_t \text{ does not occur.} \tag{15}$$

Lemma 32. *Let $p_c \in [0,1)$ be a constant, $p \leq p_F$, and $n \geq n_0(p_c)$. Recall $a_n = (20\delta+1)\varepsilon n + (20\delta+3)r_n + 20\delta$.*

(a) Cloning rarely fails. $\mathrm{P}[\mathcal{Y}_t \mid \mathcal{F}_t] \leq e^{-c_s n}$ *for every $t \geq 0$.*

(b) One regular transition. *Let the transition from t to $t+1$ be regular. If $s_t \geq \sigma n/2$, then*

$$s_{t+1} \;\leq\; (1-\delta)s_t + a_n, \qquad \textit{where} \quad a_n \leq \tfrac{3}{8}\delta\sigma n. \tag{16}$$

If $s_t < \sigma n/2$, then

$$s_{t+1} \;\leq\; s_t + Z_t + 3C_t \;<\; \tfrac{3}{4}\sigma n. \tag{17}$$

(c) Retention. *If the transition from t to $t+1$ is regular and $s_t \leq \sigma n$, then $s_{t+1} \leq \sigma n$.*

(d) Reaching the bound. *If the d_0 transitions from t to $t+1, \ldots,$ from $t+d_0-1$ to $t+d_0$ are regular, then, for every value of s_t,*

$$s_{t+d_0} \;\leq\; \sigma n.$$

(e) Non-Pareto individuals. *If $Z_t \leq \varepsilon n$ and $s_t \leq \sigma n$, then $J_t \leq (2\varepsilon+\sigma)n < n/100$.*

Proof. (a). On the event $\{Z_t \leq \varepsilon n,\; C_{t-1} \leq r_n,\; s_t \geq \sigma n/2\}$, Lemma 30(a) gives $B_t \leq \sigma n/4$, hence $s_t - B_t \geq \sigma n/4$. Now apply (13).

(b), first case. Let $s_t \geq \sigma n/2$. The transition is regular, so $X_t \geq \delta(s_t - B_t)$. We obtain

$$\begin{aligned} s_{t+1} &\leq s_t - \delta(s_t - B_t) + Z_t + 3C_t && \text{(by (14))} \\ &= (1-\delta)s_t + \delta B_t + Z_t + 3C_t \\ &\leq (1-\delta)s_t + 20\delta(\varepsilon n + r_n + 1) + \varepsilon n + 3r_n && \text{(by (12) and (15))} \\ &= (1-\delta)s_t + a_n. \end{aligned}$$

It remains to bound a_n. By the definition of ε,

$$(20\delta + 1)\,\varepsilon n = \frac{\delta\sigma n}{8}.$$

By $p \le p_F$ and the definition of p_F,

$$(20\delta + 3)\,r_n \;\le\; (20\delta + 3)\big(8Np_F + 8\ln n\big) = \frac{\delta\sigma N}{16} + 8(20\delta + 3)\ln n.$$

The right-hand side plus the last term 20δ of a_n is $\frac{1}{16}\delta\sigma(n+1) + O(\ln n)$, which is at most $\frac{1}{4}\delta\sigma n$ for large n. Hence $a_n \;\le\; \frac{1}{8}\delta\sigma n + \frac{1}{4}\delta\sigma n \;=\; \frac{3}{8}\delta\sigma n$.

(b), second case. Let $s_t < \sigma n/2$. Drop $X_t \ge 0$ in (14) and use (15):

$$s_{t+1} \;\le\; s_t + Z_t + 3C_t \;<\; \frac{\sigma n}{2} + \varepsilon n + 3r_n.$$

Here $\varepsilon n < \delta\sigma n/8 \le \sigma n/512$ by (3) and $\delta \le \frac{1}{64}$. Further $p_F < \delta\sigma/384 \le \sigma/24576$, so

$$3r_n \;\le\; 24Np_F + 24\ln n \;\le\; \frac{\sigma N}{1024} + 24\ln n \;\le\; \frac{\sigma n}{512} \qquad \text{for large } n.$$

Hence $\varepsilon n + 3r_n < \sigma n/4$, which gives (17).

(c). If $s_t < \sigma n/2$, use (17). If $\sigma n/2 \le s_t \le \sigma n$, use (16): $s_{t+1} \;\le\; (1-\delta)\sigma n + \frac{3}{8}\delta\sigma n \;<\; \sigma n$.

(d). Suppose first that $s_{t+i} < \sigma n/2$ for some $i \in \{0, \dots, d_0\}$. If $i = d_0$ we are done. Otherwise $s_{t+i} \le \sigma n$, and (c) keeps $s \le \sigma n$ along the remaining regular transitions. Suppose now that $s_{t+i} \ge \sigma n/2$ for all $i \in \{0, \dots, d_0\}$. Then (16) applies to all d_0 transitions. Unrolling the recursion,

$$\begin{aligned}
s_{t+d_0} &\le (1-\delta)^{d_0} s_t + a_n \sum_{i=0}^{d_0-1} (1-\delta)^i \\
&\le e^{-\delta d_0}\, N + \frac{a_n}{\delta} && (s_t \le N, \text{ geometric sum}) \\
&\le e^{-9} \cdot 2n + \tfrac{3}{8}\sigma n && (d_0 \ge 9/\delta,\ N \le 2n,\ (16)) \\
&\le \tfrac{1}{8}\sigma n + \tfrac{3}{8}\sigma n \;<\; \sigma n && (2e^{-9} < \sigma/8 = 1/3200).
\end{aligned}$$

(e). By Lemma 26, $J_t \le 2Z_t + s_t \le (2\varepsilon + \sigma)n$, and $2\varepsilon + \sigma < \frac{1}{100}$ by (3). □

E.5 Bounding the Number of Corrupted Evaluations

Lemmas 30 and 32 assume $C_{t-1}, C_t \le r_n$. The next lemma shows that this holds for all generations up to an exponential horizon.

Lemma 33. *Let $K > 0$, $p \le 1/K$, and let $H' \ge 1$ be an integer with $H' \le e^{Knp^2}$. Then, with $r_n = 8Np + 8\ln n$,*

$$\mathrm{P}\Big[\max_{-1 \le t < H'} C_t > r_n\Big] \;\le\; (H'+1)\exp\big(-\tfrac{3}{2}Np - 3\ln n\big) \;\le\; 2n^{-3}.$$

Proof. Step 1: one generation. Each C_t is stochastically dominated by $\xi \sim \mathrm{Bin}(4N, p)$; for $t = -1$ this holds because $\mathrm{Bin}(2N, p) \preceq \mathrm{Bin}(4N, p)$. Put $m := \mathbb{E}\xi = 4Np$ and $\lambda := m + 8\ln n$. Then $m + \lambda = r_n$, $\mathrm{Var}[\xi] \le m$ and $m \le \lambda$. By Lemma 23(b) with $v = m$,

$$\mathrm{P}[\xi \ge r_n] \;\le\; \exp\Big(-\frac{\lambda^2}{2(m + \lambda/3)}\Big) \;\le\; \exp\Big(-\frac{3\lambda}{8}\Big) = \exp\big(-\tfrac{3}{2}Np - 3\ln n\big),$$

where the second inequality uses $2(m + \lambda/3) \le \frac{8}{3}\lambda$.

Step 2: union bound. There are $H' + 1$ indices t. This gives the first inequality. For the second, $Kp \le 1$ gives $Knp^2 \le np \le \frac{3}{2}Np$, hence $H' e^{-\frac{3}{2}Np} \le 1$. Therefore $(H'+1)e^{-\frac{3}{2}Np} \le 2$, and $e^{-3\ln n} = n^{-3}$. □

E.6 Proof of Theorem 8

Fix a constant $p_c \in [0,1)$ and the constants of Table 2. Let $n \geq n_0(p_c)$ and let p satisfy $C_2/n \leq p^2 \leq p_{\max}^2$, that is, $L_n \geq C_2$ and $p \leq p_{\max}$. We show $\mathrm{P}[T \leq H] \leq \frac{1}{2}$.

Proof of Theorem 8. Step 0: the parameters fit together. We use the definitions of $p_{\max}$ and C_2.

(P1) $p \leq p_1$ and $c_L L_n \geq 1$, so Lemma 25 applies. $p \leq p_F$, so Lemmas 30 and 32 apply. $p \leq 1/\kappa$ and $H \leq e^{\kappa n p^2}$, so Lemma 33 applies with $K = \kappa$ and $H' = H$.

(P2) $p^2 \leq 2c_{d_0}/c_L$ gives $W \leq \frac{c_L}{4} n p^2 \leq \frac{1}{2} c_{d_0} n$. With (4) and (3),

$$W \;<\; c_{d_0} n \;<\; c_0 n \;<\; \varepsilon n \;<\; n/400. \tag{18}$$

(P3) $L_n \geq 16/c_L$ gives $W \geq 4$ and $b \geq 2$. Moreover,

$$\tfrac{c_L}{8} L_n - 1 \;\leq\; b \;\leq\; \tfrac{c_L}{8} L_n. \tag{19}$$

The lower bound holds because $W \geq \frac{c_L}{4} L_n - 1$ and $b \geq \frac{W}{2} - \frac{1}{2}$.

(P4) $p^2 \leq \min\{\gamma_*, c_s\}/(2\kappa)$ gives

$$\ln H \;\leq\; \kappa L_n \;\leq\; \tfrac{1}{2} \min\{\gamma_*, c_s\}\, n. \tag{20}$$

Step 1: the regularity event $\mathcal{G}$. Let $\mathcal{G} := \mathcal{G}_{\mathrm{init}} \cap \mathcal{G}_{\mathrm{occ}} \cap \mathcal{G}_{\mathrm{cap}} \cap \mathcal{G}_{\mathrm{clone}}$, where

- $\mathcal{G}_{\mathrm{init}}$: P_0 contains no Pareto string. Then $Z_0 = N$. Each of the μ initial strings is a Pareto string with probability $N 2^{-n}$. So $\mathrm{P}[\mathcal{G}_{\mathrm{init}}^c] \leq \mu N 2^{-n}$.
- $\mathcal{G}_{\mathrm{occ}}$: for every $0 \leq t < H$ and every $c \in \{\varepsilon, c_0, c_1, \ldots, c_{d_0-1}\}$,

$$Z_t \geq cn \quad \Longrightarrow \quad Z_{t+1} \geq \theta(c)\, Z_t. \tag{21}$$

These are $H(d_0+1)$ implications. Corollary 2 applies to each of them, because $n \geq n_0(p_c) \geq 4/c_{d_0-1}$. So each fails with probability at most $e^{-\gamma(c)n} \leq e^{-\gamma_* n}$, and $\mathrm{P}[\mathcal{G}_{\mathrm{occ}}^c] \leq H(d_0+1)e^{-\gamma_* n}$.

- $\mathcal{G}_{\mathrm{cap}}$: $C_t \leq r_n$ for all $-1 \leq t < H$. By Lemma 33, $\mathrm{P}[\mathcal{G}_{\mathrm{cap}}^c] \leq 2n^{-3}$.
- $\mathcal{G}_{\mathrm{clone}}$: cloning does not fail, that is, no $\mathcal{Y}_t$ occurs for $0 \leq t < H$. By Lemma 32(a), $\mathrm{P}[\mathcal{G}_{\mathrm{clone}}^c] \leq H e^{-c_s n}$.

Adding up, and using (20) in the second line,

$$\begin{aligned} \mathrm{P}[\mathcal{G}^c] \;\leq\; \varepsilon_{\mathrm{reg}} &:= \mu N 2^{-n} + H(d_0+1)\big(e^{-\gamma_* n} + e^{-c_s n}\big) + 2n^{-3} \\ &\leq \mu N 2^{-n} + 2(d_0+1)\, e^{-\min\{\gamma_*, c_s\} n/2} + 2n^{-3}. \end{aligned} \tag{22}$$

The right-hand side does not depend on p and tends to 0 as $n \to \infty$.

Step 2: periods close to coverage on $\mathcal{G}$. A *near-coverage period* is a maximal interval $[\tau, \tau']$ of generations with $0 < Z_t \leq \varepsilon n$. We only consider such periods with $\tau < T$. On $\mathcal{G}_{\mathrm{init}}$ we have $Z_0 = N > \varepsilon n$, so $\tau \geq 1$. Since $\tau - 1 < T$, we have $Z_{\tau-1} \neq 0$. By maximality,

$$Z_{\tau-1} > \varepsilon n. \tag{23}$$

Claim 1 (the first d_0 generations of a period stay above W). On $\mathcal{G}$, for every $0 \leq j \leq d_0$ with $\tau + j \leq H$,

$$Z_{\tau+j} \;\geq\; c_j n \;>\; W.$$

Proof. The second inequality is (18). We prove the first by induction on j. For $j = 0$, use (21) with $c = \varepsilon$ at time $\tau - 1$, and (23):

$$Z_\tau \;\geq\; \theta(\varepsilon)\, Z_{\tau-1} \;>\; \theta(\varepsilon)\, \varepsilon n \;=\; c_0 n.$$

Let $j < d_0$, $\tau + j + 1 \leq H$ and $Z_{\tau+j} \geq c_j n$. Use (21) with $c = c_j$ at time $\tau + j$: $Z_{\tau+j+1} \geq \theta(c_j)\, Z_{\tau+j} \geq \theta(c_j)\, c_j n = c_{j+1} n$.

Claim 2 (later generations have few non-Pareto individuals). On $\mathcal{G}$, let $t \le H$ belong to the period and satisfy $t \ge \tau + d_0$. Then

$$s_t \le \sigma n \qquad \text{and} \qquad J_t < n/100.$$

Proof. Let $\tau \le u < t$. Then $Z_u \le \varepsilon n$ because u belongs to the period, $C_{u-1}, C_u \le r_n$ by $\mathcal{G}_{\mathrm{cap}}$, and $\mathcal{Y}_u$ does not occur by $\mathcal{G}_{\mathrm{clone}}$. So the transition from u to $u+1$ is regular in the sense of (15). Lemma 32(d), applied at τ, gives $s_{\tau+d_0} \le \sigma n$. Lemma 32(c), applied to the transitions from $\tau + d_0$ to t, gives $s_t \le \sigma n$. Lemma 32(e) gives $J_t < n/100$.

Step 3: no jump below b. For $0 \le t < H$ define the $\mathcal{F}_t$-measurable event

$$\mathcal{M}_t := \{W < Z_t \le n/100,\ \ J_t \le n/100\}, \qquad \mathcal{E}_{\mathrm{jump}} := \bigcup_{t<H} \big(\mathcal{M}_t \cap \{Z_{t+1} < b\}\big).$$

On $\mathcal{M}_t$ the hypothesis (5) of Lemma 25 holds. By (8) with $x = b$, by (19), and by $H \le e^{(c_L/64)L_n}$,

$$\mathrm{P}[\mathcal{E}_{\mathrm{jump}}] \ \le\ H\, e^{\,b - \frac{c_L}{2}L_n} \ \le\ H\, e^{-\frac{3}{8}c_L L_n} \ \le\ e^{-\frac{23}{64}c_L L_n}. \tag{24}$$

Step 4: one entry into the band rarely ends in coverage. For $1 \le u \le H$ define the $\mathcal{F}_u$-measurable *entry event* $\mathcal{E}_{\mathrm{entry}}(u) := \{Z_{u-1} > W,\ \ b \le Z_u \le W,\ \ J_u \le n/100\}$, the stopping time $\eta_u := \min\{t \ge u :\ Z_t = 0 \ \text{ or } \ Z_t > W \ \text{ or } \ J_t > n/100\} \wedge H$, and $\mathcal{E}_{\mathrm{cover}}(u) := \mathcal{E}_{\mathrm{entry}}(u) \cap \{Z_{\eta_u} = 0\}$.

Claim 3. $\mathrm{P}[\mathcal{E}_{\mathrm{cover}}(u)] \le e^{-b}$.
Proof. Let $u \le t < \eta_u$. Then $1 \le Z_t \le W < n/100$ by (18), and $J_t \le n/100$. So (5) holds. By (2) and (7),

$$\begin{aligned}
\mathbb{E}\big[e^{-(Z_{t+1}-Z_t)} \,\big|\, \mathcal{F}_t\big] &= \mathbb{E}\big[e^{\Gamma_t - \Lambda_t} \,\big|\, \mathcal{F}_t\big] \\
&\le \exp\big(Z_t - \tfrac{c_L}{2}L_n\big) && \text{(by (7))} \\
&\le \exp\big(W - \tfrac{c_L}{2}L_n\big) \ \le\ e^{-\frac{c_L}{4}L_n} \ <\ 1 && (Z_t \le W \le \tfrac{c_L}{4}L_n).
\end{aligned}$$

Hence $S_t := \exp(-Z_{t\wedge\eta_u})$, $t \ge u$, is a supermartingale with values in $[0,1]$. The stopping time η_u is bounded by H. By optional stopping, on $\mathcal{E}_{\mathrm{entry}}(u)$,

$$\mathrm{P}\big[Z_{\eta_u} = 0 \,\big|\, \mathcal{F}_u\big] \ \le\ \mathbb{E}\big[e^{-Z_{\eta_u}} \,\big|\, \mathcal{F}_u\big] \ \le\ e^{-Z_u} \ \le\ e^{-b}.$$

Step 5: coverage before H forces one of the bad events. We claim

$$\{T \le H\} \ \subseteq\ \mathcal{G}^c \ \cup\ \mathcal{E}_{\mathrm{jump}} \ \cup\ \bigcup_{u=1}^{H} \mathcal{E}_{\mathrm{cover}}(u). \tag{25}$$

Assume that $\mathcal{G}$ occurs, that $\mathcal{E}_{\mathrm{jump}}$ does not occur, and that $T \le H$. We have $Z_0 = N > W$. Let $t^* := \max\{t < T :\ Z_t > W\}$, $\quad u := t^* + 1 \le T \le H$. Then $Z_{t^*} > W$, and $Z_t \le W$ for $u \le t \le T$.

(i) $Z_{t^*} \le \varepsilon n$. Otherwise (21) with $c = \varepsilon$ gives $Z_u \ge \theta(\varepsilon) Z_{t^*} > c_0 n > W$, a contradiction.

(ii) $J_{t^*} < n/100$ *and* $Z_u \ge b$. By (i), t^* belongs to a near-coverage period; let $\tau \le t^*$ be its start. If $u \le \tau + d_0$, Claim 1 gives $Z_u > W$, a contradiction. So $t^* \ge \tau + d_0$, and Claim 2 gives $J_{t^*} < n/100$. Together with $W < Z_{t^*} \le \varepsilon n < n/100$ this shows that $\mathcal{M}_{t^*}$ occurs. Since $\mathcal{E}_{\mathrm{jump}}$ does not occur, $Z_u \ge b$.

(iii) *The bound on J_t holds until coverage.* Since $b \ge 2$, we have $u < T$. For $u \le t < T$ we have $0 < Z_t \le W < \varepsilon n$. So t belongs to the same period and $t \ge \tau + d_0$. Claim 2 gives $J_t < n/100$.

(iv) By (ii) and (iii), $\mathcal{E}_{\mathrm{entry}}(u)$ occurs. By (iii), none of the three stopping conditions of η_u holds for $u \le t < T$. At time $T \le H$ we have $Z = 0$. Hence $\eta_u = T$ and $Z_{\eta_u} = 0$. So $\mathcal{E}_{\mathrm{cover}}(u)$ occurs.

This proves (25).

Step 6: the numbers. By (25), (22), (24) and Claim 3,

$$\mathrm{P}[T \le H] \ \le\ \varepsilon_{\mathrm{reg}} + e^{-\frac{23}{64}c_L L_n} + (H+1)\, e^{-b}.$$

By (19) and $H+1 \le 2H \le 2e^{(c_L/64)L_n}$,

$$(H+1)\,e^{-b} \;\le\; 2e^{\frac{c_L}{64}L_n}\cdot e^{\,1-\frac{c_L}{8}L_n} \;=\; 2e\cdot e^{-\frac{7}{64}c_L L_n}.$$

Now use $L_n \ge C_2 \ge \frac{64}{7c_L}\ln(64e)$. Then $e^{\frac{7}{64}c_L L_n} \ge 64e$, so

$$2e\cdot e^{-\frac{7}{64}c_L L_n} \le \frac{1}{32}, \qquad e^{-\frac{23}{64}c_L L_n} \le (64e)^{-23/7} \le \frac{1}{32}.$$

By (22), $\varepsilon_{\mathrm{reg}} \le \frac{1}{4}$ for all $n \ge n_0(p_c)$. Hence

$$\mathrm{P}[T \le H] \;\le\; \frac{1}{4}+\frac{1}{32}+\frac{1}{32} \;<\; \frac{1}{2}.$$

Since T is an integer and $H = \lfloor e^{(c_L/64)L_n}\rfloor$, the event $\{T > H\}$ equals $\{T > \exp(\frac{c_L}{64}np^2)\}$. Moreover,

$$\mathbb{E}[T] \ge (H+1)\mathrm{P}[T>H] \ge \frac{H+1}{2} \ge \frac{1}{2}\lfloor e^{(c_L/64)np^2}\rfloor,$$

which gives the expectation claim in Theorem 3. □

F Proof of Theorem 4(a): Runtime Upper Bound for the Reevaluation Strategy on PairedCOCZ

This appendix defines PairedCOCZ formally and proves Theorem 10, the precise form of Theorem 4(a): at $p = 1$ the *reevaluating* NSGA-II covers the front in $O(n^2\mu^{4M-7})$ expected generations, from every initial population.

At $p = 1$ every evaluation misjudges the defect of a string by exactly one, and survival selection orders the combined population by the *observed* defect. So true defects that differ by at least three are never confused, which alone brings an arbitrary population to one with at most γn defects per member. From then on *favorable* evaluations, which underestimate the defect, are rare, and the defects decrease level by level until the population is *saturated*: all but $\varepsilon\mu$ members are Pareto optima. The front is then built one class at a time: a neighbouring class arises from a double bit flip inside one pair, at a cost of n^2 rather than n^{2r}.

Notation. The algorithm is the NSGA-II of Algorithm 1 with the reevaluation strategy, an even population size μ, crossover probability p_c, and one-bit prior noise with $p = 1$. Each crowding computation uses the stable sorting rule specified in Section 2.3. The survival arguments below concern the endpoints of the current blocks, not identities retained across generations. The main notation and sufficient parameter choices are collected here.

Table 3: Notation and parameter choices for the reevaluation proof on PairedCOCZ.

Symbol	Meaning
n, r, M, W	string length; number of pairs ($2r < n$); $M := r+1$ is the size of the front; $W := r+1$ is the penalty per defect
$x = (w, u)$	$w \in \{0,1\}^{2r}$ is read as r pairs, $u \in \{0,1\}^{n-2r}$
c_{11}, c_{00}, bad	numbers of pairs equal to 11, to 00, and to 01 or 10
$d(x)$, $\rho(x)$	defect $d := \mathrm{bad} + \#\mathrm{zeros}(u)$; number of repairing bits $\rho := \#\mathrm{zeros}(u) + 2\,\mathrm{bad}$; so $d \le \rho \le 2d$
optimum, class	a string with $d = 0$; its class is $j = c_{11} \in \{0,\dots,r\}$
V_j, U_i	one-defect values $V_j := (j, r-j) - (W, W)$ and $U_i := (i, r-1-i) - (W, W)$
η	$:= 2r/n$, the probability that the noise bit falls into w
P_t, R_t, $\mathcal{F}_t$	population; combined population of generation t (2μ members); history up to P_t
$\tilde{d}$, favorable	observed defect after evaluation; an evaluation is favorable if $\tilde{d} = d-1$

Symbol	Meaning
$N_a(S)$	number of members of the multiset S with true defect at most a
I_a	number of members of R_t with a favorable evaluation and true defect $a+1$ or $a+2$
J_t, NO_t	number of members of R_t with a favorable evaluation and true defect 1 or 2; number of non-optimal members of P_t
$\theta, \varepsilon, \gamma$	$\theta := (1-p_c)/4, \quad \varepsilon := \theta/16, \quad \gamma := \theta/256$
a_0, low defect	$a_0 := \lfloor \gamma n \rfloor - 2$; P_t has low defect if all its members have true defect at most a_0
saturated	P_t is saturated if $N_0(P_t) \geq (1-\varepsilon)\mu$
D_t	$:= \min\{b : N_b(R_t) \geq \mu\}$, the μ-th smallest true defect in R_t
c_1, c_*	$c_1 := 1/288; \quad c_* := (1-p_c)(\gamma/8)^7/e$
C_0	$:= \max\{((2r+2)M+4)/(c_1\theta),\ 6/c_*\}$; the hypothesis is $\mu \geq C_0 \ln n$
good, τ_{bad}, δ_{bad}	Definition 3; first generation that is not good; $\delta_{\mathrm{bad}} := n\, e^{-c_1\theta\mu}$
q_a	$:= (1-p_c)\, a/(4en)$, success rate of the descent at level a
L_0, L_1	$L_0 := \mu + (8/\theta)(2+\ln\mu) + 4$ (one epoch); $L_1 := 2\mu n + 8en\ln(n)/(1-p_c)$ (descent)
H, π	$H := n^{(2r+2)M+2}; \quad \pi := \frac{1}{4}\big[(1-p_c)/(4\mu^2 n^{2r})\big]^M$
$c_j(s)$, $\mathcal{S}_s$, A_s	number of optima of class j in P_s; $\mathcal{S}_s := \{j : c_j(s) \geq 1\}; \quad A_s := \#\{j : c_j(s) = 1\}$
loss, gain	loss at s: $\mathcal{S}_s \not\subseteq \mathcal{S}_{s+1}$; gain at s: $\mathcal{S}_s \supsetneq \mathcal{S}_{s+1}$
g, q, b	$g := (1-p_c)/(32\mu n^2); \quad q := \theta/(8\mu); \quad b := 50Mr\mu/n$
Λ, L_2	$\Lambda := 16(4Mr^2+40); \quad L_2 := \lfloor n^2/(\Lambda\mu^3) \rfloor$, the length of one stage
T	first generation t with $\mathcal{S}_t = \{0, \ldots, r\}$

"For n large" means $n \geq n_0(r, p_c)$. We use three elementary facts without further mention: $(1-1/n)^n \geq 1/4$ for $n \geq 2$; $(1-1/n)^{n-1} \geq 1/e$; and, for $x \in [0,1]$ and an integer $k \geq 1$,

$$1-(1-x)^k \;\geq\; 1-e^{-kx} \;\geq\; \min\{\tfrac{1}{2},\ \tfrac{kx}{2}\}. \tag{26}$$

We use Theorem 5 in the following two forms. Let X be a sum of independent indicator variables. If $\mathbb{E}[X] \geq m$, or if $\mathbb{E}[X] \leq m'$, then respectively

$$\mathrm{P}\big[X \leq \tfrac{1}{2}m\big] \;\leq\; e^{-m/8}, \qquad\qquad \mathrm{P}\big[X \geq \tfrac{3}{2}m'\big] \;\leq\; e^{-m'/12}. \tag{27}$$

These are the two bounds of Theorem 5 with $\delta = 1/2$. They remain valid with a bound in place of the exact mean, because X then dominates (is dominated by) a sum of independent indicators with mean exactly m (exactly m').

Histories and restarts. Here $\mathcal{F}_t$ contains the complete state at the start of generation t, including stored fitness values, ranks and crowding distances. Let $\mathcal{H}_t$ additionally reveal the offspring strings of generation t, but not its fresh evaluations or survival tie-breaking choices. Then $\mathcal{F}_t \subseteq \mathcal{H}_t \subseteq \mathcal{F}_{t+1} \subseteq \mathcal{H}_{t+1}$. In particular, D_t is $\mathcal{H}_t$-measurable, not necessarily $\mathcal{F}_t$-measurable. At a finite $\mathcal{F}$-stopping time the subsequent variation, evaluation and tie-breaking choices have their prescribed fresh laws; at an $\mathcal{H}$-stopping time this applies to the still-unrevealed choices of that generation and to later generations. To justify these claims, partition on the possible deterministic values of the stopping time and use independence of the unrevealed random choices on each event. All conditional bounds below are uniform over the corresponding revealed history, so the same bounds apply at these restarts.

The proofs also allow an arbitrary fixed initial stored state, with any initial fitness values, ranks and crowding distances for which tournament comparisons are well-defined. The first variation step may use these data. Its survival step evaluates the entire combined population afresh. In Lemma 40, every comparison of stored ranks occurs after that refresh; the subsequent progress arguments either sample both tournament participants from the desired set or sample one position twice. Thus they require no distributional assumption on the initial ranks and crowding distances.

F.1 THE FUNCTION PAIREDCOCZ

Definition 2. Let $r \geq 1$ and $n > 2r$. Write $x \in \{0,1\}^n$ as $x = (w, u)$, where $w = (x_1, \ldots, x_{2r})$ is read as the r pairs $(x_1x_2), (x_3x_4), \ldots, (x_{2r-1}x_{2r})$ and $u = (x_{2r+1}, \ldots, x_n)$. Let $c_{11}(x)$, $c_{00}(x)$ and $\mathrm{bad}(x)$ be the numbers of pairs equal to 11, equal to 00, and equal to 01 or 10. The *defect* of x is $d(x) := \mathrm{bad}(x) + \#\{\text{zeros of } u\}$. With the penalty $W := r + 1$, the function is

$$\textsc{PairedCOCZ}(x) \;=\; f(x) \;:=\; \big(c_{11}(x) - W\,d(x),\;\; c_{00}(x) - W\,d(x)\big),$$

and both objectives are maximised. A string with $d(x) = 0$ is called an *optimum*; its *class* is $j := c_{11}(x) \in \{0, \ldots, r\}$. The runtime T is the first generation t such that P_t contains an optimum of every class.

Lemma 34. *If $d(x) < d(y)$, then $f(x)$ strictly dominates $f(y)$; in fact $f_i(x) \geq f_i(y) + 1$ for $i = 1, 2$.*

Proof. Since $0 \leq c_{11} \leq r$ and $d(y) - d(x) \geq 1$,

$$f_1(x) - f_1(y) \;=\; c_{11}(x) - c_{11}(y) + W\big(d(y) - d(x)\big) \;\geq\; -r + W \;=\; 1.$$

The same computation with c_{00} gives $f_2(x) - f_2(y) \geq 1$. □

Lemma 34 identifies the Pareto set as

$$\mathcal{P} := \{x : d(x) = 0\} = \{00, 11\}^r \times \{1^{n-2r}\}, \qquad |\mathcal{P}| = 2^r.$$

Class j contains $\binom{r}{j}$ optima, all with value $(j, r - j)$. Thus the front has $M = r + 1$ points, namely $\{(j, r - j) : 0 \leq j \leq r\}$. Every front value strictly dominates every off-front value of f.

The tail u gives both objectives the same direction of improvement: changing a zero to a one reduces $d(x)$ and improves both objectives. In contrast, changing a pair from 00 to 11 increases $c_{11}(x)$ but decreases $c_{00}(x)$. The penalty $r+1$ is larger than the range of either count, so reducing $d(x)$ always improves both objectives, regardless of the pair counts.

Examples with $n = 20$. Let $r = 3$. The first six bits form three pairs, the tail has 14 bits, and the penalty is $r + 1 = 4$. Below, commas separate the pairs and a semicolon separates them from the tail; 1^{14} denotes fourteen ones. There are $2^3 = 8$ Pareto-optimal strings, but only four front values:

$$\{(0, 3), (1, 2), (2, 1), (3, 0)\}.$$

For example, the following five strings are Pareto-optimal:

x	$c_{11}(x)$	$c_{00}(x)$	$d(x)$	$f(x)$
$(00, 00, 00; 1^{14})$	0	3	0	$(0, 3)$
$(11, 00, 00; 1^{14})$	1	2	0	$(1, 2)$
$(00, 11, 00; 1^{14})$	1	2	0	$(1, 2)$
$(11, 11, 00; 1^{14})$	2	1	0	$(2, 1)$
$(11, 11, 11; 1^{14})$	3	0	0	$(3, 0)$

The second and third strings are different solutions with the same fitness $(1, 2)$. The four front values have $1, 3, 3, 1$ optimal strings, respectively. Covering the front requires one true optimum for each value, not all eight strings. Increasing n with $r = 3$ only lengthens the all-ones tail of these optima; the Pareto front still has four values.

For comparison, the following strings are nonoptimal:

x	$c_{11}(x)$	$c_{00}(x)$	$d(x)$	$f(x)$
$(01, 00, 00; 1^{14})$	0	2	1	$(-4, -2)$
$(11, 00, 00; 1^{13}0)$	1	2	1	$(-3, -2)$
$(10, 00, 11; 1^{12}00)$	1	1	3	$(-11, -11)$

The first has one mixed pair, the second has one zero in the tail, and the third has both a mixed pair and two tail zeros, giving $d(x) = 1+2 = 3$. Noise can make $(01, 00, 00; 1^{14})$ receive the front value $(1, 2)$ by flipping its first bit during evaluation. The reverse flip makes the optimum $(11, 00, 00; 1^{14})$

receive $(-4, -2)$. Neither evaluation changes the actual string. Thus a nonoptimal string can receive better noisy fitness than a true optimum, even though every true front value dominates its true fitness.

The defect also measures distance to the Pareto set:

$$d(x) = \min_{z \in \mathcal{P}} \mathrm{Ham}(x, z).$$

Indeed, reaching $\mathcal{P}$ requires changing at least one bit in each bad pair and each zero of u. These disjoint repairs use exactly $d(x)$ bit flips. If $d(x) > 0$, performing any one repair reduces the defect by one and hence strictly improves both objectives by Lemma 34. Thus no nonoptimum is locally optimal under one-bit moves, despite the noisy algorithm's absorbing states.

For two optima, Hamming distance is twice the number of pairs on which their choices of 00 and 11 differ. Optima from classes j and k therefore have distance at least $2|j - k|$, with equality when the smaller set of 11 pairs is contained in the larger. Flipping an entire pair moves to an adjacent class; under these moves the Pareto set forms an r-dimensional hypercube. This two-bit connectivity contrasts with the one-bit isolation used by the noise mechanism below.

Lemma 35. *Every Hamming neighbour of an optimum has defect exactly one.*

Proof. An optimum has $u = 1 \cdots 1$ and only pairs equal to 00 or 11. Flipping a bit of u creates one zero of u. Flipping a bit of a pair turns this pair into 01 or 10. In both cases the defect becomes one. □

Lemma 36. *The values of f on strings of defect one are exactly the $2r + 1$ values*

$$V_j := (j, r-j)-(W, W) \quad (0 \le j \le r), \qquad U_i := (i, r-1-i)-(W, W) \quad (0 \le i \le r-1).$$

The value V_j is attained when the defect is a zero of u and $c_{11} = j$; the value U_i is attained when the defect is a bad pair and $c_{11} = i$. No one-defect value strictly dominates any V_j. Each U_i is strictly dominated by V_i and by V_{i+1}.

Proof. If the single defect is a zero of u, all r pairs are consistent, so $c_{11} + c_{00} = r$ and the value is V_j with $j = c_{11}$. If the single defect is a bad pair, then $c_{11} + c_{00} = r - 1$ and the value is U_i with $i = c_{11}$. The two coordinates of V_j sum to $r - 2W$, and the coordinates of every one-defect value sum to at most $r - 2W$. A value that strictly dominates V_j has a strictly larger coordinate sum, so it is not a one-defect value. Finally $U_i = V_i - (0, 1) = V_{i+1} - (1, 0)$. □

F.2 Noisy Fitness at $p = 1$

At $p = 1$ an evaluation of x returns the *noisy fitness value* $f(x \oplus e_I)$ for a uniformly random position I. We call $\tilde{d}(x) := d(x \oplus e_I)$ the *observed defect*. The $\rho(x)$ *repairing bits* of x are the zeros of u and both bits of every bad pair. The evaluation is *favorable* if I is a repairing bit.

Lemma 37. *Consider one generation t of the reevaluating NSGA-II at $p = 1$.*

(R1) Off by one. *For every member x of R_t,*

$$\tilde{d}(x) = \begin{cases} d(x) - 1 & \text{with probability } \rho(x)/n \le 2d(x)/n \quad \text{(favorable)}, \\ d(x) + 1 & \text{otherwise.} \end{cases}$$

The 2μ noise positions of generation t are independent of each other, of $\mathcal{F}_t$, and of the variation step of generation t.

(R2) Selection by observed defect. *An individual with a smaller observed defect has a strictly smaller rank than an individual with a larger one. Hence, for every integer δ: if at least μ members of R_t have observed defect at most δ, then every member of P_{t+1} has observed defect at most δ; if at most μ members of R_t have observed defect at most δ, then all of them belong to P_{t+1}.*

(R3) Possible noisy fitness values. *An optimum of class j has noisy fitness V_j if its noise falls into u, which has probability $1 - \eta$; noise in a 11 pair gives U_{j-1} (only if $j \ge 1$), while noise in a 00 pair gives U_j (only if $j \le r - 1$). An individual has observed defect 0 only if it has true defect one and a favorable evaluation. The individuals of observed defect 1*

are exactly the optima and the individuals of true defect two with a favorable evaluation. In particular, an individual with noisy fitness V_j is an optimum of class j with noise in u, or it has true defect two and a favorable evaluation.

(R4) Crowding. *In a layer of the non-dominated sorting, call the set of members with one common noisy fitness value a* block. *Its* endpoints *are its first and last members in the input order inherited by the layer; the two objective sorts have the same endpoints. The endpoints have positive normalized crowding distance, and all other members of the block have crowding distance zero. Consequently, if a layer contains at most k distinct fitness values and at least $2k$ individuals remain to be selected when this layer is processed, then all endpoints of all its blocks belong to P_{t+1}.*

Proof. (R1). Flipping a zero of u removes this defect. Flipping either bit of a bad pair makes the pair consistent and removes this defect. Flipping a one of u creates a new zero, and flipping a bit of a consistent pair creates a new bad pair. So a repairing bit lowers the defect by one and every other bit raises it by one. The position I is uniform, and $\rho \leq 2d$. Independence holds because the evaluations use fresh independent noise (Algorithm 1).

(R2). The noisy fitness of x is the true fitness of $x \oplus e_I$, whose defect is the observed defect. By Lemma 34, a smaller observed defect yields a fitness value that strictly dominates every fitness value with a larger observed defect. In non-dominated sorting, an individual whose noisy fitness is strictly dominated lies in a strictly later layer than the dominating individual. So every layer consists of individuals of one observed defect, and the layers are ordered by observed defect. Algorithm 1 takes whole layers in this order and selects the remaining individuals from the next layer. Both claims follow.

(R3). By Lemma 35 an optimum has noisy fitness equal to a one-defect value. If its noise falls into u, the noisy string has a zero in u and the same pairs, so its fitness is V_j (Lemma 36). If the noise hits a pair 11, its fitness is U_{j-1}; if it hits a pair 00, its fitness is U_j. By (R1), a string of true defect one has observed defect 0 or 2, a string of true defect two has observed defect 1 or 3, and a string of true defect at least three has observed defect at least 2. This gives the remaining claims.

(R4). Distinct noisy fitness values in a non-dominated layer differ in both coordinates. Thus either the layer has one block and both objective ranges are zero, or both ranges are positive. The two stable objective sorts preserve the same input order within each block and hence have the same endpoint individuals. In the positive-range case, an endpoint is a whole-layer boundary or has a neighbour outside its block, giving an infinite contribution or a strictly positive gap divided by a positive range. Interior members have zero gaps in both objectives. In the single-block case, the endpoints are whole-layer boundaries and receive $+\infty$; all other contributions are zero by the zero-range convention. There are consequently at most $2k$ positive-crowding members, all selected before any zero-crowding member when at least $2k$ individuals remain to be selected. □

The next two lemmas are deterministic statements about one survival selection. They hold for every realisation of the noise.

Lemma 38. *Fix $a \geq 0$ and a generation t. Let I_a be the number of members of R_t with a favorable evaluation whose true defect is $a+1$ or $a+2$. Then*

$$N_a(P_{t+1}) \;\geq\; \min\{\, N_a(R_t),\; \mu - I_a \,\}.$$

Proof. Let G be the set of individuals with observed defect at most $a+1$. By Lemma 37(R1), a member with true defect at most a belongs to G. A member with true defect at least $a+1$ belongs to G only if its true defect is $a+1$ or $a+2$ and its evaluation is favorable. Hence $|G| \;=\; N_a(R_t) + I_a$, and at most I_a members of G have true defect above a. If $|G| \leq \mu$, then all of G survives by (R2), so $N_a(P_{t+1}) = N_a(R_t)$. If $|G| > \mu$, then by (R2) the survivors are μ members of G, and at least $\mu - I_a$ of them have true defect at most a. □

When $d = \Theta(n)$, an evaluation is favorable with constant probability, so Lemma 38 does not guarantee progress. However, the observed defect differs from the true defect by exactly one, so *defects that differ by at least three are never confused*. Recall $D_t := \min\{b : N_b(R_t) \geq \mu\}$.

Lemma 39. *For every generation t:*

(E1) *every member of P_{t+1} has true defect at most $D_t + 2$;*

(E2) *if $D_t \geq a+3$, then every member of R_t with true defect at most a belongs to P_{t+1}.*

Proof. (E1). At least μ members of R_t have true defect at most D_t, hence observed defect at most D_t+1. By (R2) all survivors have observed defect at most D_t+1, and by (R1) true defect at most D_t+2.

(E2). $D_t \geq a+3$ means that fewer than μ members of R_t have true defect at most $a+2$. By (R1) every individual of observed defect at most $a+1$ is one of them. So there are fewer than μ such individuals, and by (R2) all of them survive. Every member of true defect at most a has observed defect at most $a+1$ and therefore survives. □

F.3 Reaching a Population with Low Defect

This section uses only Lemma 39. The quantity D_t is lowered by one in each *epoch*. Within an epoch, strings that are five levels below D_t are created, they multiply because they win their tournaments, and they fill half of the population.

Lemma 40. *Let n be large. Fix a generation t, condition on the history up to and including the creation of R_t, and suppose $D_t = q \geq a_0 - 1 = \lfloor \gamma n \rfloor - 3$. Let $\sigma := \min\{s > t : D_s \leq q-1\}$. Then there is a failure event $\mathcal{F}$ with*

$$\mathrm{P}[\mathcal{F}] \leq e^{-c_*\mu/2} + \Big(\frac{8}{\theta}+2\Big)e^{-\theta\mu/64} \qquad \text{and} \qquad \mathbb{E}\big[(\sigma - t)\,\mathbf{1}_{\mathcal{F}^c}\big] \leq L_0 = \mu + \frac{8}{\theta}(2+\ln\mu) + 4.$$

The bounds are conditional on this history and remain valid when t is any finite $\mathcal{H}$-stopping time.

Proof. An active trial at generation $s > t$ is started only if the epoch has not ended before; this is the $\mathcal{F}_s$-measurable condition $\sigma \geq s$. After creating R_s, the trial either ends the epoch ($D_s \leq q-1$) or has $D_s \geq q$, in which case Lemma 39 controls selection. We do not condition the offspring law on the latter outcome. The first child of a round is, with probability $1-p_c$, a mutant of the winner of the first tournament. With probability at least $\theta = (1-p_c)/4$ it is an unchanged copy of this winner (no crossover, no bit flipped). The first children of the $\mu/2$ rounds are independent given $\mathcal{F}_s$.

Step 1: generation $t+1$ creates a string of true defect at most $q-5$. By (E1) every member of P_{t+1} has true defect $d' \leq q+2$. Let d' be the true defect of the winner of a first tournament. If $d' \leq q-5$, an unchanged copy has defect at most $q-5$; this has probability at least $\theta \geq c_*$. Otherwise put $k := d' - (q-5) \in [1,7]$. Choose k of the d' defects and one repairing bit in each. If mutation flips exactly these k bits, the child has defect $q-5$. Since $d' \geq q-4 \geq \gamma n - 8$, this has probability at least

$$\binom{d'}{k}\frac{1}{n^k}\Big(1-\frac{1}{n}\Big)^{n-1} \geq \Big(\frac{d'}{kn}\Big)^k\frac{1}{e} \geq \Big(\frac{\gamma n - 8}{7n}\Big)^7\frac{1}{e} \geq \Big(\frac{\gamma}{8}\Big)^7\frac{1}{e}$$

for $n \geq 64/\gamma$. Together with the probability $1-p_c$ of no crossover, each first child of generation $t+1$ has true defect at most $q-5$ with probability at least $c_* = (1-p_c)(\gamma/8)^7/e$, in both cases. Let $\mathcal{F}_1$ be the event that none has. Then $\mathrm{P}[\mathcal{F}_1] \leq (1-c_*)^{\mu/2} \leq e^{-c_*\mu/2}$. Outside $\mathcal{F}_1$, either the epoch has ended at generation $t+1$, or (E2) with $a = q-5$ applies there because $D_{t+1} \geq q = (q-5)+5$ and puts this child into P_{t+2}. On $\mathcal{F}_1$ we abandon the following analysis stages; their trial counts are then defined to be zero.

Step 2: Stage 1, while $N_{q-3}(P_s) \leq \mu/2$. Let $G_s := N_{q-5}(P_s)$, so $G_{t+2} \geq 1$ if this stage is entered. Let Y_s be the number of first children of generation s that are unchanged copies of a member of P_s with true defect at most $q-5$. If the epoch does not end at generation s, (E2) with $a = q-5$ gives $G_{s+1} \geq G_s + Y_s$. The stored ranks of P_s come from the survival selection of generation $s-1$. A member of true defect at most $q-5$ had observed defect at most $q-4$ there. In Stage 1 at least $\mu/2$ members have true defect at least $q-2$; they had observed defect at least $q-3$. By (R2) the former have a strictly smaller stored rank. So a tournament returns a member of true defect at most $q-5$ whenever its first sample is such a member and its second sample has true defect at least $q-2$; this has probability at least $(G_s/\mu)\cdot\frac{1}{2}$. By (26),

$$\mathrm{P}[Y_s \geq 1 \mid \mathcal{F}_s] \geq 1 - \Big(1 - \frac{\theta G_s}{2\mu}\Big)^{\mu/2} \geq \min\Big\{\frac{1}{2}, \frac{\theta G_s}{8}\Big\}.$$

Thus each active trial either ends the epoch or increases G_s, with conditional probability at least $\min\{1/2, \theta G_s/8\}$. While neither happens, G_s cannot decrease. By iterating this conditional

bound, the expected number of active trials with $G_s = g$ is at most $\max\{2, 8/(\theta g)\}$. This counts the terminal trial as well and does not condition on any future failure event. In Stage 1, $G_s \le N_{q-3}(P_s) \le \mu/2$. Summing over $1 \le g \le \mu/2$, the expected length of Stage 1 is at most

$$\sum_{g=1}^{\mu/2}\Big(2 + \frac{8}{\theta g}\Big) \;\le\; \mu + \frac{8}{\theta}\,(1 + \ln \mu).$$

Step 3: Stage 2, once $N_{q-3}(P_s) > \mu/2$. The first entry into this stage is a generation-start stopping time. In each active trial, either the epoch ends at generation s or (E2) with $a = q-3$ gives $N_{q-3}(P_{s+1}) = N_{q-3}(R_s) \ge N_{q-3}(P_s)$. Thus the stage is never left before the epoch ends. Both samples of a first tournament have true defect at most $q-3$ with probability at least $1/4$, and then so has the winner. So a first child is an unchanged copy of such a member with probability at least $\theta/4$, and the expected number of such first children is at least $\theta\mu/8$. By (27), at least $\theta\mu/16$ first children are such copies, except with probability $e^{-\theta\mu/64}$. Let $\mathcal{F}_2$ be the event that this fails in one of the first $\lceil 8/\theta \rceil$ active trials of Stage 2. Whether such a trial is entered is known at its generation start, so conditional bounds and a union bound give $\mathrm{P}[\mathcal{F}_2] \le (8/\theta + 1)\, e^{-\theta\mu/64}$. Outside $\mathcal{F}_2$, after $\lceil 8/\theta \rceil$ generations of Stage 2 the combined population satisfies

$$N_{q-3}(R_s) \;>\; \frac{\mu}{2} + \Big\lceil \frac{8}{\theta} \Big\rceil \cdot \frac{\theta\mu}{16} \;\ge\; \mu,$$

that is $D_s \le q-3$, and σ has occurred. So Stage 2 lasts at most $8/\theta + 1$ generations.

Step 4: total. Put $\mathcal{F} := \mathcal{F}_1 \cup \mathcal{F}_2$. Outside $\mathcal{F}$, $\sigma - t$ is at most 2 (Step 1) plus the length of Stage 1 plus $8/\theta + 1$. Taking expectations,

$$\mathbb{E}\big[(\sigma - t)\mathbf{1}_{\mathcal{F}^c}\big] \;\le\; 2 + \mu + \frac{8}{\theta}(1 + \ln \mu) + \frac{8}{\theta} + 1 \;\le\; L_0.$$

The positive conditional progress rates make an infinite Stage 1 have probability zero. On $\mathcal{F}^c$, Stage 2 has the displayed deterministic length bound. All estimates are uniform in the history at the start, establishing the stopping-time version as well. □

Lemma 41. *Let n be large. Start from an arbitrary population and let $\tau_0 := \min\{t : D_t \le a_0 - 2\}$. Then P_{τ_0+1} has low defect. Moreover there is a failure event $\mathcal{F}_G$ with*

$$\mathrm{P}[\mathcal{F}_G] \;\le\; n\Big(e^{-c_*\mu/2} + \Big(\frac{8}{\theta} + 2\Big)e^{-\theta\mu/64}\Big) \qquad \textit{and} \qquad \mathbb{E}\big[\tau_0\, \mathbf{1}_{\mathcal{F}_G^c}\big] \;\le\; n\, L_0.$$

If $\mu \ge C_0 \ln n$, then $\mathrm{P}[\mathcal{F}_G] \le 1/n$.

Proof. By (E1) every member of P_{τ_0+1} has true defect at most $D_{\tau_0} + 2 \le a_0$, so P_{τ_0+1} has low defect. The time τ_0 is an $\mathcal{H}$-stopping time, so $\tau_0 + 1$ is a generation-start stopping time. Define epochs by $t_0 := 0$ and, as long as $D_{t_k} \ge a_0 - 1$, let t_{k+1} be the time σ of Lemma 40 started at t_k. Then $D_{t_{k+1}} \le D_{t_k} - 1$. Defects are at most n, so $D_{t_0} \le n$ and there are at most n epochs before $D \le a_0 - 2$. Stop the epoch construction at the first revealed failure. Every entered epoch starts at an $\mathcal{H}$-stopping time, so the conditional bounds of Lemma 40 apply there. Let $\mathcal{F}_G$ be the union of their failure events. The conditional union bound gives its probability bound. On $\mathcal{F}_G^c$, τ_0 is the sum of at most n successful epoch lengths. For each entered epoch, its length times the indicator of its own success has conditional expectation at most L_0; dropping subsequent success indicators can only increase this nonnegative quantity. Summation gives the expectation bound without conditioning on $\mathcal{F}_G^c$.

For the last claim, $\mu \ge C_0 \ln n$ and $C_0 \ge 6/c_*$ give $n\, e^{-c_*\mu/2} \le n^{-2} \le 1/(2n)$. Also $C_0 \ge 288 \cdot 12/\theta$, so $\theta\mu/64 \ge 54 \ln n$ and $n(8/\theta + 2)e^{-\theta\mu/64} \le 1/(2n)$ for n large. □

F.4 Reaching a Saturated Population

From now on defects are at most $a_0 + 2 \le \gamma n$ wherever it matters, so favorable evaluations are rare and Lemma 38 becomes effective.

Definition 3. Generation t is *good* if

(G1) $I_a \le \varepsilon\mu$ for every $0 \le a \le a_0$, and

(G2) in case P_t is saturated, at least $\theta\mu/8$ offspring of generation t are optima.

τ_{bad} is the first generation that is not good.

Lemma 42. *For every generation t, every $0 \le a \le a_0$, and every state of the algorithm,*

$$\mathrm{P}\big[I_a > \varepsilon\mu \mid \mathcal{F}_t\big] \;\le\; e^{-c_1\theta\mu}, \qquad c_1 = \tfrac{1}{288},$$

and

$$\mathrm{P}\big[\textit{generation } t \textit{ is not good} \mid \mathcal{F}_t\big] \;\le\; \delta_{\mathrm{bad}} \;=\; n\,e^{-c_1\theta\mu}.$$

If $\mu \ge C_0 \ln n$, then $\delta_{\mathrm{bad}} \le n^{-(2r+2)M-3}$.

Proof. (G1). Condition on R_t. A member of true defect $a+1$ or $a+2$ with $a \le a_0$ has a favorable evaluation with probability at most

$$\frac{2(a+2)}{n} \;\le\; \frac{2\lfloor\gamma n\rfloor}{n} \;\le\; 2\gamma \;=\; \frac{\theta}{128},$$

independently of the others, by (R1). So I_a is a sum of at most 2μ independent indicators and

$$\mathbb{E}[I_a \mid R_t] \;\le\; 2\mu\cdot\frac{\theta}{128} \;=\; \frac{\varepsilon\mu}{4} \;\le\; \frac{2\varepsilon\mu}{3} =: m'.$$

By (27), $\mathrm{P}[I_a \ge \varepsilon\mu \mid R_t] = \mathrm{P}[I_a \ge \tfrac{3}{2}m' \mid R_t] \le e^{-m'/12} = e^{-\theta\mu/288}$.

(G2). Let P_t be saturated. Both samples of a first tournament are optima with probability at least $(1-\varepsilon)^2$, and then the winner is an optimum. So a first child is an unchanged copy of an optimum with probability at least $(1-\varepsilon)^2\theta \ge \theta/2$. The first children of the $\mu/2$ rounds are independent, so the expected number of such copies is at least $m := \theta\mu/4$. By (27), fewer than $m/2 = \theta\mu/8$ offspring are optima with probability at most $e^{-m/8} = e^{-\theta\mu/32}$.

Union bound. There are a_0+1 values of a, and $a_0+2 \le \gamma n \le n$. So generation t is not good with probability at most $(a_0+1)e^{-\theta\mu/288} + e^{-\theta\mu/32} \le n\,e^{-c_1\theta\mu}$. Finally $\mu \ge C_0\ln n$ and $C_0 \ge ((2r+2)M+4)/(c_1\theta)$ give $e^{-c_1\theta\mu} \le n^{-(2r+2)M-4}$. □

Lemma 43. *Let $0 \le a \le a_0$ and let generation t be good.*

(i) *If $N_a(P_t) \ge (1-\varepsilon)\mu$, then $N_a(P_{t+1}) \ge (1-\varepsilon)\mu$. In particular, a saturated population is followed by a saturated population.*

(ii) *Let $a \ge 1$, let Y_t be the number of offspring of generation t with true defect at most $a-1$, and put $X_t := \min\{N_{a-1}(P_t), (1-\varepsilon)\mu\}$. Then*

$$X_{t+1} \;\ge\; \min\big\{X_t + Y_t,\; (1-\varepsilon)\mu\big\}.$$

Moreover, for every generation t (good or not) with $N_a(P_t) \ge (1-\varepsilon)\mu$ and $1 \le a \le a_0$,

$$\mathrm{P}\big[Y_t \ge 1 \mid \mathcal{F}_t\big] \;\ge\; \min\Big\{\frac{1}{2},\; \frac{\mu q_a}{2}\Big\}, \qquad q_a = \frac{(1-p_c)\,a}{4en}. \tag{28}$$

Proof. (i). Apply Lemma 38. Here $N_a(R_t) \ge N_a(P_t) \ge (1-\varepsilon)\mu$, and $\mu - I_a \ge (1-\varepsilon)\mu$ by (G1).

(ii). Apply Lemma 38 at threshold $a-1$. Here $N_{a-1}(R_t) = N_{a-1}(P_t) + Y_t$ and $\mu - I_{a-1} \ge (1-\varepsilon)\mu$ by (G1).

The rate (28). Consider the first child of one round. Both samples of the first tournament have true defect at most a with probability at least $(1-\varepsilon)^2 \ge \tfrac{1}{2}$, and then so has the winner z. With probability $1-p_c$ there is no crossover. If $d(z) = a$, the child has defect $a-1$ when mutation flips exactly one of the $\rho(z) \ge a$ repairing bits and no other bit; this has probability at least

$$a\cdot\frac{1}{n}\Big(1-\frac{1}{n}\Big)^{n-1} \;\ge\; \frac{a}{en}.$$

If $d(z) \le a-1$, the child is an unchanged copy with probability at least $\frac{1}{4} \ge \frac{a}{en}$, because $a \le \gamma n$. So the first child has true defect at most $a-1$ with probability at least

$$\frac{1}{2}\,(1-p_c)\,\frac{a}{en} \;=\; 2q_a.$$

The $\mu/2$ rounds are independent given $\mathcal{F}_t$, and (26) gives $1-(1-2q_a)^{\mu/2} \ge \min\{\frac{1}{2}, \mu q_a/2\}$. □

Lemma 44. *Let n be large, let P_0 have low defect, and let τ be the first generation in which the population is saturated. Then*

$$\mathbb{E}\big[\tau \wedge \tau_{\mathrm{bad}}\big] \;\le\; L_1 \;=\; 2\mu n + \frac{8en\ln n}{1-p_c},$$

and P_t is saturated for every t with $\tau \le t \le \tau_{\mathrm{bad}}$.

Proof. Step 1: levels. For $t \le \tau_{\mathrm{bad}}$ let ℓ_t be the smallest $a \le a_0$ with $N_a(P_t) \ge (1-\varepsilon)\mu$. Since P_0 has low defect, $N_{a_0}(P_0) = \mu$ and $\ell_0 \le a_0$. All generations before τ_{bad} are good, so Lemma 43(i) shows that ℓ_t is non-increasing for $t \le \tau_{\mathrm{bad}}$. Moreover τ is the first time with $\ell_t = 0$. This proves the last claim.

Step 2: stopped trials at one level $a \ge 1$. A trial is active when no earlier generation was bad and the current level is a. This event is $\mathcal{F}_t$-measurable; whether generation t itself is good is not. Stop the trials when the level drops or when a bad generation is detected, including that terminal generation in the trial count. Set $K_a := \lceil (1-\varepsilon)\mu \rceil$ and use the residual potential $V_t := K_a - N_{a-1}(P_t) > 0$ during an active trial, with value zero after stopping. Initially $V_t \le \mu$. On a good generation, Lemma 43(ii) makes V_t non-increasing and decreases it by at least one if $Y_t \ge 1$, unless it reaches zero. On a bad generation the stopped potential is set to zero. Consequently, from every active generation-start history,

$$\mathbb{E}[V_t - V_{t+1} \mid \mathcal{F}_t] \;\ge\; \mathrm{P}[Y_t \ge 1 \mid \mathcal{F}_t] \;\ge\; p_a := \min\{1/2, \mu q_a/2\}.$$

This uses the unconditional offspring rate (28), not its law conditioned on the current generation being good. Additive drift bounds the expected number of active trials, including a possible terminal bad trial, by

$$\frac{\mu}{\min\{1/2,\ \mu q_a/2\}} \;=\; \max\Big\{2\mu,\ \frac{2}{q_a}\Big\} \;\le\; 2\mu + \frac{2}{q_a}.$$

Each level is entered at most once before the first bad generation. These trial counts therefore dominate the number of generations up to $\tau \wedge \tau_{\mathrm{bad}}$.

Step 3: sum over the levels. Since $a_0 \le \gamma n$ and $e\gamma \le 1$, we have $\sum_{a=1}^{a_0} 1/a \le 1 + \ln(\gamma n) \le \ln n$. Therefore

$$\mathbb{E}[\tau \wedge \tau_{\mathrm{bad}}] \;\le\; \sum_{a=1}^{a_0}\Big(2\mu + \frac{8en}{(1-p_c)\,a}\Big) \;\le\; 2\mu n + \frac{8en\ln n}{1-p_c}.$$

□

The last lemma of this section shows that, in the saturated regime, non-optimal members are replaced in every generation and only a few strings with a favorable evaluation survive.

Lemma 45. *Let generation t be good and let P_t be saturated. Then R_t contains more than μ optima, and every member of P_{t+1} is an optimum or has true defect 1 or 2 and a favorable evaluation. Hence $\mathrm{NO}_{t+1} \le J_t$. Moreover $\mathbb{E}[J_t \mid R_t] \le 8\mu/n$ for every R_t. Consequently, with $\mathcal{G}_t$ the event that generation t is good and P_t is saturated,*

$$\mathbb{E}\big[\mathrm{NO}_{t+1}\mathbf{1}_{\mathcal{G}_t}\big] \;\le\; \frac{8\mu}{n}. \tag{29}$$

Proof. By (G2), R_t contains at least $(1-\varepsilon)\mu + \theta\mu/8 = (1+\theta/16)\mu > \mu$ optima. By (R3) each of them has observed defect 1. By (R2) every member of P_{t+1} has observed defect at most 1, and by (R3) it is an optimum, or it has true defect two or one and a favorable evaluation. So the non-optimal members of P_{t+1} are counted by J_t. The combined population has 2μ members, and a member of true defect 1 or 2 has a favorable evaluation with probability $\rho/n \le 4/n$ by (R1). This gives $\mathbb{E}[J_t \mid R_t] \le 2\mu \cdot 4/n$. Finally $\mathrm{NO}_{t+1}\mathbf{1}_{\mathcal{G}_t} \le J_t$, which gives (29). □

F.5 A Simple Polynomial Bound

One optimum suffices to cover the whole front in a single generation, with a small but polynomial probability. Together with the early phase and the descent this already gives a polynomial bound. The proof also introduces the restart argument that is used again for Theorem 10.

Lemma 46. *Let $2M \le \mu \le n/32$ and $n \ge 4rM$. If P_t contains an optimum, then*

$$\mathrm{P}\big[P_{t+1} \text{ covers the front} \mid \mathcal{F}_t\big] \;\ge\; \pi \;=\; \frac{1}{4}\Big[\frac{1-p_c}{4\mu^2 n^{2r}}\Big]^M.$$

Proof. Let $x \in P_t$ be an optimum and let h be its class.

Step 1: variation creates one optimum y_j of every class j. Use round $j+1$ for class $j = 0, \dots, r$; this is possible because $M \le \mu/2$. In this round we require three things. Both samples of the first tournament are the position of x; this has probability μ^{-2}, and then x wins whatever the stored ranks and crowding distances are. There is no crossover; this has probability $1 - p_c$. The mutation of the first child flips exactly both bits of $|h-j|$ prescribed pairs of x (pairs 00 if $j > h$, pairs 11 if $j < h$, nothing if $j = h$); this has probability

$$n^{-2|h-j|}\Big(1-\frac{1}{n}\Big)^{n-2|h-j|} \;\ge\; \frac{n^{-2r}}{4}.$$

Then the first child y_j is an optimum of class j. The rounds are independent, so Step 1 has probability at least $\big[(1-p_c)/(4\mu^2 n^{2r})\big]^M$.

Step 2: the noise. Condition on R_t. Let B_1 be the event that the noise of each y_j falls into u, and B_2 the event that no member of R_t of true defect 1 or 2 has a favorable evaluation. By (R3) and a union bound, $\mathrm{P}[B_1] \ge 1 - M\eta = 1 - 2rM/n \ge \frac{1}{2}$. A member of true defect 1 or 2 has a favorable evaluation with probability at most $4/n$, so $\mathrm{P}[B_2] \ge 1 - 8\mu/n \ge \frac{3}{4}$. The two events concern different members of R_t, so they are independent and $\mathrm{P}[B_1 \cap B_2] \ge \frac{1}{4}$.

Step 3: selection on $B_1 \cap B_2$. By (R3) and B_2, no individual has observed defect 0, and the individuals of observed defect 1 are exactly the optima in R_t. By B_1, each y_j has noisy fitness V_j. Every other noisy fitness value of an optimum is some U_i, which is strictly dominated by V_i (Lemma 36). Hence the first layer consists exactly of individuals with noisy fitness in $\{V_0, \dots, V_r\}$, and the block of V_j consists of optima of class j only. The layer has M distinct fitness values and is processed first, when $\mu \ge 2M$ individuals remain to be selected. By (R4) both endpoints of every block belong to P_{t+1}. So P_{t+1} contains an optimum of every class. □

Theorem 9. *Let $r \ge 1$ and $p_c \in [0,1)$ be constants, $M = r+1$, $p = 1$, and let n be large. Let μ be even with*

$$C_0 \ln n \;\le\; \mu \;\le\; \frac{n}{32} \qquad \text{and} \qquad \mu \ge 2M,$$

where $C_0 = C_0(r, p_c)$ is the constant of the notation table. Then the reevaluating NSGA-II on PairedCOCZ(n, r) *satisfies,* from every initial population *and with $H := n^{(2r+2)M+2}$,*

$$\mathrm{P}[\,T \le H\,] \;\ge\; 1 - \frac{4}{n} \qquad \text{and} \qquad \mathbb{E}[T] \;\le\; 2H.$$

Proof. We split the budget H into $H/4$ generations for the early phase, $H/4$ for the descent, and $H/2$ for the coverage. For n large, $2n^2 L_0 + 1 \le H/4$ and $2nL_1 \le H/4$, because $L_0 = O(n)$, $L_1 = O(n^2)$ and $H \ge n^{10}$.

Step 1: early phase. By Lemma 41 and Markov's inequality,

$$\mathrm{P}\big[\tau_0 > 2n \cdot nL_0\big] \;\le\; \mathrm{P}[\mathcal{F}_G] + \frac{\mathbb{E}[\tau_0 \mathbf{1}_{\mathcal{F}_G^c}]}{2n^2 L_0} \;\le\; \frac{1}{n} + \frac{1}{2n}.$$

Otherwise the population at time $\tau_0 + 1 \le H/4$ has low defect. Restart the clock there.

Step 2: no bad generation. By Lemma 42 and a union bound,

$$\mathrm{P}[\tau_{\mathrm{bad}} < H] \;\le\; H\,\delta_{\mathrm{bad}} \;\le\; n^{(2r+2)M+2} \cdot n^{-(2r+2)M-3} \;=\; \frac{1}{n}.$$

Step 3: descent. By Lemma 44 and Markov's inequality, $\mathrm{P}[\tau \wedge \tau_{\mathrm{bad}} > 2nL_1] \le 1/(2n)$. Outside the events of Steps 2 and 3 we have $\tau \le H/4$, and every P_t with $\tau \le t \le H$ is saturated. A saturated population contains at least $(1-\varepsilon)\mu \ge 1$ optima.

Step 4: coverage. By Lemma 46 and $\mu \le n$, each generation that starts with an optimum in the population covers the front with conditional probability at least

$$\pi \;\ge\; \frac{1}{4}\Big(\frac{1-p_c}{4}\Big)^M n^{-(2r+2)M} \;=\; \frac{\theta^M}{4}\, n^{-(2r+2)M}.$$

The early phase and descent end at an integer time at most $\lfloor H/2 \rfloor$, leaving at least $\lceil H/2 \rceil$ generations within the budget. The probability that all these generations start with an optimum but none covers the front is at most

$$(1-\pi)^{\lceil H/2 \rceil} \;\le\; \exp(-\pi H/2) \;\le\; \exp\big(-\theta^M n^2/8\big) \;\le\; \frac{1}{n}.$$

Step 5: conclusion. The four failure probabilities sum to $\frac{3}{2n} + \frac{1}{n} + \frac{1}{2n} + \frac{1}{n} = \frac{4}{n}$. Outside these events, $T \le \lfloor H/2 \rfloor + \lceil H/2 \rceil = H$.

Step 6: expectation. The bound of Step 5 holds for every initial state. By the Markov property, $\mathrm{P}[T > (k+1)H \mid T > kH] \le 4/n \le \frac{1}{2}$ for every $k \ge 0$. Hence $\mathrm{P}[T > kH] \le 2^{-k}$ and $\mathbb{E}[T] \le H \sum_{k \ge 0} 2^{-k} = 2H$. □

F.6 Covering the Front Class by Class

Theorem 9 asks one generation to create all classes at once, which costs a factor n^{2r} per class. It is much cheaper to create one *neighbouring* class at a time, by a double bit flip inside one pair, which costs n^2. The price is that the classes obtained so far must be kept for $\Theta(n^2/\mu^3)$ generations. The probability of losing any present class in one generation is at most $\eta A_s + M\eta^2 + 40\mu^2/n^2$ (Lemma 47). Doubling single copies quickly controls A_s (Lemma 48); the remaining $O(\mu^2/n^2)$ loss term also enters the choice of stage length.

Throughout this section μ is even, $4M+1 \le \mu \le n/32$, and n is large. Recall $\eta = 2r/n$, the counts $c_j(s)$, the set $\mathcal{S}_s = \{j : c_j(s) \ge 1\}$ of present classes, and $A_s = \#\{j : c_j(s) = 1\}$. A *loss* at generation s means $\mathcal{S}_s \not\subseteq \mathcal{S}_{s+1}$. A *gain* at s means $\mathcal{S}_s \subsetneq \mathcal{S}_{s+1}$: a new class appears and no class is lost.

Lemma 47. *Let $\mu \ge 4M+1$. Fix a generation s, condition on the strings of R_s, and let j be a class. Suppose that the noise of generation s satisfies*

(i) *some optimum of class j in R_s has its noise in u;*

(ii) *at most one member of R_s of true defect 2 has a favorable evaluation;*

(iii) *at most one member of R_s of true defect 1 has a favorable evaluation.*

Then P_{s+1} contains an optimum of class j. Consequently,

$$\mathrm{P}\big[\textit{loss at } s \mid \mathcal{F}_s\big] \;\le\; p_s \;:=\; \eta A_s + M\eta^2 + \frac{40\mu^2}{n^2}. \tag{30}$$

Proof. Step 1: all individuals with noisy fitness in $\{V_0, \dots, V_r\}$ lie in one layer $\mathcal{L}^*$, preceded by at most one individual. By (R3) and (iii) at most one individual has observed defect 0. If it exists, its noisy fitness strictly dominates that of every other individual (Lemma 34), so it forms the first layer alone. No fitness value of observed defect 1 strictly dominates V_j (Lemma 36), and V_j strictly dominates every fitness value of observed defect at least 2. So all individuals with noisy fitness in $\{V_0, \dots, V_r\}$ lie in the first layer $\mathcal{L}^*$ after any individual of observed defect 0.

Step 2: the endpoints of all blocks of $\mathcal{L}^*$ survive. All individuals in $\mathcal{L}^*$ have observed defect 1. By Lemma 36 their noisy fitness takes at most $2r+1$ distinct values, so $\mathcal{L}^*$ has at most $2r+1$ blocks. When $\mathcal{L}^*$ is processed, at least $\mu - 1 \ge 4M \ge 2(2r+1)$ individuals remain to be selected. By (R4) all endpoints survive.

Step 3: one endpoint of the block of V_j is an optimum of class j. By (i) the block of V_j contains an optimum of class j. By (R3) every other member of this block is an optimum of class j or has true defect two and a favorable evaluation, and by (ii) there is at most one member of the second kind. So at most one of the two endpoints is not an optimum of class j. If the block has a single member, it is the optimum of (i).

Step 4: the bound (30). Let $k_j \geq c_j(s)$ be the number of optima of class j in R_s. By (R1) and (R3), condition (i) fails for class j with probability $\eta^{k_j} \leq \eta^{c_j(s)}$. This is η if $c_j(s) = 1$ and at most η^2 if $c_j(s) \geq 2$. Summing over $j \in \mathcal{S}_s$ gives at most $\eta A_s + M\eta^2$. A string of true defect 2 has a favorable evaluation with probability at most $4/n$, and one of true defect 1 with probability at most $2/n$. So (ii) or (iii) fails with probability at most

$$\binom{2\mu}{2}\Big[\Big(\frac{4}{n}\Big)^2 + \Big(\frac{2}{n}\Big)^2\Big] \;\leq\; \frac{40\mu^2}{n^2}.$$

If (ii), (iii), and (i) for every $j \in \mathcal{S}_s$ hold, there is no loss. □

Lemma 48. *Let μ be even, $4M + 1 \leq \mu \leq n/32$, and n large. For every generation s and every state of the algorithm the following holds.*

(a) Gain. *If $\emptyset \neq \mathcal{S}_s \neq \{0, \ldots, r\}$, let j be a missing class that is adjacent to a present class. Then*

$$\mathrm{P}\big[\, j \in \mathcal{S}_{s+1} \text{ and no loss at } s \;\big|\; \mathcal{F}_s \big] \;\geq\; g \;=\; \frac{1 - p_c}{32\mu n^2}.$$

In particular a gain occurs at s with conditional probability at least g.

(b) Doubling a single copy. *If $c_j(s) = 1$, then*

$$\mathrm{P}\big[c_j(s+1) \geq 2 \mid \mathcal{F}_s\big] \;\geq\; q \;=\; \frac{\theta}{8\mu}.$$

(c) Classes with a single copy. *With NO_s the number of non-optimal members of P_s,*

$$\mathbb{E}\big[A_{s+1} \mid \mathcal{F}_s\big] \;\leq\; (1 - q)\, A_s + \frac{17Mr\mu}{n} + 4\,\mathrm{NO}_s.$$

Proof. In all three parts the variation step and the noise of generation s are independent by (R1).

(a), variation. Let $z \in P_s$ be an optimum of a class h with $|h - j| = 1$. In one round, both samples of the first tournament are the position of z with probability μ^{-2}; there is no crossover with probability $1 - p_c$; and mutation flips exactly the two bits of one prescribed pair of z (a pair 00 if $j = h + 1$, a pair 11 if $j = h - 1$) with probability $n^{-2}(1 - 1/n)^{n-2} \geq 1/(4n^2)$. Then the first child y is an optimum of class j. By (26), some of the $\mu/2$ rounds does this with probability at least

$$\frac{1}{2} \cdot \frac{\mu}{2} \cdot \frac{1 - p_c}{4\mu^2 n^2} \;=\; \frac{1 - p_c}{16\mu n^2} \;=\; 2g.$$

(a), noise. Require that the noise of y falls into u, and that (i)–(iii) of Lemma 47 hold for all classes of $\mathcal{S}_s$. By the computation in Step 4 of that proof, this fails with probability at most $\eta + M\eta + 40\mu^2/n^2 \leq \frac{1}{2}$ for n large, because $\mu \leq n/32$. On this event, Lemma 47 applies to every class of $\mathcal{S}_s$ and also to class j, with y as the optimum of condition (i). So no class is lost and $j \in \mathcal{S}_{s+1}$. The total probability is at least $2g \cdot \frac{1}{2} = g$.

(b). Let z be the unique optimum of class j in P_s. In one round, the first child is an unchanged copy of z with probability at least θ/μ^2 (both samples are the position of z, no crossover, no bit flipped). By (26) some round does this with probability at least

$$1 - \Big(1 - \frac{\theta}{\mu^2}\Big)^{\mu/2} \;\geq\; \frac{\theta}{4\mu}.$$

Then R_s contains two optima z, z' of class j. Require that both have their noise in u and that no member of R_s of true defect 1 or 2 has a favorable evaluation. This fails with probability at most $2\eta + 2\mu \cdot 4/n = 2\eta + 8\mu/n \leq \frac{1}{2}$ for n large. On this event no individual has observed defect 0,

so the layer $\mathcal{L}^*$ of Lemma 47 is the first layer. By (R3) the block of V_j consists of optima of class j only, and it has at least two members. Its two endpoints are therefore two optima of class j. As in Step 2 of Lemma 47, both survive. So $c_j(s+1) \geq 2$ with probability at least $\frac{\theta}{4\mu} \cdot \frac{1}{2} = q$.

(c). A class j counted by A_{s+1} has $c_j(s) = 1$, or $c_j(s) \geq 2$, or $c_j(s) = 0$. We bound the three contributions.

(c1): $c_j(s) = 1$. By (b), $c_j(s+1) = 1$ has probability at most $1 - q$. These classes contribute at most $(1-q)A_s$.

(c2): $c_j(s) \geq 2$. Then R_s contains $k_j \geq 2$ optima of class j. Suppose that at least two of them have their noise in u, that no member of true defect 2 has a favorable evaluation, and that (iii) of Lemma 47 holds. Then the block of V_j consists of at least two optima of class j, and both endpoints survive by Step 2 of Lemma 47; so $c_j(s+1) \geq 2$. The first requirement fails with probability at most 2η: the failure probability is largest for $k_j = 2$, where it equals $2\eta - \eta^2$. The second fails with probability at most $2\mu \cdot 4/n$, and the third with probability at most $\binom{2\mu}{2}(2/n)^2 \leq 8\mu^2/n^2$. Using $\mu \leq n/32$,

$$2\eta + \frac{8\mu}{n} + \frac{8\mu^2}{n^2} \;\leq\; \frac{4r}{n} + \frac{8\mu}{n} + \frac{\mu}{4n} \;\leq\; \frac{14r\mu}{n}.$$

These classes contribute at most $M \cdot 14r\mu/n$.

(c3): $c_j(s) = 0$. Then some offspring is an optimum of a class that is absent from P_s. We bound the expected number of such offspring. Consider first a round in which all four tournament samples are optima. Both parents have $u = 1 \cdots 1$, so a crossover with a cut point outside w returns the two parents. The cut point falls inside w with probability at most $2r/(n-1)$. Otherwise the child is a mutant of an optimum of a present class, and it is an optimum of another class only if mutation flips both bits of some pair, which has probability at most r/n^2. So each such offspring counts with probability at most $2r/(n-1) + r/n^2 \leq 3r/n$, in total at most $3r\mu/n$. A round has a non-optimal tournament sample with probability at most $4\mathrm{NO}_s/\mu$, and it produces two offspring. Over the $\mu/2$ rounds this contributes at most $4\mathrm{NO}_s$.

Adding (c1)–(c3) gives the claim, because $14Mr\mu/n + 3r\mu/n \leq 17Mr\mu/n$. □

A *stage* is a window of $L_2 = \lfloor n^2/(\Lambda\mu^3) \rfloor$ generations with $\Lambda = 16(4Mr^2 + 40)$. The next lemma needs $L_2 \geq 1$, that is $\Lambda\mu^3 \leq n^2$.

Lemma 49. *Let $r \geq 1$ and $p_c \in [0, 1)$ be constants, n large, and μ even with*

$$C_0 \ln n \;\leq\; \mu \;\leq\; \frac{n}{32}, \qquad \mu \geq 4M + 1, \qquad \Lambda\mu^3 \leq n^2.$$

Start at time 0 from a saturated population that does not cover the front. Let Z be the event that for some $s < L_2$: the generations $0, \ldots, s$ are good, no loss occurs at the generations $0, \ldots, s-1$, and a gain occurs at s. Then

$$\mathrm{P}[Z] \;\geq\; \frac{gL_2}{2}.$$

On Z, let s be the first generation witnessing this event. Then P_{s+1} is saturated and $\mathcal{S}_{s+1} \supsetneq \mathcal{S}_0$.

Proof. Let G_s be the event that the generations $0, \ldots, s-1$ are good; G_0 is the sure event and G_s is $\mathcal{F}_s$-measurable. On G_s the population P_s is saturated by Lemma 43(i), so $\mathcal{S}_s \neq \emptyset$. By Lemma 42, $\mathrm{P}[G_s^c] \leq s\,\delta_{\mathrm{bad}}$. For n large we have

$$\delta_{\mathrm{bad}} \leq \frac{g}{16}, \qquad L_2\,\delta_{\mathrm{bad}} \leq \frac{1}{16}, \qquad gL_2 \leq \frac{1-p_c}{32\Lambda\mu^4} \leq \frac{1}{32}, \tag{31}$$

because $\delta_{\mathrm{bad}} \leq n^{-11}$, $g \geq (1-p_c)n^{-3}$ and $L_2 \leq n^2$.

Step 1: few classes have a single copy. Put $\alpha_s := \mathbb{E}[A_s \mathbf{1}_{G_s}]$. Trivially $\alpha_s \leq M$. Let $s \geq 1$. On G_s generation $s-1$ is good and P_{s-1} is saturated, so $\mathbb{E}[\mathrm{NO}_s \mathbf{1}_{G_s}] \leq 8\mu/n$ by (29). Since $G_{s+1} \subseteq G_s$, Lemma 48(c) gives

$$\alpha_{s+1} \;\leq\; \mathbb{E}\big[\mathbf{1}_{G_s}\mathbb{E}[A_{s+1} \mid \mathcal{F}_s]\big] \;\leq\; (1-q)\alpha_s + \frac{17Mr\mu}{n} + \frac{32\mu}{n} \;\leq\; (1-q)\alpha_s + b,$$

with $b = 50Mr\mu/n$, because $Mr \geq 2$. Hence $\alpha_s \leq M(1-q)^{s-1} + b/q$ for $s \geq 1$, and

$$\sum_{s<L_2} \alpha_s \;\leq\; M + \frac{M}{q} + \frac{L_2 b}{q} \;\leq\; \frac{2M}{q} + \frac{L_2 b}{q}.$$

We multiply by $\eta = 2r/n$ and use $1/q = 8\mu/\theta$, $L_2 \leq n^2/(\Lambda\mu^3)$ and $\Lambda \geq 64Mr^2$:

$$\eta \cdot \frac{2M}{q} \;=\; \frac{32rM\mu}{\theta n} \;\leq\; \frac{1}{16}, \qquad \text{since } \mu \leq n^{2/3} \text{ and } n \text{ is large;}$$

$$\eta \cdot \frac{L_2 b}{q} \;\leq\; \frac{2r}{n} \cdot \frac{n^2}{\Lambda\mu^3} \cdot \frac{50Mr\mu}{n} \cdot \frac{8\mu}{\theta} \;=\; \frac{800Mr^2}{\Lambda\theta\mu} \;\leq\; \frac{1}{16}, \quad \text{since } \mu \geq C_0 \ln n \text{ and } n \text{ is large.}$$

So $\eta \sum_{s<L_2} \alpha_s \leq \frac{1}{8}$.

Step 2: a loss is unlikely. By (30), the expected number of generations $s < L_2$ such that G_s holds and a loss occurs at s is at most

$$\sum_{s<L_2} \mathbb{E}[p_s \mathbf{1}_{G_s}] \;\leq\; \eta \sum_{s<L_2} \alpha_s + L_2\Big(M\eta^2 + \frac{40\mu^2}{n^2}\Big)$$
$$\leq\; \frac{1}{8} + \frac{4Mr^2 + 40\mu^2}{\Lambda\mu^3} \;\leq\; \frac{1}{8} + \frac{1}{16\mu} \;\leq\; \frac{1}{4}.$$

Let NL be the event that G_{L_2} holds and no loss occurs at any $s < L_2$. By Markov's inequality and (31),

$$\mathrm{P}[NL] \;\geq\; 1 - L_2\delta_{\mathrm{bad}} - \frac{1}{4} \;\geq\; \frac{11}{16}.$$

Step 3: the first gain. Let E_s be the event that G_{s+1} holds, $\mathcal{S}_0 = \mathcal{S}_1 = \cdots = \mathcal{S}_s$, and a gain occurs at s. The events E_s are disjoint, and $Z = \bigcup_{s<L_2} E_s$: on Z the first change of $\mathcal{S}$ is not a loss, so it is a gain. Let NC_s be the $\mathcal{F}_s$-measurable event that G_s holds and $\mathcal{S}_0 = \cdots = \mathcal{S}_s$. On NC_s the set $\mathcal{S}_s = \mathcal{S}_0$ is neither empty nor full, so Lemmas 48(a) and 42 give $\mathrm{P}[E_s \mid \mathcal{F}_s] \geq g - \delta_{\mathrm{bad}} \geq \frac{15}{16} g$. The events NC_s decrease in s. With $x := \mathrm{P}[Z]$,

$$x \;=\; \sum_{s<L_2} \mathrm{P}[E_s] \;\geq\; \frac{15}{16}\, g \sum_{s<L_2} \mathrm{P}[NC_s] \;\geq\; \frac{15}{16}\, gL_2\, \mathrm{P}[NC_{L_2}].$$

On NL, either $\mathcal{S}$ never changes before L_2, or its first change is a gain and Z occurs. So $NL \subseteq NC_{L_2} \cup Z$ and $\mathrm{P}[NC_{L_2}] \geq \frac{11}{16} - x$. Therefore

$$x \;\geq\; \frac{15}{16}\, gL_2\Big(\frac{11}{16} - x\Big), \qquad \text{so} \qquad x \;\geq\; \frac{165}{256} \cdot \frac{gL_2}{1 + gL_2} \;\geq\; \frac{165}{256} \cdot \frac{32}{33}\, gL_2 \;=\; \frac{5}{8}\, gL_2,$$

where we used $gL_2 \leq 1/32$ from (31).

Step 4: the state after the gain. On E_s generation s is good and P_s is saturated, so P_{s+1} is saturated by Lemma 43(i). Also $\mathcal{S}_{s+1} \supsetneq \mathcal{S}_s = \mathcal{S}_0$. □

F.7 Proof of the Main Bound

The following theorem gives the precise form of Theorem 4(a).

Theorem 10. *Let $r \geq 1$ and $p_c \in [0, 1)$ be constants, $M = r + 1$, $p = 1$, and let n be large. Let μ be even with*

$$C_0 \ln n \;\leq\; \mu \;\leq\; \frac{n}{32}, \qquad \mu \geq 4M + 1, \qquad \Lambda\mu^3 \leq n^2,$$

where $C_0 = C_0(r, p_c)$ and $\Lambda = 16(4Mr^2 + 40)$ are as in the notation table. Put

$$K_0 := (M-1)\, L_2\Big(\frac{128\Lambda\mu^4}{1 - p_c}\Big)^{M-1}, \qquad K := \Big\lceil 8\big(nL_0 + 1 + L_1 + K_0\big)\Big\rceil.$$

Then $K \leq c(r, p_c)\, n^2\mu^{4M-7}$, and the reevaluating NSGA-II on PairedCOCZ(n, r) *satisfies,* from every initial population*,*

$$\mathbb{E}[T] \;\leq\; 2K \;=\; O\big(n^2\mu^{4M-7}\big) \qquad \textit{and} \qquad \mathrm{P}[\,T > kK\,] \;\leq\; 2^{-k} \quad \textit{for every integer } k \geq 1.$$

The bounds also hold conditionally on every generation-start history, including an arbitrary fixed initial stored state as specified above Section F.1; in that case T is the remaining runtime. For $\mu = \Theta(\log n)$ this is $n^2 \operatorname{polylog}(n)$.

Proof. Step 0: the size of K. Since $L_2 \le n^2/(\Lambda\mu^3)$,

$$K_0 \;\le\; (M-1)\,\frac{n^2}{\Lambda\mu^3}\Big(\frac{128\Lambda}{1-p_c}\Big)^{M-1}\mu^{4M-4} \;=\; c'(r,p_c)\,n^2\mu^{4M-7}.$$

Moreover $nL_0+1+L_1 = O(n\mu+n\ln n)$ with a constant that depends on p_c only, and $n\mu+n\ln n \le 2n^2\mu^{4M-7}$ because $4M-7 \ge 1$. In particular $K \le n^{4M}$ for n large.

Step 1: early phase. By Lemma 41 and Markov's inequality,

$$\mathrm{P}[\tau_0 > 8nL_0] \;\le\; \mathrm{P}[\mathcal{F}_G] + \frac{\mathbb{E}[\tau_0 \mathbf{1}_{\mathcal{F}_G^c}]}{8nL_0} \;\le\; \frac{1}{n} + \frac{1}{8}.$$

Otherwise the population at time $\tau_0 + 1 \le 8nL_0 + 1$ has low defect. Restart the clock there.

Step 2: no bad generation. By Lemma 42, $\mathrm{P}[\tau_{\mathrm{bad}} < K] \le K\delta_{\mathrm{bad}} \le n^{4M}\cdot n^{-(2r+2)M-3} \le 1/n$, because $2r+2 \ge 4$.

Step 3: descent. By Lemma 44 and Markov's inequality, $\mathrm{P}[\tau \wedge \tau_{\mathrm{bad}} > 8L_1] \le \frac{1}{8}$. Moreover P_t is saturated for $\tau \le t \le \tau_{\mathrm{bad}}$.

Step 4: stages and attempts. If coverage or a bad generation has already been detected, there is nothing to count. Otherwise start at time τ from a saturated population. A *stage* succeeds at the first witness to the event Z in Lemma 49: generations $0,\dots,s$ of its local clock are good, none of the earlier generations loses a class, and a gain occurs at $s < L_2$. It ends at generation start $s+1$, with saturation preserved and at least one additional class. Stop earlier if coverage or a bad generation is detected; otherwise declare failure after L_2 generations. All these endpoints are generation-start stopping times, since generation s and its goodness are fully revealed in $\mathcal{F}_{s+1}$.

An *attempt* ends at the first failed stage, at coverage, or at detection of a bad generation. At most $M-1$ successful stages cover all classes, so an attempt lasts at most $(M-1)L_2$ generations. At every stage start, its conditional probability of success or earlier coverage or bad-generation detection is at least $gL_2/2$ by Lemma 49. Iterating these bounds, which hold from every history, an attempt reaches coverage or detects a bad generation with conditional probability at least

$$\varsigma \;:=\; \Big(\frac{gL_2}{2}\Big)^{M-1} \;\ge\; \Big(\frac{1-p_c}{128\Lambda\mu^4}\Big)^{M-1}.$$

Here we used $L_2 \ge n^2/(2\Lambda\mu^3)$, which holds because $n^2/(\Lambda\mu^3) \ge 1$. If an attempt fails without either terminal event, all its generations were good and the population is still saturated (Lemma 43(i)); the next attempt starts from that revealed state. Let T_{att} count generations from τ through coverage or bad-generation detection (zero if already stopped). Attempt number $i+1$ is started with probability at most $(1-\varsigma)^i$, so the expected number of attempts is at most $1/\varsigma$ and

$$\mathbb{E}[T_{\mathrm{att}}] \;\le\; \frac{(M-1)L_2}{\varsigma} \;\le\; K_0.$$

By Markov's inequality, $\mathrm{P}[T_{\mathrm{att}} > 8K_0] \le \frac{1}{8}$.

Step 5: one block of K generations. The failure probabilities of Steps 1–4 sum to at most $\frac{2}{n}+\frac{3}{8} \le \frac{1}{2}$. Outside these events no generation among the first K after the restart is bad. Hence $\tau \le 8L_1$, the attempts are not stopped by a bad generation before time $8L_1 + 8K_0 \le K$, and the front is covered at the latest at time $(8nL_0+1)+8L_1+8K_0 \;\le\; K$. So $\mathrm{P}[T > K] \le \frac{1}{2}$ from every initial population.

Step 6: restart. The early phase and all subsequent estimates hold for every fixed stored state, as explained in the histories-and-restarts paragraph. Apply Step 5 conditionally at generation starts kK: on $\{T > kK\}$ its remaining failure probability is at most $1/2$. Hence $\mathrm{P}[T > kK] \le 2^{-k}$ and $\mathbb{E}[T] \le K\sum_{k\ge 0} 2^{-k} = 2K$. □

G Proof of Theorem 4(b): Runtime Lower Bound for the No-Reevaluation Strategy on PairedCOCZ

We first quantify a permanent trap at $p = 1$, then transfer a short trapping event to $p = 1-q$. All trapping arguments allow crossover and arbitrary ordering of equal stored values. The finite upper

bound for no reevaluation at $p < 1$ uses stable sorting with parents before offspring, as specified in Section 2.3. The final sections also transfer the reevaluation upper bound and assemble Theorem 4 for $p < 1$.

G.1 Absorbing Populations at $p = 1$

Let $n > 2r \geq 2$, let $\mu \geq 2$ be even, and put $M = r + 1$. At $p = 1$ an individual stores the true fitness of a Hamming neighbour. By Lemmas 34 and 35, a stored front value strictly dominates every off-front value, and an optimum always has noisy fitness equal to a one-defect value. Hence every individual that stores a front value is a one-defect, nonoptimal string whose evaluation repaired its defect. Let H_t be the number of individuals in P_t that store a front value, let N_t count such individuals among the offspring of generation t, and set $\tau_{\mathrm{abs}} = \inf\{t : H_t = \mu\}$.

Lemma 50. *For every $p_c \in [0, 1]$ and every rule for ordering equal stored values,*

$$H_{t+1} = \min\{H_t + N_t, \mu\} \geq H_t.$$

Once $H_t = \mu$, all subsequent populations consist of individuals storing front values and contain no optimum. Thus $T = \infty$ if coverage has not occurred before τ_{abs}.

Proof. The individuals storing front values form the first non-dominated layer and strictly dominate every other member of the combined population. Selection therefore retains all of them if there are at most μ, and only such individuals otherwise. Stored values never change, so this property persists. None of them is optimal at $p = 1$. □

Theorem 11. *At $p = 1$, from every initial population that does not cover the front, for every $p_c \in [0, 1]$ and every rule for ordering equal stored values,*

$$\mathrm{P}[T = \infty] \geq a_n := n^{-(n+1)\mu} > 0, \qquad \mathbb{E}[T] = \infty.$$

In particular, $\mathrm{P}[T > B] \geq a_n$ for every finite $B \geq 0$.

Proof. Fix a one-defect string h. Whatever the selected parents and crossover outcome, bit-wise mutation produces h with probability at least n^{-n}: every bit takes its prescribed value with probability at least $1/n$. Its evaluation repairs the defect with probability at least $1/n$. Thus the probability that all μ offspring of generation 0 store front values is at least $n^{-(n+1)\mu}$. On this event $H_1 = \mu$, and coverage did not occur at time 0, so Lemma 50 gives $T = \infty$. For every finite $B > 0$, $\mathbb{E}[T] \geq B\mathrm{P}[T \geq B] \geq Ba_n$; letting $B \to \infty$ gives the expectation claim. □

G.2 Absorption from Uniform Initialization

The preceding bound may be very small. A better bound from uniform initialization suffices for an exponential expected runtime at noise levels slightly below 1.

Lemma 51. *Assume $p = 1$, $p_c < 1$, uniform initialization, and $(\mu - 1)2^{r-n} \leq 1/2$. Put*

$$\alpha_n := \frac{2^{-n}}{2n}(1 - p_c)^{\mu/2}(4\mu n)^{-\mu}.$$

There is an event E, determined by initialization and generation 0, with $\mathrm{P}[E] \geq \alpha_n$, on which P_0 contains no optimum and every member of P_1 stores a front value. In particular, $\mathrm{P}[T = \infty] \geq \alpha_n$.

Proof. Fix a string h whose w-pairs are all 00 and whose u-part has exactly one zero. Require the first initial individual to be h and its noise bit to repair that zero, and require all other initial strings to be nonoptimal. There are 2^r optima, so the probability of these requirements is

$$2^{-n}n^{-1}(1 - 2^{r-n})^{\mu-1} \geq \frac{2^{-n}}{2n}.$$

There is now at least one individual storing a front value and no optimum. Each tournament selects such an individual with probability at least $1/\mu$. In each round of generation 0, require both selected parents to be such individuals, no crossover, no mutation, and a repairing noise bit for each child. Conditionally on the initial state, all these requirements have probability at least

$$(1 - p_c)^{\mu/2}\mu^{-\mu}(1 - 1/n)^{n\mu}n^{-\mu} \geq (1 - p_c)^{\mu/2}(4\mu n)^{-\mu}.$$

All offspring store front values, hence so does every member of P_1. Lemma 50 completes the proof. □

G.3 Expected Runtime without Reevaluation for $p < 1$

Write $p = 1 - q$ with $0 < q < 1$. Couple the p-run with a $p = 1$ run using the same initial strings, variation randomness, tournament and survival tie-breaking choices, and noise-bit positions. Both runs use the same deterministic stable sorting rule. Independently mark each evaluation as noisy with probability p. The p-run flips its noise bit only when marked, whereas the $p = 1$ run always does. The runs agree until the first unmarked evaluation.

Proposition 1. *Under the initialization and parameter assumptions of Lemma 51, but with $p = 1-q$, for every integer $B \geq 1$,*

$$\mathrm{P}[T > B] \geq \alpha_n p^{\mu(B+1)}, \qquad \mathbb{E}[T] \geq \frac{\alpha_n p^{2\mu}}{1 - p^{\mu}}.$$

If $\mu q \leq 1/4$, then $\mathbb{E}[T] \geq \alpha_n/(2\mu q)$. For fixed r and fixed $p_c < 1$, if $\mu \log n = o(n)$ and $0 < q \leq 2^{-2n}$, this gives

$$\mathbb{E}[T] \geq 2^{n - o(n)}.$$

Proof. The trap event E of Lemma 51 depends only on randomness used by the $p = 1$ run, hence is independent of the extra marks. If the first $\mu(B + 1)$ marks are noisy, the two runs coincide through P_B. On E, neither run has covered the front by then. Consequently $\mathrm{P}[T > B] \geq \mathrm{P}[E]p^{\mu(B+1)}$. For a nonnegative integer-valued runtime, summing tail probabilities gives

$$\mathbb{E}[T] = \sum_{B=0}^{\infty} \mathrm{P}[T > B] \geq \alpha_n \sum_{B=1}^{\infty} p^{\mu(B+1)} = \frac{\alpha_n p^{2\mu}}{1 - p^{\mu}}.$$

If $\mu q \leq 1/4$, Bernoulli's inequality yields $p^{2\mu} \geq 1 - 2\mu q \geq 1/2$ and $1 - p^{\mu} \leq \mu q$. Finally,

$$-\log_2 \alpha_n = n + \log_2(2n) + \mu \log_2(4\mu n) + \frac{\mu}{2} \log_2 \frac{1}{1 - p_c} = n + o(n)$$

under the stated assumptions, and the exponential bound follows. □

Proposition 2. *Consider NSGA-II without reevaluation on* PairedCOCZ(n, r) *with $p = 1 - q$, $0 < q \leq 1$, $\mu \geq 2M$, and the stable sorting rule of Section 2.3. From every initial population, for every $p_c \in [0, 1]$,*

$$\mathbb{E}[T] \leq (qn^{-n})^{-\mu} < \infty.$$

Proof. Fix any revealed history at the start of a generation before coverage. Prescribe μ offspring strings that are all optima and include every front class, and require all their creation evaluations to be noise-free. Conditional on the history, the parent choices and the crossover outcomes, mutation produces each prescribed child with probability at least n^{-n}. The mutation bits and evaluation noises of different children are independent. Thus this event has conditional probability at least $s_0 := (qn^{-n})^{\mu}$.

On this event every class has at least one optimum with an uncorrupted evaluation among the offspring. Every off-front stored value is strictly dominated by every front value, so the first layer consists of the M front-value blocks. In each block, all parents precede all offspring, and every offspring is an optimum of that class. Its last member y_j is therefore an optimum, even if some parents in the block hold an incorrect stored value. The two stable objective sorts have the same last member y_j in each block. There are at most $2M \leq \mu$ positive-crowding members in the first layer, and each y_j has positive normalized crowding by the same argument as (R4), including the zero-range case. Consequently all y_j survive, and the next population covers the front.

The conditional success probability in each generation is therefore at least s_0, after every unsuccessful history. Hence, with $s = s_0$, $\mathrm{P}[T > k] \leq (1 - s)^k$ for every integer $k \geq 0$, and summing the tail gives $\mathbb{E}[T] \leq s^{-1} = (qn^{-n})^{-\mu}$. □

G.4 Runtime Upper Bound with Reevaluation for $p < 1$

Theorem 12. *Assume the conditions on r, p_c, n, μ in Theorem 10, and let $E_{\max} := 2K = O(n^2\mu^{4M-7})$ be its expected-runtime bound at $p = 1$, which holds from every generation-start*

history. Put $L = \lceil 4(E_{\max} + 1) \rceil$. *If*

$$0 \le q \le \frac{1}{8\mu(L+1)},$$

then at $p = 1 - q$ *the reevaluating NSGA-II satisfies* $\mathbb{E}[T] \le 2L = O(n^2 \mu^{4M-7})$ *from every initial population and conditionally on every generation-start history. In the conditional statement,* T *is the remaining runtime from that start.*

Proof. Fix a generation-start history of the p-run and start a reference $p = 1$ continuation from the same stored state, without rerunning initialization. The full-state bound of Theorem 10 gives expected remaining runtime at most $E_{\max}$ for this continuation. Markov's inequality therefore gives $\mathrm{P}[T^{(1)} \le L] \ge 3/4$. Couple the two continuations using the same variation, tie-breaking choices and noise-bit positions, and use the same stable sorting rule. The next L generations use $2\mu L$ fresh evaluations, so their success probabilities differ by at most $2\mu L q \le 2\mu(L+1)q \le 1/4$. Thus every block of L generations succeeds with conditional probability at least $1/2$. Iterating this bound yields $\mathrm{P}[T > kL] \le 2^{-k}$ and $\mathbb{E}[T] \le 2L$. □

G.5 Proof of Theorem 4

Proof of Theorem 4. Fix r, $p_c < 1$ and $0 < \delta < 2/3$. For sufficiently large $C = C(r, p_c)$ and $n \ge n_0(r, p_c, \delta)$, the range $C \ln n \le \mu \le n^{2/3-\delta}$ implies $\mu \ge 4M+1$, $\mu \le n/32$ and $\Lambda\mu^3 \le n^2$. Part (a) follows from Theorem 10; part (b) is Theorem 11. With $L = \lceil 4(E_{\max} + 1) \rceil$ as in Theorem 12, $\mu(L+1) = O(n^2\mu^{4M-6})$ is polynomial in n, because M is fixed and $\mu \le n^{2/3-\delta}$. Hence, for all sufficiently large n, $2^{-2n} \le 1/(8\mu(L+1))$, so that theorem preserves the reevaluating upper bound for $0 < q \le 2^{-2n}$. The same population bound implies $(\mu-1)2^{r-n} \le 1/2$, $\mu \log n = o(n)$, and $\mu q \le 1/4$. Under uniform initialization, Proposition 1 therefore gives the lower bound $2^{n-o(n)}$ without reevaluation, while Proposition 2 shows that its expected runtime is finite. □

H Experimental Details and Additional Results

H.1 Setup

All runs use Algorithm 1: binary tournaments on the stored ranks and crowding distances, one-point crossover with probability $p_c = 0.5$ and a uniform cut, and bit-wise mutation with rate $1/n$. Noise changes the evaluated string, never the stored one. Without reevaluation, each offspring is evaluated once, when it is created; with reevaluation, all 2μ individuals of R_t are evaluated once per generation, and offspring are not evaluated at creation. In each non-dominated layer, parents precede offspring and both objective sorts are stable; the initial population keeps the order in which it was generated. Crowding distances are normalized within each layer, tournament ties are decided by a fair coin, and ties at the crowding cutoff uniformly at random. Every run starts from a uniformly random population.

A run is complete once the current population contains a true optimum of every front value; noisy fitness values and optima that were found and later lost do not count. Completion is checked with the true fitness, at initialization and after every survival step, and these checks are not counted as evaluations. A run that completes in generation T therefore uses $\mu(1+T)$ evaluations without reevaluation and $\mu(1+2T)$ with it. Table 4 lists the 32 settings; each is run 1000 times with each strategy, with a separate seed for every run. Pilot runs, which are not part of the reported data, were used to choose sizes and budgets such that the runs without reevaluation on LOTZ and the runs with reevaluation on PairedCOCZ finish well within the budget. All settings, seeds and budgets were fixed before the reported runs. The settings illustrate the two effects; they are not a systematic search for a noise threshold.

H.2 Statistics

For each setting we report the number of completed runs, the empirical median and 90th percentile of the runtime over all runs, and the capped mean, that is, the mean of $\min\{T, H\}$, with its standard

Table 4: Experimental settings. Each setting is run 1000 times with and 1000 times without reevaluation, with $p_c = 0.5$ and uniform initialization; H is the limit in generations and $p_{\text{test}}(n)$ is defined in Section H.3.

Problem	n	μ	p	H
LOTZ	$8, 12, 16, 24, 32, 64$	$2n+2$	0.5	$20n^2$
LOTZ	$16, 32, 64, 128$	$2n+2$	$1/(8\ln n)$, $1/(8n^2)$, $p_{\text{test}}(n)$	$20n^2$
PairedCOCZ, $r=3$	40	32	0.5, 0.9, 0.99, 0.999, 1	10^4
PairedCOCZ, $r=3$	24	32	0.999, 1	10^4
PairedCOCZ, $r=3$	$24, 40, 64$	$2\lceil 4\ln n\rceil$	1	$10n^2$
PairedCOCZ, $r=3$	$128, 256$	18	$1, 1-2^{-2n}$	$10n^2$

error. A quantile that is not reached within the budget is marked $> H$, and one that is infinite because of runs certified as trapped (Section H.4) is marked ∞. The capped mean equals H when no run completes; it then only records that the budget was exhausted.

Completion curves carry simultaneous 95% confidence bands. By the Dvoretzky–Kiefer–Wolfowitz–Massart inequality (Massart, 1990) and a union bound over all $K = 64$ combinations of setting and strategy, each with $R = 1000$ runs, the empirical fraction of completed runs is within

$$\varepsilon_R = \sqrt{\frac{\ln(2K/\delta)}{2R}} \approx 0.0626, \qquad \delta = 0.05,$$

of the completion probability, simultaneously for all budgets up to H and all combinations. Runs that are unfinished at H do not affect this. For the difference between the two strategies we subtract the two bands. A budget of b evaluations corresponds to $\lfloor (b-\mu)/\mu \rfloor$ generations without reevaluation and $\lfloor (b-\mu)/(2\mu) \rfloor$ with it. All estimates use exact event counts; the straight segments in the figures only connect them.

H.3 Results on LOTZ

Table 5 gives the results at $p = 0.5$. All 6000 runs without reevaluation complete; at $n = 64$ the largest runtime is 3655 generations, far below $H = 81\,920$, and the median and 90th percentile (2175 and 2707) are close to the mean (2215). With reevaluation, the fraction of completed runs drops from 0.956 at $n = 8$ to 0.043 at $n = 12$ and is zero for $n \geq 16$, where the capped means equal H. The means without reevaluation grow roughly like n^2, well below the bound $O(n^4)$ of Theorem 1, although $p = 0.5$ is outside its noise range.

Table 5: LOTZ, $p = p_c = 0.5$ and $\mu = 2n+2$, with 1000 runs per row. R and NR denote reevaluation and no reevaluation. The capped mean is shown with one standard error; all runtimes and budgets are in generations. The quantiles refer to all runs, and $> H$ means that the corresponding quantile was not reached within the budget.

n	H	Policy	Completed runs	Capped mean $\pm$ SE	Median	90th pct.
8	1 280	NR	1000/1000	37 ± 1	33	62
8	1 280	R	956/1000	397 ± 11	295	927
12	2 880	NR	1000/1000	87 ± 1	82	133
12	2 880	R	43/1000	$2\,808 \pm 12$	$> H$	$> H$
16	5 120	NR	1000/1000	154 ± 2	147	225
16	5 120	R	0/1000	$5\,120 \pm 0$	$> H$	$> H$
24	11 520	NR	1000/1000	339 ± 3	327	465
24	11 520	R	0/1000	$11\,520 \pm 0$	$> H$	$> H$
32	20 480	NR	1000/1000	596 ± 4	580	783
32	20 480	R	0/1000	$20\,480 \pm 0$	$> H$	$> H$
64	81 920	NR	1000/1000	$2\,215 \pm 12$	2 175	2 707
64	81 920	R	0/1000	$81\,920 \pm 0$	$> H$	$> H$

Table 6 uses noise levels that decrease with n as in Theorems 1 and 2; all runs complete. For $p = 1/(8 \ln n)$, reevaluation needs 4.8–8.3% more generations on average. Here $p^2 n < 0.1$ for all tested n, which is far from the range of Theorem 3. For $p = 1/(8n^2)$, the two means differ by at most 1.1%. The factor $1/8$ was chosen for the experiments and does not meet the constants of either upper bound.

Table 6: LOTZ with decreasing noise, $\mu = 2n + 2$, $p_c = 0.5$ and $H = 20n^2$. Each strategy has 1000 runs per row; all complete. Means and standard errors are in generations.

Noise schedule	n	No reevaluation		Reevaluation	
		Completed	Mean ± SE	Completed	Mean ± SE
$1/(8 \ln n)$	16	1000	86 ± 1	1000	91 ± 1
$1/(8 \ln n)$	32	1000	339 ± 3	1000	355 ± 3
$1/(8 \ln n)$	64	1000	$1\,275 \pm 7$	1000	$1\,364 \pm 8$
$1/(8 \ln n)$	128	1000	$4\,756 \pm 19$	1000	$5\,148 \pm 21$
$1/(8n^2)$	16	1000	81 ± 1	1000	82 ± 1
$1/(8n^2)$	32	1000	322 ± 2	1000	326 ± 2
$1/(8n^2)$	64	1000	$1\,236 \pm 7$	1000	$1\,237 \pm 7$
$1/(8n^2)$	128	1000	$4\,677 \pm 19$	1000	$4\,666 \pm 19$

As a noise-free control, Table 7 uses

$$p_{\text{test}}(n) := \frac{1}{8847360(2n+2)n},$$

half the explicit noise bound of Theorem 7 at $p_c = 1/2$. The expected total number of corrupted evaluations in these 8000 runs is about 0.024, and none occurred. The means of the two strategies differ by less than 1%, as expected, since both strategies have the same runtime distribution at $p = 0$.

Table 7: LOTZ with $p = p_{\text{test}}(n)$, inside the explicit noise range of Theorem 7; $\mu = 2n + 2$, $p_c = 0.5$, $H = 20n^2$. Each strategy has 1000 runs per row; means and standard errors are in generations.

n	p_{test}	No reevaluation		Reevaluation	
		Completed	Mean ± SE	Completed	Mean ± SE
16	2.078×10^{-10}	1000	82 ± 1	1000	82 ± 1
32	5.352×10^{-11}	1000	330 ± 2	1000	330 ± 3
64	1.359×10^{-11}	1000	$1\,229 \pm 7$	1000	$1\,233 \pm 7$
128	3.423×10^{-12}	1000	$4\,650 \pm 18$	1000	$4\,656 \pm 20$

H.4 Results on PairedCOCZ

Table 8 varies the noise level at $n = 40$ and $\mu = 32$. Counted in generations, no reevaluation is faster at $p = 0.5$ and slower at $p = 0.9$, while both strategies complete every run. At $p = 0.99$ and $p = 0.999$, at most 11.3% of the runs without reevaluation complete, whereas the 90th percentile with reevaluation stays below 1400 generations. At $n = 24$ and $p = 0.999$, 169 runs without reevaluation and all 1000 runs with it complete. For $p < 1$, a population in which every individual stores a front value can still accept a true optimum after a noise-free evaluation, so the unfinished runs in this table are ordinary timeouts.

At $p = 1$, a population of nonoptimal individuals that all store front values is absorbing without reevaluation (Appendix G). We test for this state at initialization, at generations $1, 65, 129, \ldots$ and at the end of the budget, and count a run as a certified trap once it is detected, up to 63 generations after the population enters the state. Table 9 shows that 836 of the 1000 runs at $n = 24$ and 942 at $n = 40$ are trapped and all others complete, whereas reevaluation completes every run. With $\mu = 2\lceil 4 \ln n \rceil$ (Table 10), the fraction of trapped runs rises from 0.898 at $n = 24$ to 0.985 at $n = 64$, and reevaluation again completes every run.

Table 8: PairedCOCZ below maximal noise, with $r = 3$, $\mu = 32$, $p_c = 0.5$, $H = 10^4$ generations and 1000 runs per row. Notation is as in Table 5.

n	p	Policy	Completed runs	Capped mean ± SE	Median	90th pct.
40	0.5	NR	1000/1000	465 ± 9	398	818
40	0.5	R	1000/1000	600 ± 15	473	1 209
40	0.9	NR	1000/1000	784 ± 13	716	1 339
40	0.9	R	1000/1000	637 ± 15	518	1 258
40	0.99	NR	113/1000	$9\,191 \pm 79$	$> H$	$> H$
40	0.99	R	1000/1000	672 ± 16	544	1 342
40	0.999	NR	58/1000	$9\,424 \pm 73$	$> H$	$> H$
40	0.999	R	1000/1000	643 ± 15	518	1 250
24	0.999	NR	169/1000	$8\,317 \pm 118$	$> H$	$> H$
24	0.999	R	1000/1000	218 ± 5	175	433

Table 9: PairedCOCZ at $p = 1$, with $r = 3$, $\mu = 32$, $p_c = 0.5$ and $H = 10^4$; 1000 runs per row. Every run either completes or is a certified trap, and ∞ marks quantiles that are infinite because of the trapped runs.

n	Policy	Completed runs	Certified traps	Capped mean ± SE	Median	90th pct.
24	NR	164/1000	836/1000	$8\,367 \pm 117$	∞	∞
24	R	1000/1000	0/1000	215 ± 6	170	435
40	NR	58/1000	942/1000	$9\,425 \pm 73$	∞	∞
40	R	1000/1000	0/1000	640 ± 15	527	1 247

Table 10: PairedCOCZ with $p = 1$, $r = 3$, $p_c = 0.5$, $\mu = 2\lceil 4 \ln n \rceil$ and $H = 10n^2$. Each strategy has 1000 runs per row. Means are capped at H, with one standard error, in generations.

n	μ	No reevaluation			Reevaluation	
		Completed	Trapped	Mean ± SE	Completed	Mean ± SE
24	26	102	898	$5\,177 \pm 55$	1000	291 ± 7
40	30	51	949	$15\,188 \pm 111$	1000	704 ± 16
64	34	15	985	$40\,348 \pm 157$	1000	$1\,462 \pm 32$

Table 11 compares $p = 1$ with $p = 1 - 2^{-2n}$ at $n \in \{128, 256\}$ and $\mu = 18$. We sample a noise-free evaluation exactly, as the event that $2n$ random bits are all zero, rather than rounding $1 - 2^{-2n}$ to 1. Reevaluation completes all runs at both noise levels. Without reevaluation, all runs are trapped at $p = 1$, and all runs at $p = 1 - 2^{-2n}$ reach the budget. A run at $p = 1 - 2^{-2n}$ can differ from a coupled run at $p = 1$ within the budget only if one of its at most $\mu(1 + 2H)$ evaluations is noise-free; summed over all runs, this probability is below 10^{-67} at $n = 128$ and smaller at $n = 256$. The two noise levels therefore cannot be told apart at these sizes, and these runs do not probe the expected-runtime bound for $p < 1$.

None of the PairedCOCZ settings satisfies the population conditions of the explicit proof in Appendix F: it requires $\mu \geq C_0 \ln n$ with $C_0 \geq 6e\, 2^{91}/(1 - p_c)^8 \approx 1.0 \times 10^{31}$ at $p_c = 1/2$, as well as $\mu \leq n/32$ and $\Lambda\mu^3 \leq n^2$, and the populations $2\lceil 4 \ln n \rceil$ exceed $n^{2/3}$ at the tested sizes.

Table 12 compares the strategies at equal generation and equal evaluation budgets. At $p = 0.9$, the difference after 1000 generations is -7.5 points and its interval contains zero; after 32 032 evaluations, which is 500 generations with reevaluation, it is $+27.0$ points and its interval excludes zero. The mean costs of the completed runs are 25 131 evaluations without and 40 823 with reevaluation. At $p = 0.99$ and $p = 0.999$, reevaluation is better under both comparisons.

Table 11: PairedCOCZ at and exponentially close to maximal noise: $r = 3$, $\mu = 18$, $p_c = 0.5$, $H = 10n^2$, with 1000 runs per row. NR/R denote no reevaluation/reevaluation. Trapped runs, which occur only at $p = 1$, are certified never to complete; censored runs are unfinished at H. Capped means and standard errors are in generations.

n	p	Policy	Completed	Trapped	Censored	Capped mean ± SE
128	1	NR	0	1000	0	$163\,840 \pm 0$
128	1	R	1000	0	0	$11\,339 \pm 271$
128	$1 - 2^{-2n}$	NR	0	0	1000	$163\,840 \pm 0$
128	$1 - 2^{-2n}$	R	1000	0	0	$12\,003 \pm 292$
256	1	NR	0	1000	0	$655\,360 \pm 0$
256	1	R	1000	0	0	$45\,490 \pm 1\,123$
256	$1 - 2^{-2n}$	NR	0	0	1000	$655\,360 \pm 0$
256	$1 - 2^{-2n}$	R	1000	0	0	$47\,381 \pm 1\,173$

Table 12: PairedCOCZ, $n = 40$, $r = 3$, $\mu = 32$: completion-probability differences (no reevaluation minus reevaluation) in percentage points, with simultaneous 95% confidence intervals. Each setting has 1000 runs per strategy.

p	Equal generations: $t = 1000$	Equal evaluations: $b = 32\,032$
0.5	$+10.9$ $[-1.6, +22.9]$	$+40.5$ $[+28.0, +52.5]$
0.9	-7.5 $[-20.0, +5.0]$	$+27.0$ $[+14.5, +39.5]$
0.99	-74.9 $[-87.0, -62.4]$	-40.0 $[-52.1, -27.5]$
0.999	-76.3 $[-88.4, -63.8]$	-41.6 $[-53.7, -29.1]$

H.5 Runtime Distributions and Reproducibility

Figure 2 shows completion curves in generations for six settings. Unfinished runs remain in the denominator, and each curve ends at its budget. The supplement contains the source code, the configuration of every setting, the seed and outcome of every run, and the scripts that produce all tables and figures from these records.

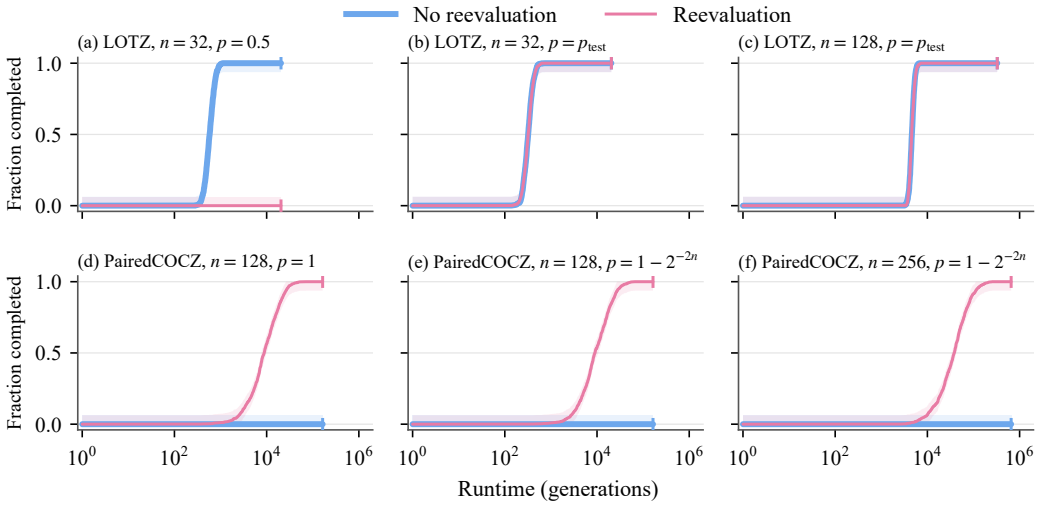

Figure 2: Fraction of runs that have covered the Pareto front within each generation budget, with simultaneous 95% confidence bands; 1000 runs per strategy and setting. Thick lines show no reevaluation and thin lines reevaluation; end marks indicate the budgets. The remaining parameters are as in Tables 5, 7 and 11.