Stochastic Queens Elimination

Artin Kavoosi

School of Industrial Engineering, College of Engineering, University of Tehran, Tehran, Iran

alireza.kavoosi@ut.ac.ir

Abstract

This research introduces the **Stochastic Sequential Queens Elimination Problem**, where on the n-queens board, each activated queen simultaneously attempts to eliminate all queens in her unblocked neighborhood, each independently succeeding with probability p. The objective is to minimize the expected cumulative conflict count over the trajectory. This research proposes a Markov decision process for this problem, as well as a heuristic method for $p < 1$, and a mixed-integer programming model when $p = 1$.

Key words: n-Queens problem; probabilistic sequential problem; Markov decision processes; mixed-integer programming; attack graph dynamics.

---

1. Introduction

The classical n-queens problem, as described beautifully by Bowtell and Keevash [6], is how to set n queens on a board so that no two can attack each other. Assuming $n = 8$ this problem has 92 solutions. It was first posed for the 8×8 board by the chess composer Max Bezzel in 1848 [5], and Nauck settled it two years later by exhibiting all 92 solutions by hand [28]. Yet despite this long history, almost the entire resulting literature, from these earliest enumerations through the modern asymptotic and algorithmic treatments, treats the board as a *static* combinatorial object, where queens are placed once, and the question asked is existence, counting, or optimality of a fixed configuration. Dynamic, elimination-based variants, in which the configuration itself evolves through a sequence of removals, have received much less attention.

Many real-world problems, by contrast, involve sequential and stochastic elimination instead of static placement, and are more naturally modeled as controlled stochastic processes than as one-shot constraint-satisfaction problems. Two established bodies of theory supply complementary language for such processes. On one hand, Markov decision process and stochastic-shortest-path theory [2, 31] give a rigorous account of sequential decision-making under uncertainty toward an absorbing goal, which is exactly the structure needed once "elimination" is treated as a controlled transition rather than a single static move. On the other hand, models of stochastic spread and removal on graphs, including the independent cascade model of influence maximization [21] and the contact process of interacting particle systems [18, 24], supply the probabilistic vocabulary of independent Bernoulli activation of neighbors and almost-sure extinction that is needed to describe how a stochastic kill can propagate through a population of competing agents. Consider a set of competing agents so that the removal of an activated agent removes a conflicting rival with probability $p \in (0, 1]$. The n-queens board, whose fixed attack relation already furnishes a well-understood conflict graph that has been exploited by local-search and metaheuristic solvers for the static problem [27, 33], is a clean, combinatorially rich model on which to instantiate such interactions.

The **Stochastic Sequential Queens Elimination Problem** (S-SQEP) proceeds as follows: Assuming a $n \times n$ board, each capture attempt by an activated queen succeeds independently with probability p. Our objective is to minimize expected cumulative conflict count over the trajectory. [1]

[1] The source code and additional discussion are available at `https://akavosi.github.io/blog/posts/sqep.html`.

Crucially, activation is a simultaneous multi-kill event when queen σ_t is activated, every queen in her unblocked neighborhood $N^{(t)}(\sigma_t)$ is independently targeted with probability p, so the completed kill set $Y(t)$ may contain zero, one, or multiple queens in a single step. This differs from sequential single-target capture models, in which an activated agent selects exactly one rival and attempts to eliminate it; the multi-kill structure introduces combinatorial dependence among conflict pairs and richer transition distributions. A toy example of this model is shown in Figure 1.

The rest of this paper is organized as follows. Section 2 explains the model description. Section 3 describes the structural properties of the proposed S-SQEP model. Section 4 presents a heuristic algorithm for $p < 1$. Section 5 discusses the proposed mixed-integer programming when $p = 1$. Section 6 collects further theoretical results, including kill-opportunity counts and cumulative-cost bounds. Section 7 gives some numerical results, and Section 8 concludes the paper.

2. Model

2.1. Markov decision process

Let $\mathscr{B} = \{1,\dots,n\}^2$. Let $Q^{(0)} = \{q_1,\dots,q_n\}$ be the initial set of queens, with injective positions $\ell^{(0)} : Q^{(0)} \to \mathscr{B}$. Queens q_i, q_j are *in mutual attack* at time t, written $q_i \overset{\text{atk}}{\sim} q_j$, if they share a row, column, or diagonal. Assume the intermediate-cell set $\text{Block}(q_i,q_j) \subset \mathscr{B}$ as the (possibly empty) open segment between their cells. Then

$$\text{CanKill}^{(t)}(q_i,q_j) \;=\; \mathbf{1}[q_i \overset{\text{atk}}{\sim} q_j] \cdot \mathbf{1}\big[\text{Block}(q_i,q_j)^{(t)}(q_i,q_j) = \emptyset\big].$$

State, policy, and objective. The *board state* is $S^{(t)} = (Q^{(t)}, \ell^{(t)})$. A *(Markov) activation policy* is a measurable map $\pi : \mathscr{S} \to Q$ with $\sigma_t = \pi(S^{(t)}) \in Q^{(t)}$, making σ_t measurable with respect to the filtration $\mathscr{F}_{t-1} = \sigma\Big(S^{(0)}, \xi^{(1)}, \dots, \xi^{(t-1)}\Big)$.

Definition 1 (S-SQEP)**.** *Fix $p \in (0,1]$ and an initial configuration $(Q^{(0)}, \ell^{(0)})$, the* Stochastic SQEP *is the Markov Decision Process (MDP) $\mathscr{M} = (S, \mathscr{A}, \mathscr{T}, R, H)$ with state space $\mathscr{S} = \{(Q,f) : Q \subseteq Q^{(0)}, f : Q \to \mathscr{B}$ injective$\}$; action space $\mathscr{A}(S^{(t)}) = Q^{(t)}$; transition $\mathscr{T}$ defined by the stochastic kill mechanism above; The stage cost is $c\big(S^{(t)}, \sigma_t, S^{(t+1)}\big) := \kappa\big(S^{(t)}\big)$, where for state $s = (Q,f) \in \mathscr{S}$. For state $s = (Q,f)$, the* conflict count *is the number of unordered unblocked attacking pairs:*

$$\kappa(s) \;:=\; \Big|\Big\{\{i,j\} : i < j,\ \text{CanKill}^{(t)}(q_i,q_j) = 1\Big\}\Big|.$$

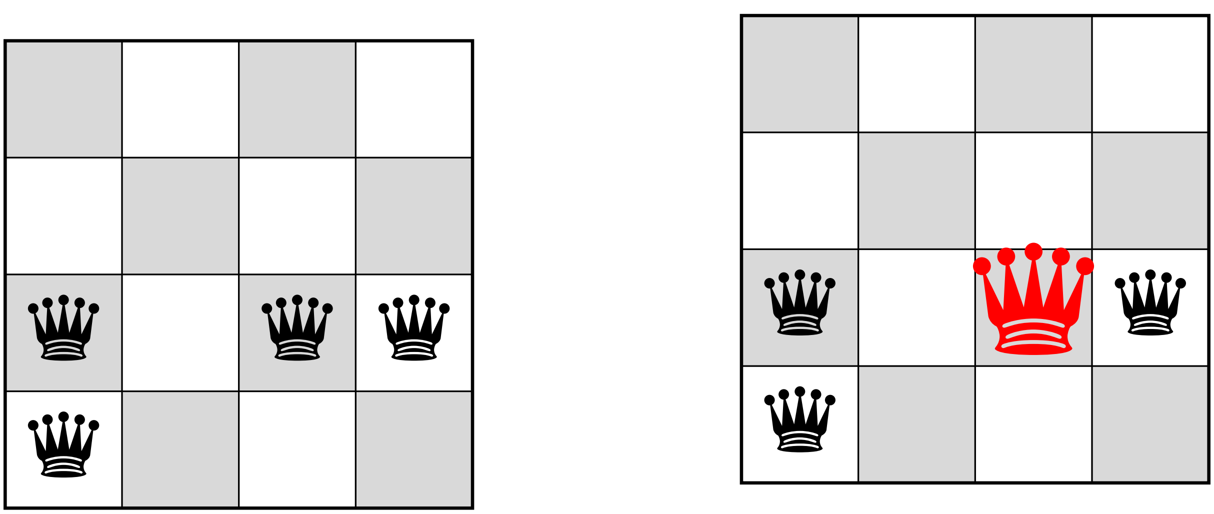

(a) Initial step on the 4×4 board.

(b) Step one. A queen is activated and independently attacks each neighboring queen.

Figure 1: Illustration of the S-SQEP board. At time step t, one queen is activated, and each of its neighboring queens is independently eliminated with probability p.

Equivalently, $\kappa(s) = \frac{1}{2}|\mathscr{E}^t|$ *since* CanKill *is symmetric. Define the stopping time* $\tau := \inf\{t \geq 0 : S^{(t)} \in S_{\text{abs}}\}$.

For a policy π*, the total cost is* $\Psi(\pi, p) := \sum_{t=0}^{\tau-1} \kappa\big(S^{(t)}\big)$*, where the absorbing set is* $S_{\text{abs}} := \{s : \kappa(s) = 0\}$*, and the value function is*

$$\text{(S-SQEP)} \qquad V^*(s) := \inf_{\pi} \mathbb{E}_{\pi}\Big[\Psi(\pi, p) \mid S^{(0)} = s\Big].$$

The *attack graph* at time t is $G^{(t)} = (Q^{(t)}, \mathscr{E}^t)$ where $\mathscr{E}^t = \{(q_i, q_j) : \text{CanKill}^{(t)}(q_i, q_j) = 1\}$. The attack graph $G^{(t)}$ has directed edges and $|\mathscr{E}^t| = 2\kappa(s)$; the cost κ counts each conflict pair once.

Remark 2. *Since queen positions never change, the position map* $\ell^{(t)}$ *is always the restriction* $\ell^{(0)}|_{Q^{(t)}}$*, so the reachable state space is* $\mathscr{S}_{\text{reach}} = \{(Q, \ell^{(0)}|_Q) : Q \subseteq Q^{(0)}\}$*, with* $|\mathscr{S}_{\text{reach}}| = 2^n$*; for* $n = 8$*, this gives* 256 *reachable states. Nevertheless, the transition distributions under* $p < 1$ *involve* $2^{|N^{(t)}(\sigma_t)|}$ *possible kill outcomes per step, and the number of (state, action) pairs scales as* $n \cdot 2^n$*, motivating policy-based rather than exhaustive value-iteration methods for moderate* n.

Remark 3. *Setting* $p = 1$ *states the deterministic SQEP.*

Lemma 4. *For every* $t' > t$*, we have* $Q^{(t')} \subseteq Q^{(t)}$.

Definition 5 (Adaptive Kill Step)**.** *Fix a probability* $p \in (0, 1]$ *and an adaptive policy* π*. At each step* t *we have the following.*

(i) ***Observation.*** *The policy observes the current board state* $S^{(t-1)} = (Q^{(t-1)}, \ell^{(t-1)})$.

(ii) ***Activation.*** *The policy selects* $\sigma_t = \pi(S^{(t-1)}) \in Q^{(t-1)}$.

(iii) ***Neighborhood computation.*** *The unblocked kill neighborhood of* σ_t *is computed:*

$$N^{(t)}(\sigma_t) \;=\; \Big\{q_j \in Q^{(t-1)} \setminus \{\sigma_t\} \;:\; \text{CanKill}^{(t)}(\sigma_t, q_j) = 1\Big\}.$$

(iv) ***Stochastic kill.*** *For each* $q_j \in N^{(t)}(\sigma_t)$*, draw* $\xi_j^{(t)} \sim \text{Bern}(p)$ *independently. The realized kill set is*

$$Y^{(t)} \;=\; \big\{q_j \in N^{(t)}(\sigma_t) : \xi_j^{(t)} = 1\big\}.$$

(v) ***State transition.*** *The board updates as* $Q^{(t)} = Q^{(t-1)} \setminus Y^{(t)}$*,* $\ell^{(t)}(q) = \ell^{(t-1)}(q)$ *for all* $q \in Q^{(t)}$.

(vi) ***Stopping check.*** *The process terminates at step* t *if* $|\mathscr{E}^{(t)}| = 0$ *or* $N^{(t)}(q) = \emptyset$ *for all* $q \in Q^{(t)}$.

Remark 6. *We call a policy* π admissible *if it never activates an idle queen when an active one exists; an idle step occurs when* $|N^{(t)}(\sigma_t)| = 0$*. Formally, if* $\exists q \in Q^{(t-1)}$ *with* $|N^{(t)}(q)| \geq 1$*, then* $|N^{(t)}(\pi(S^{(t-1)}))| \geq 1$.

2.2. Geometric interpretation of Path and *CanKill*

Definition 7. *Let* $q_i, q_j \in Q^{(t)}$ *be distinct queens whose positions are* $\ell^{(t)}(q_i) = (r_i, c_i)$*, and* $\ell^{(t)}(q_j) = (r_j, c_j)$*. Assume that* q_i *and* q_j *share a row, column, or diagonal.*

Define $\delta_{ij} := \big(\text{sgn}(r_j - r_i), \text{sgn}(c_j - c_i)\big) \in \{-1, 0, 1\}^2$ *and* $d_{ij} := \max\big\{|r_j - r_i|,\ |c_j - c_i|\big\}$*. The open segment between the two queen positions is*

$$\text{Path}(q_i, q_j) := \Big\{(r_i + k\delta_{ij,1},\, c_i + k\delta_{ij,2}) : k = 1, \ldots, d_{ij} - 1\Big\} \subseteq B.$$

Equivalently, $\text{Path}(q_i, q_j)$ *is the set of all board cells strictly between the two queen positions along their common attack line. Since surviving queens never change position, this set is time-invariant for any fixed pair of surviving queens.*

Definition 8 (Blockers)**.** *For $q_i, q_j \in Q^{(t)}$ with $i \neq j$, define the blocker set*

$$\mathrm{Block}^{(t)}(q_i,q_j) := \{q_k \in Q^{(t)} \setminus \{q_i,q_j\} : \ell^{(t)}(q_k) \in \mathrm{Path}(q_i,q_j)\}.$$

Thus $\mathrm{Block}^{(t)}(q_i,q_j)$ *contains exactly those surviving queens that lie on the open segment between q_i and q_j.*

Proposition 9. *If there exists a queen $q_k \in Q^{(t)} \setminus \{q_i,q_j\}$ such that $\ell^{(t)}(q_k) \in \mathrm{Path}(q_i,q_j)$, then $CanKill^{(t)}(q_i,q_j) = 0$.*

Proof. If such a queen q_k exists, then by Definition 8 the blocker set $\mathrm{Block}^{(t)}(q_i,q_j)$ is nonempty. Hence the indicator $\mathbf{1}\left[\mathrm{Block}^{(t)}(q_i,q_j) = \emptyset\right]$ equals 0. Therefore, $CanKill^{(t)}(q_i,q_j) = \mathbf{1}[q_i \overset{\text{atk}}{\sim} q_j] \cdot 0 = 0$. ■

Remark 10. *If desired, the time index in $\ell^{(t)}$ may be replaced by $\ell^{(0)}$ for every surviving queen, because queen positions are fixed and only the set $Q^{(t)}$ changes over time. The present notation is kept only to match the state description in the main model.*

3. Structural properties

Lemma 11 (Survivor of own round)**.** *For every $t \geq 0$, $\Lambda(t) = |Q^{(t)}| \geq 1$.*

Proof. By Definition 5(iii), $\mathscr{N}^{(t)}(\sigma_t) \subseteq Q^{(t-1)} \setminus \{\sigma_t\}$, so $\sigma_t \notin Y^{(t)}$, hence $\sigma_t \in Q^{(t)} = Q^{(t-1)} \setminus Y^{(t)}$. Since $Q^{(0)}$ has $n \geq 1$ queens, induction gives $Q^{(t)} \neq \emptyset$ for all t. ■

3.1. Absorption and termination

Theorem 12 (Almost-sure absorption)**.** *For every $p > 0$ and every admissible policy π, the S-SQEP terminates almost surely.*

Proof. Let $\Lambda(t) := |Q^{(t)}|$. At every non-absorbed step $t-1 < \tau$, admissibility forces $|N^{(t)}(\sigma_t)| \geq 1$: if it did not, all queens would have empty kill-neighborhoods, which by Definition 5(iv) means $|\mathscr{E}^{(t)}| = 0$ and the process is already absorbed, a contradiction. Since the kill set $Y^{(t)}$ is the union of $|N^{(t)}(\sigma_t)| \geq 1$ independent $\mathrm{Bern}(p)$ trials, we obtain the uniform drift bound

$$\mathbb{E}[\Lambda(t-1) - \Lambda(t) \mid \mathscr{F}_{t-1}] \;=\; \mathbb{E}\Big[|Y^{(t)}| \;\Big|\; \mathscr{F}_{t-1}\Big] \;\geq\; p \cdot \mathbf{1}[t-1 < \tau] \quad \text{a.s.}$$

Define the stopped process $X(t) := \Lambda(t \wedge \tau) + p(t \wedge \tau)$. On $\{t-1 < \tau\}$ the drift $-p$ and the counter increment $+p$ exactly cancel, so $\mathbb{E}[X(t) - X(t-1) \mid \mathscr{F}_{t-1}] \leq 0$; on $\{t-1 \geq \tau\}$ the process is frozen and the increment is 0. Hence X is a non-negative supermartingale with $X(0) = n$. ■

Corollary 13 (No persistent cycles)**.** *The S-SQEP has no cyclic trajectories almost surely.*

Proof. Since $\mathbb{E}_\pi[\tau] < \infty$, the trajectory is an a.s. finite sequence. Lemma 4 makes Λ non-increasing, so no state with more surviving queens can be revisited, precluding any persistent cycle. ■

Expected stopping time.

Proposition 14 (Stopping time bound)**.** *Under any admissible policy π and any $p \in (0,1]$, $\mathbb{E}_\pi[T] \;\leq\; \frac{n-1}{p}$.*

Proof. Let $\Lambda(t) = |Q^{(t)}|$ and $\Delta(t) := \Lambda(t-1) - \Lambda(t) \geq 0$ for $t \geq 1$.

Fix $t \geq 1$ and condition on $\mathscr{F}_{t-1}$. If $t-1 < \tau$ (not yet absorbed), then $\kappa(S^{(t-1)}) > 0$, so some queen has nonempty kill-neighborhood; by admissibility, $m := |N^{(t)}(\sigma_t)| \geq 1$ as well. Conditionally on $\mathscr{F}_{t-1}$, $\Delta(t) = |Y^{(t)}| = \sum_{j \in N^{(t)}(\sigma_t)} \xi^{(t)} j$ is a sum of m i.i.d. Bernoulli(p) trials, so $\mathbb{E}[\Delta(t) \mid \mathscr{F} t-1] = mp \geq p$ on $t-1 < \tau$. If $t-1 \geq \tau$ (already absorbed), the state is frozen, so $\Delta(t) = 0$. Since $\mathbf{1}(t \leq \tau)$ is $\mathscr{F}_{t-1}$-measurable, both cases combine into the single inequality based on admissibility $\mathbb{E}[\Delta(t) \mid \mathscr{F}_{t-1}] \geq p \cdot \mathbf{1}(t \leq \tau)$ a.s.

For $T \geq 0$ define $S_T := \sum_{t=1}^{T} \big(\Delta(t) - p, \mathbf{1}(t \leq \tau)\big)$. Then, $\mathbb{E}[S_T - S_{T-1} \mid \mathscr{F}_{T-1}] \geq 0$, so $(S_T)T \geq 0$ is a submartingale with $S_0 = 0$, hence $\mathbb{E}[S_T] \geq 0$ for all T, i.e. $\mathbb{E}\Big[\sum t = 1^T \Delta(t)\Big] \geq p, \mathbb{E}[\min(T,\tau)]$.

Since $\Delta(t) = 0$ for $t > \tau$, $\sum_{t=1}^{T} \Delta(t) = \Lambda(0) - \Lambda(T \wedge \tau)$. By Lemma 11, $\Lambda(T \wedge \tau) \geq 1$, so $\sum_{t=1}^{T} \Delta(t) \leq \Lambda(0) - 1 = n-1$. Then we have $p, \mathbb{E}[\min(T, \tau)] \leq n-1$ for every T.

By Theorem 12, $\tau < \infty$ a.s., so $\min(T, \tau) \uparrow \tau$ as $T \to \infty$; by monotone convergence $\mathbb{E}[\min(T, \tau)] \to \mathbb{E}[\tau]$. Hence $\mathbb{E}_\pi[T] = \mathbb{E}[\tau] \leq \frac{n-1}{p}$. ■

Remark 15. *For $p = 1$, the bound reduces to $\mathbb{E}[T] \leq (n-1)$. As $p \to 0^+$, $\mathbb{E}[T] \to \infty$.*

3.2. Bellman optimality and policy structure

Lemma 16. *The MDP $\mathscr{M}$ of Definition 1, restricted to the finite reachable state space $\mathscr{S}_{\text{reach}}$ of Remark 2, satisfies the stochastic-shortest-path regularity conditions of Bertsekas and Tsitsiklis [2]:*

- *(A1)* ***(Existence of a proper policy.)*** *Every admissible stationary policy π (Remark 6) is* proper *and it reaches S_{abs} from every state $s \in \mathscr{S}_{\text{reach}}$ almost surely, with finite expected hitting time, $\mathbb{E}_\pi[\tau \mid S^{(0)} = s] \leq (n-1)/p < \infty$.*
- *(A2)* ***(Improper policies cost infinity.)*** *Every stationary policy π that is* not *proper (i.e. for which some $s \notin S_{\text{abs}}$ has $Pr_\pi(\tau = \infty \mid S^{(0)} = s) > 0$) satisfies $\mathbb{E}_\pi[\Psi(\pi, p) \mid S^{(0)} = s] = \infty$.*

Proof. (A1) is exactly Theorem 12 together with Proposition 14: every admissible π satisfies $Pr_\pi(\tau < \infty) = 1$ and $\mathbb{E}_\pi[\tau] \leq (n-1)/p$ from every state, which is properness with a finite expected hitting time.

For (A2), fix a non-proper π and $s \notin S_{\text{abs}}$ with $\delta := Pr_\pi(\tau = \infty \mid S^{(0)} = s) > 0$. Since κ is integer-valued and $\kappa(s') > 0$ for every $s' \notin S_{\text{abs}}$, we have $\kappa(s') \geq 1$ there. On the event $\{\tau = \infty\}$ the trajectory never enters S_{abs}, so pathwise $\Psi(\pi, p) = \sum_{t=0}^{\infty} \kappa(S^{(t)}) \geq \sum_{t=0}^{\infty} 1 = \infty$. Hence $\mathbb{E}_\pi[\Psi(\pi, p) \mid S^{(0)} = s] \geq \delta \cdot \infty = \infty$. ■

Theorem 17 (Bellman optimality)**.** *Under the regularity conditions of Lemma 16, the following hold.*

- *(i) V^* is finite everywhere on $\mathscr{S}_{\text{reach}}$ and is the* unique *function satisfying $V^*(s) = 0$ for $s \in S_{\text{abs}}$ and*

$$V^*(s) = \kappa(s) + \min_{a \in \mathscr{A}(s)} \mathbb{E}\left[V^*(S') \mid S = s, \sigma = a\right], \qquad s \notin S_{\text{abs}}.$$

- *(ii) There exists a stationary deterministic policy π^* attaining the minimum at every s; π^* is proper, and $V^{\pi^*} = V^*$.*
- *(iii) Every minimizing action $a^*(s)$ satisfies $N(s, a^*(s)) \neq \varnothing$; in particular π^* never activates an idle queen when a non-idle one exists, so π^* is admissible* without this being imposed *on the optimization.*
- *(iv) The infimum in Definition 1, taken over all stationary (possibly randomized, possibly non-admissible) policies, is attained by π^*; Restricting the competing class to admissible policies does not change V^*.*

Proof. By Lemma 16, Assumptions (A1)–(A2) of [2, Props. 1–3] hold on the finite state space $\mathscr{S}_{\text{reach}}$, with finite action sets $\mathscr{A}(s) = Q$ and one-stage costs $\kappa(s) \geq 0$ ($\kappa(s) \geq 1$ off S_{abs}). Parts (i)–(ii) are immediate specializations of [2, Prop. 2]: the optimal cost is finite and is the unique solution of the Bellman equation among functions finite on $\mathscr{S}_{\text{reach}}$, and (since $\mathscr{S}_{\text{reach}}$ is finite) direct minimization yields an optimal proper stationary deterministic policy.

For (iii), define $Q(s, a) := \kappa(s) + \mathbb{E}[V^*(S') \mid S = s, \sigma = a]$. For an idle action a_{idle} with $N(s, a_{\text{idle}}) = \varnothing$, the transition is degenerate ($Y = \varnothing$, $S' = s$ a.s.), giving $Q(s, a_{\text{idle}}) = \kappa(s) + V^*(s) > V^*(s)$ strictly whenever $s \notin S_{\text{abs}}$ (since $\kappa(s) \geq 1$ there). Since $\kappa(s) \geq 1$ implies some queen has a non-empty neighborhood, a non-idle action is always available and always strictly preferred, so a_{idle} is never a minimizer.

For (iv), by [2, Prop. 3], under (A1)–(A2) the infimum of $\mathbb{E}_\pi[\Psi(\pi, p)]$ over the full class of history-dependent randomized policies coincides with the infimum over stationary deterministic policies and equals V^*. Combined with (iii), the optimum is already attained within the class of admissible policies, so restricting Definition 1 to admissible π is without loss of generality. ■

Proposition 18 (Greedy policy suboptimality)**.** *The greedy policy $\pi_G(s) = \arg\max_{a \in \mathscr{A}(s)} |N^{(t)}(a)|$ (activating the queen with the most live attack targets) is not universally optimal for $p < 1$.*

Proof. By counterexample. Consider four queens at $A = (1,1), B = (1,2), C = (3,2), D = (4,2)$, with kill probability $p \in (0,1)$. The initial attack graph is a path $A - B - C - D$ with B, C each of degree 2 (tied for maximum) and A, D of degree 1; $B - D$ is blocked by C. Any tie-breaking rule for π_G may select C. By calculation we have $V^* = \frac{3}{p}$ against the optimal $V^{\pi_G} = \frac{7-3p}{p(2-p)}$, attained by activating C instead, a strictly positive gap $\frac{5-3p}{p(2-p)}$ for every $p \in (0,1)$. Hence π_G is not universally optimal. ■

Remark 19. *Proposition 18 highlights that the policy structure for minimizing cumulative conflicts is nontrivial. Balancing trade-offs between immediate and future conflicts is important when $p < 1$. One should design policies taking into account the entire trajectory cost $\Psi(\pi, p)$ rather than just the final survivor count.*

3.3. Stochastic dominance and monotonicity

Let $\Psi(\pi, p)$ and $\Psi(\pi, p')$ denote the terminal conflict counts under the same policy π but different kill probabilities $p \leq p'$.

Since $\ell^{(t)}(q) = \ell^{(0)}(q)$ for every surviving queen $q \in Q^{(t)}$, the *attack relation* $q_i \overset{atk}{\sim} q_j$ is a time-invariant property of the pair (q_i, q_j). Define the static *conflict skeleton*

$$\mathscr{S} \;:=\; \big\{(q_i, q_j) \in Q^{(0)} \times Q^{(0)} : q_i \overset{atk}{\sim} q_j\big\}.$$

For $Q \subseteq Q^{(0)}$, let

$$L(Q) \;:=\; \mathscr{S} \cap (Q \times Q), \qquad A(Q) \;:=\; \big\{(q_i, q_j) \in L(Q) : \mathrm{CanKill}_Q(q_i, q_j) = 1\big\}, \qquad B(Q) \;:=\; L(Q) \setminus A(Q),$$

so that $L(Q) = A(Q) \sqcup B(Q)$, and write $L^{(t)} := L(Q^{(t)})$, $A^{(t)} := A(Q^{(t)})$, $B^{(t)} := B(Q^{(t)})$. $B^{(t)}$ is the set of *currently shielded* conflicts: pairs (q_i, q_j) that are still mutually attacking and both alive, but whose line of attack is momentarily occupied by a third surviving queen.

Lemma 20 (Line decomposition of κ)**.** *Let $\mathscr{L}$ be the set of maximal lines of B (rows, columns, and diagonals of both orientations). For $Q \subseteq Q^{(0)}$ and $\ell \in \mathscr{L}$, let $m_\ell(Q) := |\{q \in Q : \ell^{(0)}(q) \in \ell\}|$. Then*

$$\kappa\big(Q, \ell^{(0)}|_Q\big) \;=\; \sum_{\ell \in \mathscr{L}} \max\big(m_\ell(Q) - 1, 0\big).$$

Consequently $Q' \subseteq Q \subseteq Q^{(0)}$ implies $\kappa\big(Q', \ell^{(0)}|_{Q'}\big) \leq \kappa\big(Q, \ell^{(0)}|_Q\big)$.

Proof. Fix ℓ and order the queens of Q on ℓ by position: $q_{(1)}, \ldots, q_{(m)}$, $m = m_\ell(Q)$. For $i < j$, $\mathrm{Path}(q_{(i)}, q_{(j)})$ lies entirely on ℓ (it is the open segment between two colinear points), so $\mathrm{Block}(q_{(i)}, q_{(j)}) = \{q \in Q : \ell^{(0)}(q) \in \mathrm{Path}(q_{(i)}, q_{(j)})\} = \{q_{(i+1)}, \ldots, q_{(j-1)}\}$. This is empty iff $j = i + 1$. So exactly the $m - 1$ consecutive pairs on ℓ are unblocked, contributing $\max(m - 1, 0)$ to κ. Every mutually-attacking pair lies on exactly one $\ell \in \mathscr{L}$ two distinct cells cannot share two of {row, column, "/"-diagonal, "∸diagonal} without coinciding so summing over ℓ double-counts nothing and misses nothing. Monotonicity follows because $Q' \subseteq Q \Rightarrow m_\ell(Q') \leq m_\ell(Q)$ for every ℓ, and $x \mapsto \max(x - 1, 0)$ is non-decreasing on $\mathbb{Z}_{\geq 0}$. ■

Lemma 21 (Monotonicity of L and B)**.** *For any $Q' \subseteq Q \subseteq Q^{(0)}$, $L(Q') \subseteq L(Q)$ and $B(Q') \subseteq B(Q)$.*

Proof. $L(Q') = \mathscr{S} \cap (Q' \times Q') \subseteq \mathscr{S} \cap (Q \times Q) = L(Q)$ is immediate from $Q' \times Q' \subseteq Q \times Q$.

For the second inclusion, fix $(q_i, q_j) \in B(Q')$, i.e. $(q_i, q_j) \in L(Q') \subseteq L(Q)$ and $\mathrm{CanKill}_{Q'}(q_i, q_j) = 0$, i.e. there exists $q_k \in Q' \setminus \{q_i, q_j\}$ with $\ell^{(0)}(q_k) \in \mathrm{Path}(q_i, q_j)$. Since $Q' \subseteq Q$, the same $q_k \in Q \setminus \{q_i, q_j\}$ witnesses $\mathrm{CanKill}_Q(q_i, q_j) = 0$. Hence $(q_i, q_j) \in L(Q)$ and $\mathrm{CanKill}_Q(q_i, q_j) = 0$, i.e. $(q_i, q_j) \in B(Q)$. ■

Remark 22 (Within-trajectory consequence)**.** *By Lemma 4, $Q^{(t)} \subseteq Q^{(t-1)}$ for every t, so Lemma 21 gives $L^{(t)} \subseteq L^{(t-1)}$ and $B^{(t)} \subseteq B^{(t-1)}$ along* any single *realization of the process: once a conflict is removed it stays removed, and once a shielded conflict becomes unblocked (or its endpoints are eliminated) it never becomes shielded again. In particular $B^{(t)} \subseteq B^{(0)}$ for all t.*

Lemma 23 (Monotone coupling in p)**.** *Fix π and $0 < p \leq p' \leq 1$. There exists a probability space $(\Omega, \mathscr{F}, Pr)$ carrying an i.i.d. family $\{U_q^{(t)}\}_{t \geq 1, q \in Q^{(0)}}$ of $\mathrm{Unif}[0,1]$ random variables such that, setting $\xi_q^{(t)}(r) := 1[U_q^{(t)} \leq r]$ for $r \in \{p, p'\}$:*

- *(i) for each $r \in \{p, p'\}$, running Definition 5 with kill indicators $\big(\xi_q^{(t)}(r)\big)_{q \in N^{(t)}(\sigma_t)}$ in place of $\xi_j^{(t)} \sim \mathrm{Bern}(r)$ reproduces the law of the S-SQEP with parameter r under policy π;*
- *(ii) $\xi_q^{(t)}(p) \leq \xi_q^{(t)}(p')$ for every $\omega \in \Omega$, $t \geq 1$, $q \in Q^{(0)}$.*

Proof. Take $\xi_q^{(t)}(r) := 1[U_q^{(t)} \le r]$. Marginally $\xi_q^{(t)}(r) \sim \mathrm{Bern}(r)$ and the family is independent across (t,q), which is exactly the law required by step (iv) of Definition 5 (the realized kill set $Y^{(t)} = \{q_j \in N^{(t)}(\sigma_t) : \xi_j^{(t)}(r) = 1\}$), proving (i). Property (ii) is immediate because $U_q^{(t)} \le p \Rightarrow U_q^{(t)} \le p'$ whenever $p \le p'$, for every ω. ■

Lemma 24 (Pathwise monotonicity of $Q^{(t)}$)**.** *Under the coupling of Lemma 23, write $Q^{(t)}(p)$, $Q^{(t)}(p')$ for the two coupled trajectories of surviving queens (each evolving by Definition 5 with its own parameter, the same policy π, and the kill indicators $\xi_\cdot^{(t)}(p)$, $\xi_\cdot^{(t)}(p')$ respectively, the latter dominating the former pointwise by Lemma 23(ii)). Then*

$$Q^{(t)}(p) \supseteq Q^{(t)}(p') \qquad \text{a.s., for every } t \ge 0.$$

Proof. By induction on t. The base case $t=0$ is an equality, $Q^{(0)}(p) = Q^{(0)}(p') = Q^{(0)}$.

Assume $Q^{(t-1)}(p) \supseteq Q^{(t-1)}(p')$. Write $Q' := Q^{(t-1)}(p') \subseteq Q := Q^{(t-1)}(p)$, with the common position map $\ell^{(0)}|_Q$. Let $\sigma_t := \pi\big(Q^{(t-1)}(p), \ell^{(0)}\big)$ be the activation chosen by π on the (larger) p-board; by admissibility (Remark 6) $N_Q(\sigma_t) \neq \varnothing$ whenever $A(Q) \neq \varnothing$.

Case 1: $\sigma_t \in Q'$. By Lemma 21 applied with $Q' \subseteq Q$, every $(q_i,q_j) \in A(Q)$ with $q_i, q_j \in Q'$ satisfies $(q_i,q_j) \in A(Q')$ as well (a pair unblocked in the larger board is unblocked in the smaller board, since the smaller board has strictly fewer potential blockers, $\mathrm{Path} \cap Q' \subseteq \mathrm{Path} \cap Q$). Hence $N_Q(\sigma_t) \cap Q' \subseteq N_{Q'}(\sigma_t)$. Using the same activation σ_t on the p'-board, the realized kill sets are

$$Y^{(t)}(p) = \big\{q \in N_Q(\sigma_t) : \xi_q^{(t)}(p) = 1\big\}, \qquad Y^{(t)}(p') = \big\{q \in N_{Q'}(\sigma_t) : \xi_q^{(t)}(p') = 1\big\}.$$

For $q \in Y^{(t)}(p) \cap Q'$ we have $q \in N_Q(\sigma_t) \cap Q' \subseteq N_{Q'}(\sigma_t)$ and $\xi_q^{(t)}(p) = 1 \Rightarrow \xi_q^{(t)}(p') = 1$ by Lemma 23(ii), so $q \in Y^{(t)}(p')$. Thus $Y^{(t)}(p) \cap Q' \subseteq Y^{(t)}(p')$, and

$$Q^{(t)}(p) = Q \setminus Y^{(t)}(p) \supseteq Q' \setminus \big(Y^{(t)}(p) \cap Q'\big) \supseteq Q' \setminus Y^{(t)}(p') = Q^{(t)}(p').$$

Case 2: $\sigma_t \notin Q'$. Then $\sigma_t \in Q \setminus Q'$ already, so σ_t takes no part in the p'-board's transition at all, and trivially $Q^{(t)}(p') = Q' \setminus Y^{(t)}(p') \subseteq Q' \subseteq Q \setminus Y^{(t)}(p) = Q^{(t)}(p)$, where the last step uses $\sigma_t \notin Q'$ so $Y^{(t)}(p) \cap Q' = \varnothing$.

In both cases $Q^{(t)}(p) \supseteq Q^{(t)}(p')$, completing the induction. ■

Proposition 25 (Monotonicity in p)**.** *Fix a policy π and $0 < p \le p' \le 1$. Under the coupling of Lemma 23,*

$$\kappa\big(S^{(t)}(p)\big) \ge \kappa\big(S^{(t)}(p')\big) \quad \text{for every } t \ge 0, \text{ a.s.},$$

hence

$$\Psi(\pi,p) \ \ge\ \Psi(\pi,p') \quad \text{a.s.}, \qquad \text{and so} \qquad \mathbb{E}_\pi[\Psi(\pi,p)] \ \ge\ \mathbb{E}_\pi[\Psi(\pi,p')].$$

In particular $p \mapsto \mathbb{E}_\pi[\Psi(\pi,p)]$ is non-increasing on $(0,1]$ and this holds pathwise, not merely in expectation.

Proof. By Lemma 20, $\kappa(Q,\cdot) = \sum_{\ell \in \mathscr{L}} \max(m_\ell(Q) - 1, 0)$ depends on Q only through the per-line counts $m_\ell(Q)$, and is non-decreasing in Q since $\max(\cdot - 1, 0)$ is non-decreasing. By Lemma 24, $Q^{(t)}(p) \supseteq Q^{(t)}(p')$ a.s. for every t under the coupling of Lemma 23, so $\kappa(S^{(t)}(p)) \ge \kappa(S^{(t)}(p'))$ a.s. for every t. Both trajectories freeze at $\kappa = 0$ past their (a.s. finite, Theorem 12) stopping times, so summing over t gives $\Psi(\pi,p) \ge \Psi(\pi,p')$ a.s., and taking expectations gives the claim. ■

Theorem 26 (Terminal monotonicity)**.** *For $0 < p \le p' < 1$ and any admissible policy π, under the coupling of Lemma 23,*

$$\big|L\big(Q^{(\tau(p))}(p)\big)\big| \ \ge\ \big|L\big(Q^{(\tau(p'))}(p')\big)\big| \quad \text{a.s.}$$

That is, the number of mutually attacking pairs surviving to termination is pathwise larger under the weaker kill probability.

Proof. By Lemma 24, $Q^{(t)}(p) \supseteq Q^{(t)}(p')$ for every t. Freeze both trajectories after their respective stopping times and let $t^* := \max\{\tau(p), \tau(p')\}$. At $t = t^*$, Lemma 21 applied to $Q^{(t^*)}(p') \subseteq Q^{(t^*)}(p)$ gives $L(Q^{(t^*)}(p')) \subseteq L(Q^{(t^*)}(p))$. Since the frozen states satisfy $L^{(\tau(\cdot))}(\cdot) = L(Q^{(t^*)}(\cdot))$, the pathwise inclusion and hence inequality of cardinalities follow. Stochastic dominance of the laws and the expectation inequality are immediate consequences. ■

Lemma 27 (Cumulative monotonicity)**.** *For $0 < p \le p' < 1$ and any admissible policy π, $\mathbb{E}_\pi[\Psi(\pi,p)] \ge \mathbb{E}_\pi[\Psi(\pi,p')]$.*

Proof. By Proposition 48, $\mathbb{E}_\pi[\Psi(\pi,p)] = \sum_{\{i,j\}}\sum_t Pr_\pi((q_i,q_j) \in A(Q^{(t)}))$. By Proposition 25, $p \mapsto \mathbb{E}_\pi[\Psi(\pi,p)]$ is non-increasing. Therefore $p \le p'$ implies $\mathbb{E}_\pi[\Psi(\pi,p)] \ge \mathbb{E}_\pi[\Psi(\pi,p')]$. ■

Remark 28. *Lemma 21 also yields, by the same substitution $Q' := Q^{(t)}(p') \subseteq Q := Q^{(t)}(p)$,*

$$B^{(t)}(p') \subseteq B^{(t)}(p) \qquad \text{a.s., for every } t,$$

i.e. at every common time the smaller-p trajectory carries at least as many shielded *conflicts as the larger-p one extra surviving queens can only add blockers, never remove them. Writing $|A^{(t)}| = |L^{(t)}| - |B^{(t)}|$, the dominance $\Psi(\pi,p) \ge \Psi(\pi,p')$ established above transfers to $\left|A^{(T(\pi,p))}(p)\right| \ge \left|A^{(T(\pi,p'))}(p')\right|$ whenever $\left|L^{(T(\pi,p))}(p)\right| - \left|L^{(T(\pi,p'))}(p')\right| \ge \left|B^{(T(\pi,p))}(p)\right| - \left|B^{(T(\pi,p'))}(p')\right|$, which holds in particular whenever $B^{(T(\pi,p))}(p) = \varnothing$, i.e. at termination no surviving queen lies on the attack line between two other surviving, mutually attacking queens.*

4. Adaptive Coverage Greedy for $p < 1$

4.1. Greedy scoring

For a state $s = (Q,\ell)$ and an action $a \in Q$, define the expected one-step conflict reduction

$$G(s,a) := \mathbb{E}\left[\kappa(s) - \kappa(s \setminus Y^{(t)}) \,\middle|\, S^{(t)} = s,\, \sigma_t = a\right],$$

where $Y^{(t)} \subseteq N^{(t)}(a)$ is the realized kill set and $s \setminus Y^{(t)}$ denotes the state $(Q \setminus Y^{(t)}, \ell)$. Since the kills are mutually independent Bernoulli(p) trials,

$$G(s,a) = \sum_{Y \subseteq N^{(t)}(a)} \left[\kappa(s) - \kappa(s \setminus Y)\right] p^{|Y|}(1-p)^{|N^{(t)}(a)|-|Y|}. \tag{1}$$

Definition 29 (ACG policy)**.** *At each step t, the Adaptive Coverage Greedy (ACG) policy activates*

$$\sigma_t = \pi_{\text{ACG}}(S^{(t)}) := \arg\max_{a \in Q^{(t)}} G(S^{(t)},a),$$

if any queen has a non-empty kill neighborhood, $G(S^{(t)},a) \ge p \cdot \kappa_{\min} > 0$ for that queen, so it is always selected over an idle queen for which $G = 0$.

Remark 30 (Computational cost)**.** *For each action a, Equation 1 sums over $2^{|N^{(t)}(a)|}$ subsets. For a board of size n, this is $O(n \cdot 2^{n-1})$ in the worst case. For small n (e.g. $n \le 8$) this is entirely practical; for larger boards, M i.i.d. samples from the kill distribution give a Monte Carlo estimate with error $O(\kappa(s)/\sqrt{M})$ by Hoeffding's inequality.*

4.2. Exact formula for $G(s,a)$

We now derive a closed-form expression for $G(s,a)$ that avoids explicit subset enumeration.

Lemma 31 (Closed-form gain)**.** *For any state s and action $a \in Q^{(t)}$,*

$$G(s,a) = \sum_{\{i,j\}:\text{CanKill}^{(t)}(q_i,q_j)=1} \Pr(\{q_i,q_j\} \text{ is resolved by activating } a), \tag{2}$$

where a conflict $\{q_i,q_j\}$ is resolved by activating a if and only if at least one of q_i,q_j belongs to $N^{(t)}(a)$ and the corresponding kill trial succeeds. Writing $a_i := \mathbf{1}[q_i \in N^{(t)}(a)]$ and $a_j := \mathbf{1}[q_j \in N^{(t)}(a)]$, the events $\{q_i \in Y^{(t)}\}$ and $\{q_j \in Y^{(t)}\}$ are independent with $\Pr(q_i \in Y^{(t)}) = p\,a_i$ and $\Pr(q_j \in Y^{(t)}) = p\,a_j$, so inclusion–exclusion for independent events gives, for each conflict pair $\{q_i,q_j\}$,

$$\Pr(\textit{resolved}) = 1 - (1-p\,a_i)(1-p\,a_j) = 1-(1-p)^{a_i+a_j}, \tag{3}$$

which recovers, case by case,

$$\Pr(\textit{resolved}) = \begin{cases} 1-(1-p)^2 & \textit{if } q_i, q_j \in N^{(t)}(a), \\ p & \textit{if exactly one of } q_i, q_j \in N^{(t)}(a), \\ 0 & \textit{if } q_i, q_j \notin N^{(t)}(a). \end{cases} \tag{4}$$

Proof. A conflict pair $\{q_i, q_j\}$ disappears from $\kappa(s \setminus Y^{(t)})$ if and only if $q_i \in Y^{(t)}$ or $q_j \in Y^{(t)}$ (or both). By Definition 5(iv), $q_k \in Y^{(t)}$ iff $q_k \in N^{(t)}(a)$ and $\xi_k^{(t)} = 1$; in particular $\Pr(q_k \in Y^{(t)}) = p\,\mathbf{1}[q_k \in N^{(t)}(a)]$, since a queen outside $N^{(t)}(a)$ is never targeted. The trials $\xi_i^{(t)}, \xi_j^{(t)}$ are mutually independent for $i \neq j$, so the events $\{q_i \in Y^{(t)}\}$ and $\{q_j \in Y^{(t)}\}$ are independent, and $\Pr(q_i \in Y^{(t)} \text{ or } q_j \in Y^{(t)}) = 1-(1-pa_i)(1-pa_j)$, giving Equation 3 and hence Equation 4. Summing over all active conflict pairs and taking expectations (which commutes with the finite sum) gives Equation 2. ∎

4.3. The key gain lower bound

The approximation guarantee rests on one structural inequality relating $\max_a G(s,a)$ to $\kappa(s)$.

Lemma 32 (Gain lower bound)**.** *For every non-absorbing state $s \notin S_{\text{abs}}$ and every $p \in (0,1]$,*

$$\max_{a \in Q^{(t)}} G(s,a) \geq \frac{2p}{n}\,\kappa(s).$$

Proof. Consider the directed attack graph $G^{(t)} = (Q^{(t)}, E^{(t)})$ with $|E^{(t)}| = 2\kappa(s)$ directed edges (each unordered pair counted in both directions). For each queen $q_k \in Q^{(t)}$, let $d_k := |N^{(t)}(q_k)|$ denote its out-degree in $G^{(t)}$. Summing over all queens $\sum_{k:\,q_k \in Q^{(t)}} d_k = |E^{(t)}| = 2\kappa(s)$.

Since $|Q^{(t)}| \leq n$, by the pigeonhole principle there exists $a^* \in Q^{(t)}$ with $d_{a^*} = |N^{(t)}(a^*)| \geq \frac{2\kappa(s)}{n}$.

Now, by Lemma 31, for any action a with $|N^{(t)}(a)| = m$, each neighbor $q_j \in N^{(t)}(a)$ contributes at least the single-kill probability p to $G(s,a)$ (its conflict with a, and also with any other surviving queen it attacks, is resolved with probability $\geq p$). To get a clean lower bound we use only the conflicts between a and each $q_j \in N^{(t)}(a)$, which, by definition, each satisfy $\text{CanKill}^{(t)}(a, q_j) = 1$ and contribute exactly p to $G(s,a)$. Hence $G(s,a^*) \geq p \cdot |N^{(t)}(a^*)| \geq \frac{2p}{n}\,\kappa(s)$, and $\max_{a \in Q^{(t)}} G(s,a) \geq G(s,a^*)$ provides the proof. ∎

Remark 33. *This bound is tight in the star-graph case; if one queen a^* attacks all $n-1$ others and no other attack edges exist, then $\kappa(s) = n-1$, $d_{a^*} = n-1$, all other $d_k = 1$, and $G(s,a^*) = p(n-1) = 2p\kappa(s)/n \cdot n/2$. In general the factor $2p/n$ cannot be improved without additional structural assumptions.*

4.4. Approximation guarantee

Theorem 34 (ACG approximation ratio)**.** *For every initial state $s_0 \notin S_{\text{abs}}$ and every $p \in (0,1]$,*

$$\mathbb{E}^{\pi_{\text{ACG}}}[\Psi(\pi_{\text{ACG}}, p)] \leq \frac{n}{2p}\,V^*(s_0).$$

Proof. We proceed in three steps.

Step 1: Geometric decay of expected conflict count. Let $\gamma := 2p/n$. For any non-absorbing state $S^{(t)} = s$, since π_{ACG} maximizes $G(s,a)$, Lemma 32 gives $G(S^{(t)}, \pi_{\text{ACG}}(S^{(t)})) \geq \gamma\,\kappa(S^{(t)})$.

Taking conditional expectations:

$$\mathbb{E}\Big[\kappa(S^{(t+1)}) \,\Big|\, S^{(t)}\Big] = \kappa(S^{(t)}) - G(S^{(t)}, \pi_{\text{ACG}}(S^{(t)})) \leq (1-\gamma)\,\kappa(S^{(t)}).$$

Iterating and taking full expectations $\mathbb{E}\Big[\kappa(S^{(t)})\Big] \leq (1-\gamma)^t\,\kappa(s_0)$. To verify Corollary 35, by Theorem 12 the process terminates a.s., so on the event $\{t \geq \tau\}$ we have $\kappa(S^{(t)}) = 0$ and the inequality holds trivially. On $\{t < \tau\}$, $S^{(t)} \notin S_{\text{abs}}$ and the bound follows from Lemma 32.

Step 2: Summing over the trajectory. By the definition of Ψ and Theorem 12 (a.s. termination):

$$\mathbb{E}^{\pi_{\mathrm{ACG}}}[\Psi(\pi_{\mathrm{ACG}},p)] = \sum_{t=0}^{\infty} \mathbb{E}\Big[\kappa(S^{(t)})\,\mathbf{1}[t<\tau]\Big] \le \sum_{t=0}^{\infty} \mathbb{E}\Big[\kappa(S^{(t)})\Big] \le \sum_{t=0}^{\infty}(1-\gamma)^t\,\kappa(s_0) = \frac{\kappa(s_0)}{\gamma} = \frac{n\,\kappa(s_0)}{2p}.$$

The interchange of sum and expectation is justified by Tonelli's theorem since all terms are non-negative.

Step 3: Relating $\kappa(s_0)$ to $V^*(s_0)$. Since the stage cost $\kappa(S^{(0)}) = \kappa(s_0)$ is incurred at $t=0$ under every admissible policy, we have $V^*(s_0) \ge \kappa(s_0)$ (this is immediate from the Bellman equation of Theorem 17: $V^*(s_0) = \kappa(s_0) + \min_a \mathbb{E}[V^*(S')] \ge \kappa(s_0)$). Therefore proof is completed. ■

Corollary 35 (Explicit bound)**.** *Since $\kappa(s_0) \le \binom{n}{2}$, $\mathbb{E}^{\pi_{\mathrm{ACG}}}[\Psi(\pi_{\mathrm{ACG}},p)] \le \frac{n^2(n-1)}{4p}$.*

Remark 36 (Tightness of the ratio)**.** *The ratio $n/(2p)$ is tight in the following sense. For the star instance of Remark 33, $V^*(s_0) = \kappa(s_0) = n-1$ (the star center resolves all conflicts in one step with probability 1 under $p=1$, but contributes $(n-1)$ to the stage cost at $t=0$). The ACG policy correctly activates the center, giving $\mathbb{E}[\Psi] = (n-1)/1 = n-1$ for $p=1$, so the ratio equals $1 < n/(2p) = n/2$. The ratio $n/(2p)$ is approached when the conflict graph is a perfect matching: $\kappa(s_0) = n/2$, each queen has exactly one neighbor, every $G(s,a) = p$, and the expected cost is $\sum_t (1-p)^{\lfloor t/n \rfloor}(n/2) \approx n/(2p)\cdot V^*(s_0)$ for $p \to 0^+$.*

4.5. Relation to the Bellman equation and optimality gap

Proposition 37 (Suboptimality characterization)**.** *Define the one-step optimality gap of ACG at state s as*

$$\Delta(s) := G(s,\pi_{\mathrm{ACG}}(s)) - G(s,\pi^*(s)),$$

where $\pi^*(s)$ is any Bellman-optimal action. Then $\Delta(s) \ge 0$ (ACG maximizes G) but the cumulative suboptimality satisfies

$$\mathbb{E}^{\pi_{\mathrm{ACG}}}[\Psi(\pi_{\mathrm{ACG}},p)] - V^*(s_0) \le \left(\frac{n}{2p} - 1\right) V^*(s_0).$$

Proof. The upper bound follows directly from Theorem 34. The term $-V^*(s_0)$ on the left follows from the trivial lower bound $\mathbb{E}^{\pi_{\mathrm{ACG}}}[\Psi] \ge V^*(s_0)$ (since V^* is the infimum over all policies). ■

Remark 38. *The policy π_G defined in Proposition 18 activates $\arg\max_a |N^{(t)}(a)|$, i.e. it maximizes degree in $G^{(t)}$. The ACG policy instead maximizes $G(s,a)$, the expected conflict reduction, which equals $p \cdot |N^{(t)}(a)|$ only when conflicts among the neighbors of a can be neglected. In general, because a pair $\{q_i,q_j\}$ with both $q_i,q_j \in N^{(t)}(a)$ contributes $1-(1-p)^2 = 2p-p^2 > p$ to $G(s,a)$, ACG correctly discounts the double-counting that makes π_G suboptimal.*

5. A mixed-integer programming formulation for $p = 1$

By Remark 2, $p=1$ collapses S-SQEP to the *deterministic* SQEP; $Y^{(t)} = N^{(t)}(\sigma_t)$ with probability 1, so an activation policy reduces to a sequence $\sigma = (\sigma_1,\dots,\sigma_H) \in (Q^{(0)})^H$, $H := n-1$ (an upper bound on the number of effective rounds by Lemma 4). Precompute, from $\ell^{(0)}$ alone, the static data are $a_{ij} := 1\big[q_i \overset{atk}{\sim} q_j\big]$, and $b_{ijk} := 1\big[\ell^{(0)}(q_k) \in \mathrm{Path}(q_i,q_j)\big]$, for all $i \ne j$, $k \notin \{i,j\}$.

Additionally, decision variables for $i,j \in [n]$, $t \in \{0,\dots,H\}$) are $x_i^t \in \{0,1\}:\ q_i \in Q^{(t)}$; and $y_i^t \in \{0,1\},\ t \ge 1:\ \sigma_t = q_i$; and $z_{ij}^t \in \{0,1\},\ t \ge 1:\ \mathrm{CanKill}^{(t)}(q_i,q_j) = 1$ on the board $Q^{(t-1)}$; Also, $u_{ij}^t \in \{0,1\},\ t \ge 1:\ y_i^t = z_{ij}^t = 1$, and $\beta_i^t \in \{0,1\}$.

The variables β_i^t and v_i^t serve exclusively to enforce admissibility (Remark 6) within the MIP: β_i^t indicates whether queen i has at least one unblocked neighbor at step t, and v_i^t indicates whether the activated queen has a non-empty kill neighborhood. The core kill mechanism for $p=1$ is encoded entirely by u_{ij}^t and constraint (11).

Model.

$$\min_{x,y,z,u} \quad \Psi_{\text{MIP}} := \sum_{t=1}^{H} \sum_{1 \le i < j \le n} z_{ij}^{t} \tag{5}$$

$$\begin{aligned}
\text{s.t.} \quad & x_i^0 = 1, && \forall i \in N, && (6)\\
& \textstyle\sum_{i=1}^{n} y_i^t = 1, \qquad y_i^t \le x_i^{t-1}, && \forall i,\ t \ge 1 && (7)\\
& z_{ij}^t \le a_{ij}, \quad z_{ij}^t \le x_i^{t-1}, \quad z_{ij}^t \le x_j^{t-1}, && \forall i < j,\ t \ge 1 && (8)\\
& z_{ij}^t \le 1 - x_k^{t-1}, && \forall k:\ b_{ijk} = 1, i < j, t \ge 1 && (9)\\
& z_{ij}^t \ge x_i^{t-1} + x_j^{t-1} + a_{ij} - 2 - \textstyle\sum_{k: b_{ijk}=1} x_k^{t-1} && \forall i < j,\ t \ge 1 && (10)\\
& u_{ij}^t \le y_i^t, \quad u_{ij}^t \le z_{\min(i,j),\max(i,j)}^t, \quad u_{ij}^t \ge y_i^t + z_{\min(i,j),\max(i,j)}^t - 1, && \forall i \ne j,\ t \ge 1 && (11)\\
& x_j^t = x_j^{t-1} - \textstyle\sum_{i \ne j} u_{ij}^t, && \forall j,\ t \ge 1 && (12)\\
& \beta_i^t \ \ge\ z_{\min(i,j),\max(i,j)}^t, && \forall, j \ne i,\ i \in [n],\ t \ge 1 && (13)\\
& \beta_i^t \le \sum_{j \ne i} z_{\min(i,j),\max(i,j)}^t, && \forall\, i \in [n], t \ge 1 && (14)\\
& v_i^t \le y_i^t, \quad v_i^t \le \beta_i^t, \quad v_i^t \ge y_i^t + \beta_i^t - 1, && \forall\, i, t && (15)\\
& \sum_{j \ne k} z_{\min(k,j),\max(k,j)}^t \ \le\ n\,(1 - x_k^{t-1})\ +\ n \sum_i v_i^t, && \forall k \in [n],\ t \ge 1 && (16)\\
& z_{ij}^{H+1} \text{ defined from } x^H \text{ via (8)–(10) with } t = H+1. && && (17)
\end{aligned}$$

Remark 39. *For $p = 1$, Φ_{MIP} computed by the formulation above equals $\Psi(\pi^\star, 1)$ for the policy $\pi^\star$ encoded by $(y^t)_{t=1}^{H}$, and an optimal solution of the MIP gives $\Psi^\star(1) = \min_\pi \Psi(\pi, 1)$; Also, $\Psi^\star(1) \le \Psi^\star(p)$ for all $p \in (0,1)$, so Ψ_{MIP} furnishes a* lower bound *usable for pruning in the branch-and-bound search for $p < 1$, the number of binary variables is $O(n^3 H) = O(n^4)$, dominated by z_{ij}^t and u_{ij}^t.*

Lemma 40 (Blocking constraints are exact)**.** *For binary $x_i^{t-1}, x_j^{t-1}, x_k^{t-1}$ and fixed data a_{ij}, b_{ijk}, any feasible integral point of* (8)–(10) *satisfies $z_{ij}^t = \text{CanKill}^{(t)}(q_i, q_j)$ exactly.*

Proof. ($\Rightarrow$) If $z_{ij}^t = 1$, then (8) forces $x_i^{t-1} = x_j^{t-1} = a_{ij} = 1$, and for *every* k with $b_{ijk} = 1$, the corresponding instance of (9) forces $x_k^{t-1} = 0$. Hence q_i, q_j are alive, mutually attacking, and no blocker survives — i.e. $\text{CanKill}^{(t)}(q_i, q_j) = 1$.

($\Leftarrow$) If $\text{CanKill}^{(t)}(q_i, q_j) = 1$, then $x_i^{t-1} = x_j^{t-1} = a_{ij} = 1$ and $x_k^{t-1} = 0$ for every blocker k, so the right-hand side of (10) equals $1 + 1 + 1 - 2 - 0 = 1$, forcing $z_{ij}^t \ge 1$, i.e. $z_{ij}^t = 1$ by binarity and (8)–(9).

If instead $\text{CanKill}^{(t)}(q_i, q_j) = 0$: either $x_i^{t-1} = 0$, $x_j^{t-1} = 0$, or $a_{ij} = 0$, in which case (8) directly forces $z_{ij}^t = 0$; or some blocker k has $x_k^{t-1} = 1$, in which case the corresponding instance of (9) forces $z_{ij}^t = 0$. In every case $z_{ij}^t = 0$. ■

6. More theoretical results

Kill-opportunity count

Definition 41. *Let $\omega \in \Omega$ be a sample path. The* activation count *is the number of* distinct *queens activated at least once during the trajectory. $\mathscr{C}(\omega) := |\{\sigma_t(\omega) : t = 1, \ldots, T(\pi, \omega)\}| \in \{1, \ldots, n\}$.*

We say that policy π achieves full activation *if $\mathscr{C}(\omega) = n$ almost surely, i.e., every queen in $Q^{(0)}$ is activated at least once on every sample path before the process terminates.*

Remark 42 (Policies achieving full activation)**.** *Full activation ($\mathscr{C} = n$ a.s.) is achievable only if no queen is eliminated before being activated. A* round-robin *policy, which cycles through a fixed ordering of $Q^{(0)}$ and skips eliminated queens, achieves $\mathscr{C} = n$ a.s. provided the process does not absorb before all queens have been visited. Under a greedy policy (max kill-neighborhood), a queen with a small kill set may be eliminated by another queen's activation before being selected itself, yielding $\mathscr{C} < n$ on some paths. The condition $\mathscr{C} = n$ is therefore a constraint on the policy, not on the model.*

Definition 43 (Kill-opportunity count and condition $\mathbf{K}_n$)**.** *For a board state $S^{(t)} = (Q^{(t)}, p^{(t)})$, the* total directed kill-opportunity count *is*

$$K^{(t)} := \sum_{q \in Q^{(t)}} \left| N^{(t)}(q) \right| = \left| \mathscr{E}^{(t)} \right|_{\text{dir}},$$

where $|\mathscr{E}^{(t)}|_{\text{dir}}$ denotes the number of directed (ordered) unblocked attack edges in the attack graph $G^{(t)}$.

We say that Condition $\mathbf{K}_n$ holds throughout the trajectory if

$$K^{(t)} \geq n \quad \text{for every non-terminated step } t \leq T(\pi).$$

Remark 44. *Unlike $|Q^{(t)}|$, the count $K^{(t)}$ is* not monotone non-increasing*; when an intermediate queen (one lying in* $\text{Path}(q_i, q_j)$ *for some pair (q_i, q_j)) is eliminated, the path becomes unblocked and $K^{(t)}$ can increase strictly. Condition $\mathbf{K}_n$ therefore does not automatically fail as queens are removed; on dense initial boards it may hold throughout the entire trajectory. Condition $\mathbf{K}_n$ asserts that there are always at least n kill-attempt edges available across all surviving queens, so that a single round of n activations covers at least n distinct kill-attempt trials in expectation.*

Remark 45. *The conditions $\mathscr{C}(\omega) = n$ and $\mathbf{K}_n$ are complementary. $\mathscr{C} = n$ is a* policy-side *condition: it ensures breadth of coverage by requiring the activation rule to visit every queen. $\mathbf{K}_n$ is a* board-side *condition. It ensures that the graph structure supports at least n kill-attempt trials per n-step round, regardless of which queen is activated.*

Together they guarantee that in every round of n activations, at least one kill occurs with probability $\geq 1 - (1-p)^n$, which is the tighter kill-phase probability that presented in Proposition 14. These conditions together form the minimal sufficient set for deriving the bound $\frac{n-1}{1-(1-p)^n} + \frac{n(n-1)}{2}$ on the expected stopping time measured in rounds.

Per-queen kill-attempt count

Definition 46 (Per-queen kill-attempt count)**.** *For each queen $q_j \in Q^{(0)}$ and a given sample path ω, the* kill-attempt count *for q_j is*

$$\widetilde{K}_j^{(\pi)}(\omega) := \left| \left\{ t \leq T(\pi, \omega) : q_j \in N^{(t)}(\sigma_t(\omega)) \right\} \right|.$$

This counts the number of steps at which an activated queen targets q_j with an unblocked kill attempt. Note that $\widetilde{K}_j^{(\pi)}$ is automatically zero for steps after q_j has been eliminated (since $q_j \notin Q^{(t)}$ implies $q_j \notin N^{(t)}(\sigma_t)$), so $\widetilde{K}_j^{(\pi)}$ counts only attempts on q_j while q_j is alive.

Lemma 47 (Survival probability via kill-attempt count)**.** *For any policy π, any queen $q_j \in Q^{(0)}$, and $p \in (0, 1]$:*

$$Pr\left(q_j \in Q^{(T)}\right) = \mathbb{E}_\pi\left[(1-p)^{\widetilde{K}_j^{(\pi)}}\right].$$

Proof. Each kill attempt on q_j at step t is an independent $\text{Bern}(p)$ trial $\xi_j^{(t)}$; queen q_j is eliminated at the first step where $\xi_j^{(t)} = 1$, if any. Thus q_j survives the entire trajectory if and only if all $\widetilde{K}_j^{(\pi)}$ trials result in failure.

Formally, let $\mathscr{F}_T$ be the σ-algebra generated by all activation choices and all kill outcomes up to step T. Conditional on $\widetilde{K}_j^{(\pi)} = k$:

$$Pr\left(q_j \in Q^{(T)} \,\middle|\, \widetilde{K}_j^{(\pi)} = k\right) = (1-p)^k,$$

since the k trials are conditionally i.i.d. $\text{Bern}(p)$ given the activation choices (which determine *when* the trials occur but not their outcomes). Taking the expectation over $\widetilde{K}_j^{(\pi)}$ via the tower property gives the result. ■

Conflict upper bound

Proposition 48 (Cumulative conflict decomposition)**.** *Let π be any admissible adaptive policy and let $p \in (0, 1]$. For the original objective, we have the decomposition $\mathbb{E}_\pi[\Psi(\pi, p)] = \sum_{\{i,j\} \in E^{(0)}} \sum_{t=0}^{\infty} Pr_\pi\left((q_i, q_j) \in A(Q^{(t)})\right)$, where*

$E^{(0)} := \{\{i,j\} : q_i \overset{\text{atk}}{\sim} q_j,\ i<j\}$ *is the set of initially attacking unordered pairs. In particular,*

$$\mathbb{E}_\pi[\Psi(\pi,p)] \le |E^{(0)}|\,\mathbb{E}_\pi[\tau] \le \binom{n}{2}\mathbb{E}_\pi[\tau].$$

Proof. By definition of the cumulative objective, $\Psi(\pi,p) = \sum_{t=0}^{\tau-1}\kappa\left(S^{(t)}\right) = \sum_{t=0}^{\tau-1}|A(Q^{(t)})|$. Since every active conflict at time t must be one of the initially attacking pairs, we have $|A(Q^{(t)})| = \sum_{\{i,j\}\in E^{(0)}}\mathbf{1}\left[(q_i,q_j)\in A(Q^{(t)})\right]$. Therefore, $\Psi(\pi,p) = \sum_{t=0}^{\tau-1}\sum_{\{i,j\}\in E^{(0)}}\mathbf{1}\left[(q_i,q_j)\in A(Q^{(t)})\right]$. All terms are nonnegative, so Tonelli's theorem and linearity of expectation give $\mathbb{E}_\pi[\Psi(\pi,p)] = \sum_{\{i,j\}\in E^{(0)}}\sum_{t=0}^{\infty}Pr_\pi\left((q_i,q_j)\in A(Q^{(t)})\right)$. For the bound, observe that $|A(Q^{(t)})| \le |E^{(0)}|$ for every t, since no new attacking pair can appear outside the initial attack skeleton. Hence $\Psi(\pi,p) = \sum_{t=0}^{\tau-1}|A(Q^{(t)})| \le \sum_{t=0}^{\tau-1}|E^{(0)}| = |E^{(0)}|\,\tau$. Taking expectations yields the stated equation. ■

Proposition 48 provides a general decomposition of the expected cumulative objective in terms of how often initially attacking pairs remain active over time. This representation is useful because it separates the objective into a structural part, determined by the conflict graph, and a stochastic part, determined by the policy through the kill-attempt counts.

To obtain more explicit bounds, we next relate the cumulative cost to the stopping time and then specialize to policies with full activation and sufficient kill opportunities. In particular, the cumulative objective is always bounded by the number of initially attacking pairs multiplied by the absorption time, and Proposition 14 yields a direct estimate for that absorption time.

Remark 49 (Bottleneck structure under cumulative cost)**.** *The main effect of a bottleneck queen under the cumulative objective is to prolong the time during which conflicts remain active. In other words, a queen that blocks many attack lines can keep the active-conflict count $|A(Q^{(t)})|$ large for many steps, so the cumulative cost can grow substantially even when the terminal configuration is eventually sparse. Thus, the bottleneck phenomenon should be interpreted primarily through its effect on the stopping time τ and on the duration of high-conflict phases, rather than through the terminal number of surviving conflicts alone. In particular, $\Psi(\pi,p) \le |E^{(0)}|\tau \le \binom{n}{2}\tau$, so any structural feature that delays eliminations increases the cumulative objective proportionally through τ.*

General cumulative-cost bound

Corollary 50. *For any admissible adaptive policy π and any $p \in (0,1]$,*

$$\mathbb{E}_\pi[\Psi(\pi,p)] \le \binom{n}{2}\mathbb{E}_\pi[\tau] \le \binom{n}{2}\left(\frac{n-1}{p} + \frac{n(n-1)}{2}\right).$$

Proof. By the definition of the cumulative objective, $\Psi(\pi,p) = \sum_{t=0}^{\tau-1}|A(Q^{(t)})|$. Since the active-conflict set at any time is a subset of the initial attack skeleton, $|A(Q^{(t)})| \le |E^{(0)}| \le \binom{n}{2}$, for all t. Therefore, $\Psi(\pi,p) \le \binom{n}{2}\tau$. Taking expectations yields $\mathbb{E}_\pi[\Psi(\pi,p)] \le \binom{n}{2}\mathbb{E}_\pi[\tau]$. ■

Sharper cumulative-cost bound under full activation and condition $\mathbf{K}_n$

Corollary 51. *Let π be an admissible round-robin policy with $\mathscr{C}(\omega) = n$ almost surely, and assume Condition $\mathbf{K}_n$ holds throughout the trajectory. Then*

$$\mathbb{E}_\pi[\Psi(\pi,p)] \le \binom{n}{2}\left(\frac{n-1}{1-(1-p)^n} + \frac{n(n-1)}{2}\right).$$

If, in addition, the process completes at least one full round before termination, then the same bound implies

$$\mathbb{E}_\pi[\Psi(\pi,p)] \le \binom{n}{2}\left(\frac{n-1}{1-(1-p)^n} + \frac{n(n-1)}{2}\right),$$

with the round term interpreted as at least one completed round.

Proof. Under the assumptions $\mathscr{C}(\omega) = n$ and Condition $\mathbf{K}_n$, Remark 45 implies that each round of n activations contains at least n kill-attempt trials, so a kill occurs in a round with probability at least $1-(1-p)^n$. Hence the number of effective rounds is stochastically dominated by a negative-binomial waiting time with success parameter $1-(1-p)^n$ and target $n-1$ eliminations, giving $\mathbb{E}_\pi[\tau] \leq \frac{n-1}{1-(1-p)^n} + \frac{n(n-1)}{2}$. Combining this with the general bound $\mathbb{E}_\pi[\Psi(\pi, p)] \leq \binom{n}{2}\mathbb{E}_\pi[\tau]$ yields the stated result. ■

Remark 52 (Hierarchy of cumulative-cost bounds)**.** *The cumulative-cost bounds satisfy*

$$\binom{n}{2}\left(\frac{n-1}{p} + \frac{n(n-1)}{2}\right) \geq \binom{n}{2}\left(\frac{n-1}{1-(1-p)^n} + \frac{n(n-1)}{2}\right) \geq \mathbb{E}_\pi[\Psi(\pi, p)].$$

The first inequality uses $1-(1-p)^n \geq p$, and the second follows from the preceding corollaries.

7. Numerical results

We now illustrate a concrete instance of the n-queens board with $n = 8$. An initial configuration $Q^{(0)}$ was drawn uniformly at random (without replacement) from the 64 cells of the board, giving the static attack data (a_{ij}, b_{ijk}) and the conflict count $\kappa(s_0) = 9$; the resulting board and its (undirected) attack graph $G^{(0)}$ are shown in Figure 2. Every method compared below (the exact backward-induction value function of Theorem 17(i), the MIP of Section 5, the ACG policy π_{ACG}, the greedy-degree policy π_G, and a uniformly random admissible baseline) was implemented against the *same* precomputed static data (a_{ij}, b_{ijk}), so that any discrepancy between the exact solver, the heuristics, and the MIP could only be attributable to a modeling error rather than to inconsistent problem instances. The MIP was solved to global optimality with the CBC branch-and-cut solver; the exact value function was computed by dynamic programming over the reachable state space $\mathscr{S}_{\text{reach}}$ of Remark 2 ($2^8 = 256$ states); and, for $p < 1$, each policy was evaluated by 4,000 independent Monte Carlo replications of the stochastic kill mechanism of Definition 5(iv).

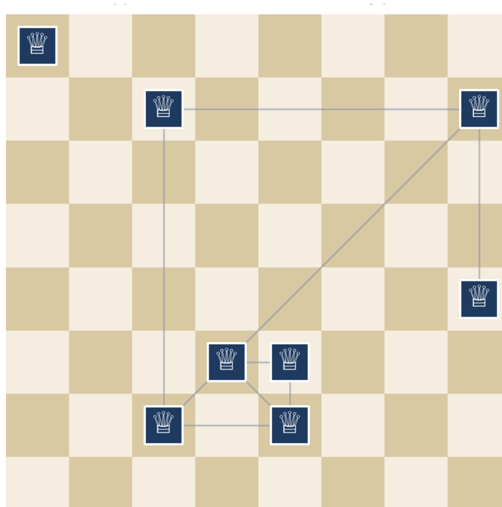

Figure 2: The $n = 8$ test instance $Q^{(0)}$ used throughout Section 7. Grey edges join pairs $\{q_i, q_j\}$ with $\text{CanKill}^{(0)}(q_i, q_j) = 1$; there are $\kappa(s_0) = 9$ such unblocked attacking pairs.

7.1. Deterministic case ($p = 1$)

Table 1 reports, for $p = 1$, the cost Ψ, the number of effective activation rounds τ, and the wall-clock solve time for exact method was 0.0029 second, and the solver time for MIP optimum Ψ_{MIP} was 0.3296 second. The MIP optimum $\Psi_{\text{MIP}} = 9$ coincides *exactly* with the backward-induction value $V^*(s_0) = 9$ obtained independently from Theorem 17(i), numerically confirming the identity $\Psi^\star(1) = \min_\pi \Psi(\pi, 1)$ of Remark 39. Replaying the MIP's optimal activation sequence $(y_i^t)_{t=1}^H$ bit-for-bit through the exact transition mechanics of Definition 5 reproduces $\Psi = 9$ in a single round ($\tau = 1$), which is exactly the cross-check licensed by Lemma 40: since $z_{ij}^t = \mathrm{CanKill}^{(t)}(q_i, q_j)$ holds at every feasible integral point, the simulator and the MIP can never disagree on this instance. On this particular board a single "hub" activation happens to resolve all nine conflicts simultaneously, so both π_{ACG} and π_G also recover the optimum; the uniformly random admissible baseline, in contrast, requires $\tau = 3$ rounds and incurs $\Psi = 16$, roughly 1.8 times the optimal cost, illustrating that even in the deterministic regime an unstructured activation order can be markedly wasteful.

Table 1: Deterministic SQEP ($p = 1$, $n = 8$): MIP optimum versus the exact value function of Theorem 17 and three activation policies.

Method	Ψ	τ
Exact $V^*(s_0)$ (backward induction)	9.00	—
MIP optimum Ψ_{MIP}	9.00	—
MIP solution replayed in simulator	9.00	1
Greedy-degree π_G rollout	9.00	1
ACG π_{ACG} rollout	9.00	1
Random admissible rollout	16.00	3

7.2. Stochastic case ($p < 1$)

For $p \in \{0.20, 0.35, 0.50, 0.70, 0.85, 1.00\}$ we recomputed $V^*(s_0)$ by backward induction and estimated $\mathbb{E}_\pi[\Psi(\pi, p)]$ for each of the three policies by Monte Carlo simulation; Table 2 reports the resulting means (with Monte Carlo standard errors in parentheses), the worst-case bound $\frac{n}{2p}V^*(s_0)$ of Theorem 34, and the ratio $\mathbb{E}_{\pi_{\mathrm{ACG}}}[\Psi]/V^*(s_0)$. Figure 3 displays the same data graphically. Panel (a) compares the three policies against $V^*(s_0)$ on a linear scale, and panel (b) checks the ACG guarantee of Theorem 34 on a logarithmic scale.

Three observations stand out. First, $V^*(s_0)$ decreases monotonically as p grows, from 40.56 at $p = 0.20$ to 9.00 at $p = 1.00$, consistent with the monotonicity of $p \mapsto \mathbb{E}_\pi[\Psi(\pi, p)]$ established in Proposition 25. Second, π_{ACG} tracks $V^*(s_0)$ extremely closely at every value of p tested — the ratio $\mathbb{E}_{\pi_{\mathrm{ACG}}}[\Psi]/V^*(s_0)$ never exceeds 1.02 — which is far tighter than the worst-case guarantee of Theorem 34: the bound $\frac{n}{2p}V^*(s_0)$ ranges from 36.00 at $p = 1$ up to 811.11 at $p = 0.20$ on this instance, confirming Remark 33's observation that the ratio $n/(2p)$ is approached only on adversarially structured conflict graphs (e.g. a perfect matching) and is essentially never attained on a generic board. Third, the greedy-degree policy π_G is uniformly at least as costly as π_{ACG}, with the largest gap at small p (e.g. 43.18 versus 41.32 at $p = 0.20$) and the two coinciding as $p \to 1$; this is consistent with the discounting argument of the discussion following Definition 29, since a pair $\{q_i, q_j\}$ with both endpoints in $N^{(t)}(a)$ contributes $1 - (1-p)^2 > p$ to

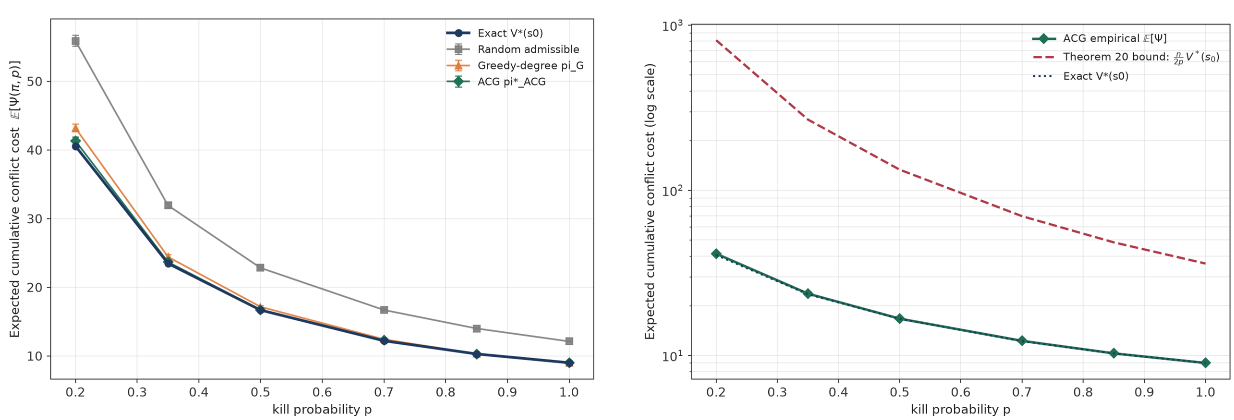

(a) Policies versus $V^*(s_0)$, linear scale. **(b)** ACG versus the Theorem 34 bound, log scale.

Figure 3: Policy comparison on the $n = 8$ instance of Figure 2 across the kill probability p. Error bars in panel (a) show ± 1.96 Monte Carlo standard errors.

Table 2: Stochastic SQEP ($n = 8$): exact value function versus Monte Carlo policy evaluation (4,000 trials per policy per p; standard errors in parentheses).

p	$V^*(s_0)$	$\mathbb{E}[\Psi]$, Random	$\mathbb{E}[\Psi]$, π_G	$\mathbb{E}[\Psi]$, π_{ACG}	Bound $\frac{n}{2p}V^*(s_0)$	$\mathbb{E}_{\pi_{\text{ACG}}}[\Psi]/V^*(s_0)$
0.20	40.56	55.88 (0.41)	43.18 (0.29)	41.32 (0.30)	811.11	1.019
0.35	23.46	31.91 (0.21)	24.42 (0.15)	23.67 (0.15)	268.15	1.009
0.50	16.67	22.81 (0.14)	17.14 (0.10)	16.69 (0.09)	133.33	1.002
0.70	12.20	16.68 (0.08)	12.41 (0.05)	12.25 (0.05)	69.70	1.004
0.85	10.28	13.99 (0.06)	10.29 (0.03)	10.29 (0.03)	48.38	1.001
1.00	9.00	12.11 (0.03)	9.00 (0.00)	9.00 (0.00)	36.00	1.000

$G(s,a)$ but only a single unit to $|N^{(t)}(a)|$, so π_{ACG} correctly rewards actions that jointly resolve many conflicts while π_G cannot distinguish them from actions with the same degree but less overlap. The uniformly random admissible baseline is markedly worse throughout, incurring between roughly 35% and 37% more expected cumulative conflict than π_{ACG} across the tested range of p (minimum gap at $p = 1$, maximum near $p = 0.5$), which quantifies the practical value of structure-aware activation even though the random policy remains admissible in the sense of Remark 6.

7.3. Conflict-count decay along sample trajectories

Figure 4 plots the average value of $\kappa(S^{(t)})$ against the round index t, at $p = 0.5$, over 800 independent trajectories per policy, together with the reference curve $\kappa(s_0)(1 - p/n)^t$. All three policies decay noticeably *faster* than this reference curve, which is expected.

Consistent with Table 2, π_{ACG} and π_G track one another closely and both dominate the random baseline throughout the trajectory, with the gap widening in the number of surviving conflicts rather than in the number of rounds elapsed.

A single random board is of course not a substitute for a broad empirical study across many instances and larger n; we also leave a systematic study over random and adversarial board families to future work.

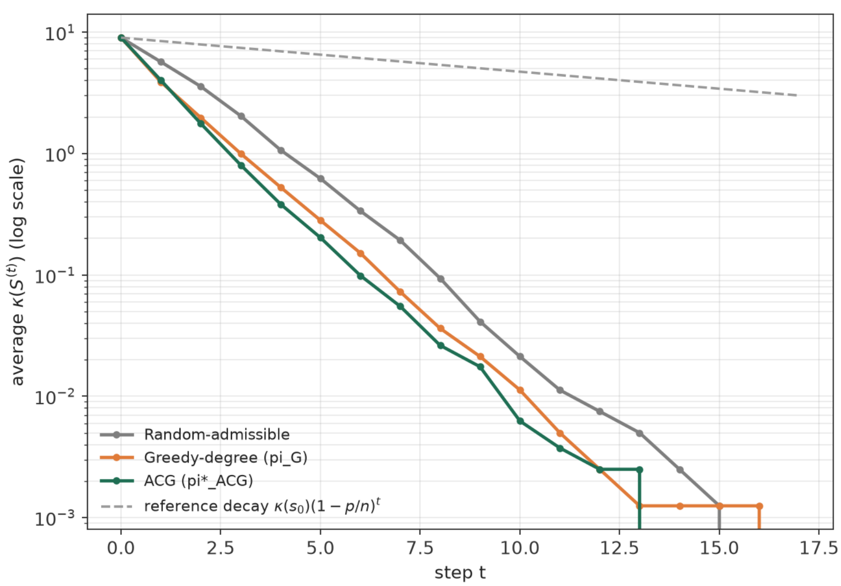

Figure 4: Average conflict count $\kappa(S^{(t)})$ per round under each policy, $p = 0.5$, $n = 8$, averaged over 800 trajectories (logarithmic scale). The dashed curve is the generic reference decay $\kappa(s_0)(1 - p/n)^t$.

8. Conclusion

This paper introduced the Stochastic Sequential Queens Elimination Problem (S-SQEP), a dynamic, elimination-based counterpart to the classical n-queens problem in which an activated queen simultaneously and independently attempts to eliminate every queen in her unblocked neighborhood, minimizing expected cumulative conflict. We cast S-SQEP as a finite-state Markov decision process, established almost-sure absorption with a bound on the expected stopping time, and showed that the natural degree-greedy heuristic is not universally optimal. We proposed the Adaptive Coverage Greedy policy, with a closed-form one-step gain and a proved worst-case approximation ratio, alongside an exact mixed-integer program for the deterministic case $p = 1$. Numerical experiments on an $n = 8$ instance confirmed these results; the mixed-integer programming optimum matched the exact value function, and Adaptive Coverage Greedy tracked it within 2% across all tested kill probabilities.

Several directions remain open. On the theoretical side, future work may focus on advancing the mathematical analysis of the problem through stronger structural characterizations, improved optimization formulations, tighter complexity and approximation results, and more efficient algorithmic paradigms for solving larger instances. On the modeling side, natural extensions include sequential (single-target) rather than simultaneous multi-kill activation, heterogeneous or state-dependent kill probabilities, partially observed variants in which p itself must be learned online, and a systematic empirical study across random and adversarial board families and larger n. We view S-SQEP as a clean and combinatorially rich test bed for these questions, bridging the static n-queens literature with the controlled stochastic-elimination processes studied in the Markov decision processes, cascade, and interacting-particle-system literatures.

Appendix A: Further discussion on related work

On the n-queens problem

This subsection expands on the brief historical remarks made in the Introduction with a fuller literature review of the n-queens problem. Much of the content in this part was inspired by the excellent work by Bowtell and Keevash [6].

The origin of the n-Queens problem is deeply embedded in recreational mathematics and chess composition. The specific 8-queens problem was first published in 1848 by the chess composer Max Bezzel [5]. In 1850, Nauck solved the 8-queens problem by discovering all 92 valid configurations [28]. Concurrently, Carl Friedrich Gauss investigated the problem, attempting to reformulate it in terms of the arithmetic representation of complex numbers, though he originally found only 72 solutions before learning of Nauck's complete enumeration, as reported in Campbell [8].

From a graph-theoretic perspective, the problem translates to finding an independent set of size n in a graph containing n^2 vertices (representing the squares), where edges connect vertices lying in the same row, column, or diagonal. Alternatively, the n-queens problem has been elegantly modeled as counting perfect matchings in a 4-partite, 4-uniform hypergraph $\mathscr{Q}(n)$, where the vertices represent the n rows, n columns, $2n-1$ forward diagonals, and $2n-1$ backward diagonals (see Bowtell and Keevash [6]).

The n-queens problem has continually acted as a proving ground for new search and optimization methodologies. *Exact algorithms* primarily utilize depth-first search (DFS) and backtracking, famously highlighted by Dijkstra [13] to demonstrate the elegance of recursive programming. From *heuristic and local search* perspective, however, the problem is to find single solutions for exceptionally large n, metaheuristics such as the Min-Conflicts algorithm and local search have been highly successful (see Sosic and Gu [33] and Masehian, Akbaripour, and Mohabbati-Kalejahi [27]). They start with an imperfect board and iteratively repair constraint violations, breaking away from the strict determinism of exhaustive search but still aiming to solve a deterministic set of constraints.

Algorithmic advances using *Probabilistic Matchings* present the state-of-the-art counting algorithms, which employ the differential equations method to analyze random greedy matching processes, as utilized by Bennett and Bohman [1].

Beyond the classical n-queens problem, several variants have been investigated. Among the most prominent is the Toroidal n-queens problem, introduced by Pólya [30], in which opposite edges of the board are identified, resulting in wrap-around diagonals. Pólya established the necessary and sufficient condition for the existence of solutions, and recent advances have completely characterized the asymptotic number of toroidal solutions (see Luria [25]). In parallel, Luria and Simkin [26] obtained asymptotically matching lower bounds for the classical counting problem. Despite their importance, these studies focus on asymptotic enumeration and toroidal board structures, which differ substantially from the problem addressed in the present paper; accordingly, they are not discussed further.

The literature reveals a distinct evolution from simple configurations to deeply mathematical structural proofs. However, the core of the problem has never deviated from strict Boolean outcomes. Almost all of this literature, classical and recent alike, concerns *static* placement, domination, or counting problems on a fixed board; dynamic, elimination-based variants such as the one studied in this paper remain comparatively unexplored, which motivates situating the present model within the four further bodies of work discussed below.

On Markov decision processes and stochastic shortest paths

The MDP formulation of Definition 1 and its analysis in Section 3 are grounded in the general theory of Markov decision processes, for which Puterman [31] remains the canonical reference for finite- and infinite-horizon formulations, the Bellman equation, and policy- and value-iteration convergence. Within this theory, the *stochastic shortest path* (SSP) framework of Bertsekas and Tsitsiklis [2] is the natural home for S-SQEP; where an SSP is an undiscounted, total-cost MDP with an absorbing goal set, and it strictly generalizes finite-horizon, discounted, and average-cost MDPs as special cases. The regularity conditions (A1)–(A2) of Bertsekas and Tsitsiklis [2], requiring properness of every admissible stationary policy and infinite cost for every improper one, are exactly the conditions verified in Lemma 16 and Theorem 17; the finiteness of the reachable state space $\mathscr{S}_{\text{reach}}$ established in Remark 2 is precisely what allows these classical conditions to be applied directly, without the more delicate machinery developed by Bertsekas and Yu [4] for SSPs with unbounded costs or state spaces that are not finite a priori.

Much of the contemporary SSP literature has moved in a different direction; rather than analyzing a fixed, fully known MDP, recent work studies SSPs under *unknown* transition probabilities, with an emphasis on regret and sample-complexity guarantees for online or decentralized learning; a representative recent instance is Chavan, Trivedi and Nandyala [10], which derives regret lower bounds for decentralized multi-agent SSPs. This is a complementary but distinct research axis from the one pursued here; where the kill probability p in S-SQEP is assumed known throughout, and the questions asked are structural and combinatorial, namely well-posedness of the value function, the suboptimality of a natural greedy policy, and a closed-form approximation ratio for the ACG policy, rather than statistical guarantees for learning the dynamics from interaction. In this sense, the present paper is closer in spirit to the classical, fully-informed treatment of Bertsekas and Tsitsiklis [2] and Puterman [31] than to the learning-theoretic strand of the modern SSP literature.

On cascade and elimination processes on graphs

The simultaneous multi-kill mechanism of Definition 5(iv), in which an activated queen independently eliminates each member of its unblocked neighborhood with probability p, is structurally the same Bernoulli-activation primitive used in the *independent cascade* (IC) model of Kempe, Kleinberg and Tardos [21]; where in the IC model, a newly active node independently activates each inactive out-neighbor with a fixed edge probability, and the diffusion process propagates in discrete rounds until no further activation is possible. The objectives differ sharply, however. Influence maximization asks which seed set S maximizes the *expected size* of the eventually active set, a monotone submodular objective admitting a $(1-1/e)$-greedy guarantee; S-SQEP instead fixes the full initial configuration $Q^{(0)}$ and sequentially chooses, at every step, *which already-alive queen to fire* to minimize the *expected cumulative* conflict count $\Psi(\pi,p)$ along the whole trajectory (Definition 1), a substantially different combinatorial-control problem despite the shared stochastic-activation primitive. The *adaptive* influence-maximization literature is closer in structure. Chen and Peng [11] studied the adaptivity gap of IC-model influence maximization under full-adoption feedback, where the seed-selection policy observes the realized cascade before choosing the next seed, mirroring the way the S-SQEP policy π observes the realized board state $S^{(t-1)}$ before selecting σ_t (Definition 5(ii)); the admissibility condition of Remark 6 and the equivalence of randomized, deterministic, and admissible optima established in Theorem 17(iv) are structural counterparts of the adaptivity results in that line of work. Greedy, degree-based heuristics for cascade control remain an active topic, illustrated by the recent homophily-aware dynamic influence-maximization heuristic of [34].

A second, closely related family is the *contact process*, introduced by Harris [18] and studied extensively as an interacting particle system (see the monograph of Liggett [24]); the contact process is a continuous-time Markov process on a graph in which each infected vertex infects each susceptible neighbor at rate λ and recovers at rate 1, and its central question is the almost-sure extinction time on finite graphs. This is a close probabilistic cousin of the absorption results of Section 3. Theorem 12 and Proposition 14 are, in spirit, discrete-time, finite-population analogues of contact-process extinction-time results. The essential distinction is control; where the contact process has no policy, since infection and recovery fire at exogenously fixed rates, whereas σ_t in S-SQEP is a decision variable chosen by

π at every step, which is precisely why an MDP/dynamic-programming treatment is required rather than a purely probabilistic extinction-time analysis. Recent contact-process work continues to refine extinction and phase-transition results on structured and random graphs; Jahnel, Lüchtrath and Mönch [20], for instance, establish sharp conditions for a non-trivial survival/extinction phase transition on one-dimensional random networks via a coupling with a random walk in a random environment that exploits the *attractiveness* (monotonicity) of the contact process under stochastic domination of initial configurations. This attractiveness property is the same qualitative idea underlying the monotone coupling in this research; where both arguments couple two realizations of a stochastic elimination process via a common source of randomness and conclude a pathwise domination of surviving configurations. The coupling in Appendix D is comparatively elementary, since it acts on a finite, discrete-time process with a simultaneous multi-kill step rather than a continuous-time infinite lattice, but the underlying monotonicity principle is the same one that drives contact-process comparison arguments.

On heuristic search, rollout, and lookahead policies

The ACG policy is naturally read as a one-step lookahead, or *rollout*, policy in the sense formalized by Bertsekas, Tsitsiklis and Wu [3]; where a rollout policy selects, at each state, the action that optimizes a one-step cost plus an estimate of the cost-to-go under a fixed base heuristic, and inherits a cost-improvement guarantee relative to that base heuristic under mild conditions. In S-SQEP, the role of the base heuristic is played not by simulating another policy but by the *exact* closed-form one-step conflict-reduction gain $G(s,a)$, so that ACG can be understood as a rollout policy with respect to an exact, rather than sampled, one-step lookahead, which is precisely what permits the closed-form approximation guarantee of Theorem 34 rather than only a Monte Carlo estimate of policy quality.

When $G(s,a)$ is too costly to evaluate exactly, Remark 30 falls back to M i.i.d. samples of the kill distribution with a Hoeffding-inequality error bound of order $O(\kappa(s)/\sqrt{M})$; this exact-versus-sampled evaluation tradeoff is the same one that motivates Monte Carlo Tree Search (MCTS): the UCT algorithm of Kocsis and Szepesvári [22] and the broader MCTS framework surveyed by [7] combine random rollouts with bandit-guided tree expansion to allocate a fixed simulation budget efficiently, and finite-sample concentration bounds for the resulting action-value estimates play exactly the role that the Hoeffding bound of Remark 30 plays for the sampled variant of ACG. Certifying such sampled lookahead policies remains an active research direction, as illustrated by recent work applying Monte Carlo tree search to general-utility MDPs in the single-trial regime, which surveys polynomial regret-concentration guarantees for MCTS-style algorithms established elsewhere in the literature as shown in Santos, Sardinha and Melo [32]; the present paper's Lemma 31 sidesteps this concentration problem altogether whenever n is small enough for the exact sum in Equation (1) to be tractable.

The dynamic-programming solution of this research, computed exactly by backward induction over the finite state space $\mathscr{S}_{\text{reach}}$ in Section 7, can equivalently be viewed as an exhaustive search of an implicit AND/OR graph; where OR nodes correspond to board states at which the policy chooses which queen to activate, and AND nodes correspond to the (up to) $2^{|N^{(t)}(\sigma_t)|}$ stochastic kill outcomes of that activation. Admissible-heuristic AND/OR search algorithms were developed precisely to avoid exhaustively expanding such graphs: the classical A* algorithm of Hart, Nilsson and Raphael [19] prunes an ordinary OR-graph using an admissible cost-to-go estimate, its AND/OR generalization AO*, Nilsson [29] extends the same admissibility-based pruning to acyclic AND/OR graphs, and Hansen and Zilberstein [17] extend AO* to cyclic AND/OR graphs via the LAO* algorithm, replacing simple backward cost propagation with value or policy iteration on the explored fragment of the graph. The $n \cdot 2^n$ growth in the number of (state, action) pairs flagged in Remark 2 as the reason to prefer policy-based methods over exhaustive value iteration for moderate n is exactly the scaling concern that motivates AO*/LAO*-style admissible pruning over exhaustive AND/OR search; recent work on learning admissible heuristics with formal sample-complexity and generalization guarantees for A*-style search in Futuhi and Sturtevant [15] suggests a concrete avenue for extending the exact dynamic-programming approach of Section 7 to values of n beyond the $n = 8$, 256-state instance studied there, by replacing exhaustive backward induction with an admissibly pruned AND/OR search.

On mixed-integer programming for blocking and visibility constraints

The MIP model of Section 5 encodes the time-indexed blocking predicate $\text{CanKill}^{(t)}(q_i, q_j)$. This is an instance of a broader tradition of encoding line-of-sight and visibility logic as combinatorial constraints, whose origin is usually traced to the art-gallery theorem of Chvátal [12]. The same per-obstacle conjunctive encoding used in (8)–(10) recurs throughout the mixed-integer motion-planning literature; where [9] formulate line-of-sight connectivity constraints for a chain of nonholonomic robots as a mixed-integer linear program, and research by Garcia, Hicks and Huchette [16] show more generally that the usual big-M formulation of "is a waypoint's line of sight blocked by an obstacle" is typically *not* the strongest possible with respect to the MIP's convex relaxation, proposing ideal combinatorial disjunctive-constraint formulations instead. Because $\mathrm{Block}(q_i, q_j)$ in Definition 7 ranges over a finite, explicitly enumerable set of candidate blocking cells on a discrete board, the exactness guaranteed by Lemma 40 is unaffected by this relaxation gap; nonetheless, the ideal-formulation techniques of Garcia, Hicks and Huchette [16] could in principle be applied to strengthen the LP relaxation of (8)–(10) without disturbing integral exactness, which is a natural direction for scaling the MIP of Section 5 beyond $n = 8$.

Within the n-queens literature specifically, integer programming has so far been applied almost exclusively to *static* formulations; Fischetti and Salvagnin [14] solve lexicographically-first and "most-beautiful" queens instances via an independent-set/clique ILP over the fixed non-attack graph, and Kunt [23] extend this line of work to higher-dimensional boards with strengthened valid inequalities that improve solver performance by an order of magnitude on benchmark instances. Both formulations impose a single, time-invariant independent-set constraint on a fixed board. The blocking predicate underlying (8)–(10), by contrast, is re-evaluated once per elimination round $t = 1, \dots, H$ against the shrinking survivor set $Q^{(t-1)}$, so the model carries a full copy of the z, u, and x variables for every round rather than a single static constraint block. To the author's knowledge, this round-indexed reformulation of the queens' visibility predicate as a dynamic, state-dependent blocking constraint has not previously appeared in the n-queens integer-programming literature, which reinforces the observation made in the Introduction that dynamic, elimination-based variants of the n-queens problem remain comparatively unexplored relative to static placement and domination formulations.

References

[1] Bennett, P. and Bohman, T. (2019). A natural barrier in random greedy hypergraph matching. *Combinatorics, Probability and Computing*, 28(6):816–825.

[2] Bertsekas, D. P. and Tsitsiklis, J. N. (1991). An analysis of stochastic shortest path problems. *Mathematics of Operations Research*, 16(3):580–595.

[3] Bertsekas, D. P., Tsitsiklis, J. N., and Wu, C. (1997). Rollout algorithms for combinatorial optimization. *Journal of Heuristics*, 3(3):245–262.

[4] Bertsekas, D. P. and Yu, H. (2013). Stochastic shortest path problems under weak conditions. Technical Report LIDS-P-2909, Laboratory for Information and Decision Systems, MIT.

[5] Bezzel, M. (1848). Proposal of 8-queens problem. *Berliner Schachzeitung*, 3:363. Submitted under the pseudonym "Schachfreund".

[6] Bowtell, C. and Keevash, P. (2021). The n-queens problem. *arXiv preprint arXiv:2109.08083*.

[7] Browne, C. B., Powley, E., Whitehouse, D., Lucas, S. M., Cowling, P. I., Rohlfshagen, P., Tavener, S., Perez, D., Samothrakis, S., and Colton, S. (2012). A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):1–43.

[8] Campbell, P. (1977). Gauss and the eight queens problem: A study in miniature of the propagation of historical error. *Historia Mathematica*, 4:397–404.

[9] Caregnato-Neto, A., Maximo, M. R. O. A., and Afonso, R. J. M. (2023). A mixed-integer approach for motion planning of nonholonomic robots under visible light communication constraints. *arXiv preprint arXiv:2306.15590*.

[10] Chavan, U., Trivedi, P., and Hemachandra, N. (2025). Regret lower bounds for decentralized multi-agent stochastic shortest path problems. In *Advances in Neural Information Processing Systems*, volume 38, pages 173948–173992. Curran Associates, Inc.

[11] Chen, W. and Peng, B. (2019). On adaptivity gaps of influence maximization under the independent cascade model with full adoption feedback. *arXiv preprint arXiv:1907.01707*.

[12] Chvátal, V. (1975). A combinatorial theorem in plane geometry. *Journal of Combinatorial Theory, Series B*, 18(1):39–41.

[13] Dijkstra, E. W. (1971). Ewd 316: A short introduction to the art of programming. Technical Report EWD316, Technische Hogeschool Eindhoven, Eindhoven, Netherlands.

[14] Fischetti, M. and Salvagnin, D. (2019). Finding first and most-beautiful queens by integer programming. *arXiv preprint arXiv:1907.08246*.

[15] Futuhi, E. and Sturtevant, N. R. (2026). Learning admissible heuristics for a*: Theory and practice. *arXiv preprint arXiv:2509.22626*.

[16] Garcia, R., Hicks, I. V., and Huchette, J. (2023). Combinatorial disjunctive constraints for obstacle avoidance in path planning. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 267–273.

[17] Hansen, E. A. and Zilberstein, S. (2001). Lao*: A heuristic search algorithm that finds solutions with loops. *Artificial Intelligence*, 129(1-2):35–62.

[18] Harris, T. E. (1974). Contact interactions on a lattice. *The Annals of Probability*, 2(6):969–988.

[19] Hart, P. E., Nilsson, N. J., and Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107.

[20] Jahnel, B., Lüchtrath, L., and Mönch, C. (2025). Phase transitions for contact processes on one-dimensional networks. *arXiv preprint arXiv:2501.16858*.

[21] Kempe, D., Kleinberg, J., and Tardos, É. (2003). Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146.

[22] Kocsis, L. and Szepesvári, C. (2006). Bandit based monte-carlo planning. In *Proceedings of the 17th European Conference on Machine Learning (ECML)*, pages 282–293, Berlin, Heidelberg. Springer Berlin Heidelberg.

[23] Kunt, T. (2025). Solving the n-queens problem in higher dimensions. In *Operations Research Proceedings 2024*, pages 205–211, Cham. Springer Nature Switzerland.

[24] Liggett, T. M. (1999). *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes*. Springer-Verlag.

[25] Luria, Z. (2017). New bounds on the number of n-queens configurations.

[26] Luria, Z. and Simkin, M. (2021). A lower bound for the *n*-queens problem.

[27] Masehian, E., Akbaripour, H., and Mohabbati-Kalejahi, N. (2013). Landscape analysis and efficient metaheuristics for solving the n-queens problem. *Computational Optimization and Applications*, 56(3):735–764.

[28] Nauck, F. (1850). Briefwechseln mit allen für alle. *Illustrirte Zeitung*, 377:182.

[29] Nilsson, N. J. (1980). *Principles of Artificial Intelligence*. Tioga Publishing.

[30] Pólya, G. (1918). Über die "doppelt-periodischen" lösungen des *n*-damen-problems. In Ahrens, W., editor, *Mathematische Unterhaltungen und Spiele*, volume 2, pages 364–374. B. G. Teubner, 2 edition.

[31] Puterman, M. L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.

[32] Santos, P. P., Sardinha, A., and Melo, F. S. (2026). Solving general-utility markov decision processes in the single-trial regime with online planning. *arXiv preprint arXiv:2505.15782*.

[33] Sosic, R. and Gu, J. (1994). Efficient local search with conflict minimization: A case study of the n-queens problem. *IEEE Transactions on Knowledge and Data Engineering*, 6(5):661–668.

[34] Wang, G., Du, S., Jiang, Y., and Li, X. (2025). A homophilic and dynamic influence maximization strategy based on independent cascade model in social networks. *Frontiers in Physics*, 12:1509905.