Adaptive Self-Preservation in Flocking Systems: How Dynamic Risk Thresholds Shape Collective Safety

Jiaman Li[1[0009−0009−9322−0400]]✉ and William G. Kennedy[1[0000−0001−9238−1215]]

George Mason University, Fairfax VA 22030, USA
jli30@gmu.edu, wkennedy@gmu.edu

Abstract. Agent-based flocking models effectively demonstrate how simple local rules produce collective behavior without central control. However, traditional models typically rely on fixed behavioral parameters, ignoring the biological reality that individuals dynamically regulate risk sensitivity based on experience. To address this, we extend the classic NetLogo flocking model with an adaptive self-preservation mechanism where each agent adjusts their personal collision threshold based on a memoryless asymmetric feedback rule ($\alpha \gg \beta$), modeling the biological principle that organisms respond quickly to danger but relax slowly when safe. Through systematic experiments (N=540 runs), we report three key findings: (1) emergent heterogeneity, where identical agents develop differentiated risk tolerances solely through interaction history; (2) the adaptation-to-decay ratio importantly determines emergent behavior across three distinct parameter regimes; and (3) robustness across a four-fold density range ($N = 100$–400). Unlike reinforcement learning (RL) approaches that require extensive training and reward engineering, our simple feedback rule produces emergent safety behaviors with constant-time per-agent threshold updates using only local binary stimuli. This work provides both a theoretical framework for understanding adaptive collective behavior and a practical tool for designing safer decentralized multi-agent systems.

Keywords: Adaptive self-preservation · Flocking · Emergent Heterogeneity · Decentralized control · Agent-based modeling

1 Introduction

In many complex systems, collective behavior emerges without any centralized leader. Examples range from bird flocks and fish schools to robot swarms and autonomous vehicle platoons. Each individual follows simple local interaction rules, yet the collective exhibits organized motion and coherent patterns [24]. This reflects a core principle of complexity science: macro-order arising from micro-interactions [20]

The classic Boids model introduced by Reynolds [19] first demonstrated how three simple rules, concerning alignment, cohesion, and separation, can produce

lifelike flocking behavior in computer simulations. In this model, each agent (the boids) aligns its direction with nearby neighbors (alignment), steers toward the local center of its group (cohesion), and turns away if others come too close (separation). No agent is designated as a leader; instead, the flock structure emerges spontaneously as each agent follows the same rules. Uri Wilensky's NetLogo Flocking model [26, 27] further popularized this framework, providing an accessible platform to experiment with how adjusting rule parameters affects group dynamics.

Traditional flocking models assume all agents use identical, fixed behavioral parameters throughout the simulation. However, real individuals (animals or people) adapt their behavior based on experience. An animal that survives a near-collision with a predator becomes more sensitive [10]; a pedestrian who nearly bumps into someone adjusts their personal space [17]. Learning and adaptation are fundamental to biological self-preservation [14].

Recent work has examined parameter sensitivity in flocking systems [7] and adaptive responses to environmental changes [18], but most standard boid agents lack cognitive flexibility to regulate risk. This raises an important question: How do individual adaptive responses affect collective dynamics? When agents can modify their own safety thresholds based on experience, do new patterns emerge? Does the population develop heterogeneous risk tolerances, even when starting from identical conditions?

We introduce an adaptive self-preservation mechanism where each agent maintains a personal collision threshold T_i that evolves based on experience. When an agent detects a neighbor within its threshold distance (a "near-miss"), it increases its threshold; when safe, the threshold slowly decays. This is formalized as a memoryless feedback rule: $T_{t+1} = T_t + \alpha \cdot I_{\text{near-miss}} - \beta \cdot (1 - I_{\text{near-miss}})$, where α is the adaptation rate, β is the decay rate, and $I_{\text{near-miss}} \in \{0, 1\}$ is a binary indicator function. The asymmetry ($\alpha \gg \beta$) models the biological principle of "once bitten, twice shy'", where agents learn quickly from danger, but forget slowly. This approach differs fundamentally from reinforcement learning. Firstly, there is no reward engineering. Agents do not optimize task performance. Secondly, there is no training phase and immediate online adaptation from time step 1. Third, threshold updates require minimal per-agent computation, which is one addition and two comparisons per step, once neighbor distances are available. What is more, it offers biological plausibility, which matches stimulus-response patterns in nature.

We address three research questions: **RQ1:** Does the adaptive mechanism produce emergent heterogeneity in risk tolerance? **RQ2:** How do adaptation parameters (α, β) influence collective behavior? And **RQ3:** Is the mechanism robust across different population densities?

We present four main contributions. We introduce a memoryless threshold adaptation mechanism for flocking agents and provide empirical evidence that it generates emergent heterogeneity from homogeneous rules. Furthermore, we conduct a rigorous parameter sensitivity analysis to identify optimal ratios α/β and validate the robustness of the system under varying environmental conditions.

The remainder of this paper is organized as follows: Section 2 reviews related work; Section 3 describes the model; Section 4 presents the experimental design; Section 5 reports results; Section 6 discusses implications; and Section 7 concludes.

2 Related Work

2.1 Classical Flocking and Collective Motion Models

Reynolds' Boids model [19] established the foundation for computational flocking, demonstrating that three local rules (alignment, cohesion, separation) are sufficient to produce realistic collective patterns. Vicsek et al. [25] introduced a minimal model showing phase transitions between ordered and disordered collective motion. Couzin et al. [3, 4] extended this work by defining interaction zones and studying how leadership emerges in complex moving groups. Wilensky's NetLogo implementation [26, 27] further popularized these models, making them accessible for interdisciplinary research.

Subsequent work has challenged some homogeneity assumptions of these classical models. Ballerini et al. [1] showed that flocking birds use topological interactions (nearest neighbors) rather than metric (distance-based). Similarly, Ling et al. [15] demonstrated that individual differences in social relationships significantly affect collective patterns, suggesting that identical agent assumptions may oversimplify real systems.

2.2 Parameter Sensitivity and Adaptive Dynamics

Extensive research examines how parameter variations affect flocking stability and emergence. Farkas and Wang [7] conducted systematic sensitivity analysis of flocking parameters including speed, distance thresholds, noise, and delays, showing that small parameter changes can induce qualitative transitions in collective behavior. This establishes a framework for understanding how threshold variations influence system-level dynamics.

Also, adaptation is observed in nature. Papadopoulou et al. [18] studied how predation pressure affects diffusion rates during collective turns in bird flocks, revealing that environmental threats change individual sensitivity to neighbors. Their findings suggest that adaptive risk responses are not merely theoretical constructs but observable phenomena. Recent studies on collective decision-making further support the importance of individual variation. Sridhar et al. [23] demonstrated that the geometry of individual decision-making shapes collective outcomes, while Harpaz et al. [11] showed that discrete modes of information processing create behavioral heterogeneity even among genetically identical fish. Our work extends this insight by modeling how such adaptive sensitivity might emerge through simple, memoryless feedback rules.

2.3 Optimization-Based Collision Avoidance

Collision avoidance is a main challenge in practical multi-agent systems. In robotics, approaches range from artificial potential fields to model predictive control [21]. These methods typically assume fixed safety parameters.

To address flexibility, recent work has explored sophisticated learning algorithms. A fundamental question is whether to use fixed rules or learned policies. Recent advances in deep Multi-Agent Reinforcement Learning (MARL) [12] enable agents to discover complex coordination strategies. Similarly, Martinez-Gil et al. [16] used MARL to learn collision avoidance policies for pedestrian models, demonstrating that learned behaviors can outperform heuristics in complex scenarios. Ebert et al. [6] demonstrated collective Bayesian decision-making in robot swarms.

However, these optimization-based approaches come with significant costs: computational overhead, extensive training data requirements, and the need for careful reward engineering. This limits their scalability in lightweight, decentralized swarms where agents lack high-end processors. In contrast, our threshold update mechanism requires only constant-time operations per agent (one addition, two comparisons) beyond the neighbor detection that any spatial model must perform, making it suitable for resource-constrained platforms.

2.4 Self-Preservation and Cognitive Heuristics

In contrast to heavy optimization, research on biological systems suggests that organisms often rely on simple, fast cognitive heuristics [9]. Moussaïd et al. [17] showed that pedestrians follow rules like "stop if another step leads to a collision" rather than computing optimal paths. Seitz et al. [22] demonstrated how such heuristics explain spatial movement patterns.

Our work occupies a middle ground: more adaptive than fixed-parameter models but simpler than a full RL. By using a minimal feedback heuristic, we aim to capture essential adaptive behaviors while maintaining interpretability and biological plausibility.

Unlike classical models [19, 25], our agents adapt their parameters online. Unlike Reinforcement Learning (RL) approaches [12], our mechanism does not require training or reward engineering. Unlike fixed sensitivity analyzes [7], our model allows the sensitivity itself to evolve. We bridge the gap between static flocking models and complex learning systems, offering a solution for safety-critical applications.

3 Model

3.1 Base Model: NetLogo Flocking

We built our simulation using NetLogo, a platform well-suited for agent-based modeling of complex systems. Our starting point is the NetLogo library's Flocking model, which implements Reynolds' boid rules [26]. This model operates

on a 2D toroidal grid with N agents (wrapping around edges) and advances in discrete time steps ("ticks"). Each agent (bird) updates its heading each tick based on local perceptions within a certain radius. By default, the flocking model includes three core behaviors:

- **Alignment:** a bird tends to turn so that it is moving in the same direction as nearby birds (to align the direction with the group).
- **Cohesion:** move toward the average position of nearby birds (to stay with the group).
- **Separation:** a bird will turn to avoid another bird that gets too close (to avoid crowding or collisions).

The separation rule takes priority: when triggered, it overrides alignment and cohesion. Each rule has a maximum turn rate that limits how quickly agents can adjust their heading per time step.

3.2 Adaptive Self-Preservation Mechanism

We extend the base flocking model by giving each agent i a dynamic personal collision threshold $T_i(t) \in [T_{\min}, T_{\max}]$. This threshold determines the agent's risk sensitivity: a near-miss is detected when the distance to the nearest neighbor falls below $T_i(t)$.

Memoryless Feedback Control: To regulate this threshold, we adopt a linear adjustment rule based on two design principles. First, the update should be memoryless, which means that only the current state matters, not the accumulation of past experiences. This reflects biological constraints: organisms often react to immediate stimulation rather than keeping detailed records of the past [9]. Second, the rule should implement asymmetric learning, rapid increase upon threat detection but gradual decay during safety, consistent with a risk-averse strategy [14]. The update rule is formalized as:

$$T_i(t+1) = T_i(t) + \alpha \cdot I_{\text{near-miss}} - \beta \cdot (1 - I_{\text{near-miss}}) \tag{1}$$

subject to $T_{\min} \leq T_i(t+1) \leq T_{\max}$, where:

- State**:** Only current threshold $T_i(t)$ persists (memoryless property)
- Stimulus**:** Binary near-miss detection $I_{\text{near-miss}} \in \{0, 1\}$
- Response**:** Increase by α if danger, decrease by β if safe

The mechanism is memoryless in the Markovian sense: $T_i(t+1)$ depends only on the current state $T_i(t)$ and current stimulus $I_{\text{near-miss}}$, requiring no history buffer of past encounters. While the threshold $T_i(t)$ itself represents accumulated experience, the update rule does not explicitly track or weight historical events. This differs from cumulative learning models (e.g., Bayesian updating) where agents integrate information across multiple timesteps. Note that the overall system complexity remains $O(N^2)$ per timestep due to nearest-neighbor calculations inherent to flocking models, but the threshold update itself is constant-time once neighbor distances are computed.

Boundary Constraints: The threshold bounds are calibrated to biological and physical constraints imposed by the base flocking model. Here, *vision* denotes each agent's perception radius, i.e. the maximum distance at which it can sense neighbors, fixed at 5 patches throughout. The lower bound $T_{\min} = 0.5$ patches (10% of the vision radius, 50% of the minimum-separation distance) represents extreme risk tolerance, ensuring agents retain basic collision avoidance. The upper bound $T_{\max} = 2.0$ patches (40% of the vision radius, 200% of minimum-separation), set in accordance with Weber's Law [8], prevents agents from reacting to ecologically irrelevant distant neighbors. The two bounds thus span a factor of four ($T_{\max}/T_{\min} = 4.0$), a range that mirrors the behavioral variability observed in animal groups [5].

Adaptation Rate (α) and Decay Rate (β): We set $\alpha = 0.05$ (baseline) based on animal learning studies. Griffin et al. [10] found birds develop lasting predator vigilance after 5–10 dangerous encounters. With $\alpha = 0.05$, agents require approximately 20 near-misses to meet the full threshold range from $T_{\min}$ to $T_{\max}$ (since $\Delta T = T_{\max} - T_{\min} = 1.5$, requiring $1.5/0.05 = 30$ increments, or $\sim$20 near-misses accounting for boundary effects).

We set $\beta = 0.01$ (baseline), yielding $\alpha/\beta = 5$. This specific ratio is grounded in Loss Aversion and Prospect Theory. Kahneman & Tversky's [13] Prospect Theory demonstrates that organisms weight losses approximately 2–3$\times$ more heavily than equivalent gains. In our context, a "loss" represents a near-miss, while a "gain" represents safety. The evolutionary theory suggests that strategies favoring false positives (over caution) over false negatives (missing a threat) typically use ratios ≥ 2. Our choice of $\alpha/\beta = 5$ is conservative but consistent with a risk-averse strategy [14]. We treat this value only as a baseline: Section 5.2 systematically sweeps α/β across the full range from near-symmetric (≈ 1) to strongly asymmetric (≈ 20) to map how the ratio shapes the collective outcome.

Technically, this asymmetry ($\alpha \gg \beta$) implements the "once bitten, twice shy" behavioral principle. With $\alpha = 0.05$ and $\beta = 0.01$, a single near-miss increases the threshold by 5%, but requires 5 consecutive safe time steps to decay by the same amount. This creates a threshold with ratchet effect: it is easy to become cautious but hard to relax, which is exactly the pattern observed in animal risk assessment under predation pressure [2].

3.3 Visual Feedback and Key Metrics

To visualize emergent heterogeneity, we use a dynamic coloring system. Agents in a near-miss state turn red, providing immediate feedback on system risk. In adaptive mode, safe agents' colors indicate their threshold level: darker colors represent higher thresholds (more cautious), lighter colors represent lower thresholds (more risk-tolerant).

We measure four outcomes:

- Near-miss rate: Percentage of agents currently in near-miss state.
- Mean threshold: Average T_i across the population.

- Threshold SD: Standard deviation of T_i, which is a proxy for behavioral heterogeneity.
- Cohesion index: Average number of neighbors per agent with vision radius.

4 Experimental Design

4.1 Overview

Table 1. Experimental Design Summary

Experiment	Purpose	Design	Runs
Exp 1	Fixed vs Adaptive	2×3 factorial	180
Exp 2	Parameter sensitivity	3×3 factorial	270
Exp 3	Density robustness	3 levels	90
Total			540

We conducted three experiments to address our research questions, totaling 540 simulation runs with 30 replications per condition and 500 time steps each. Table 1 summarizes the experimental framework.

4.2 Experiment 1: Fixed vs Adaptive Comparison

We compared fixed threshold mode (agents use constant T) with adaptive mode (agents update T_i via Equation 1) to evaluate whether the adaptive mechanism offers advantages over fixed parameter settings. We utilized a 2×3 factorial design crossing mode (Fixed vs. Adaptive) with initial threshold $T_{\text{init}} \in \{0.75, 1.0, 1.5\}$. These initial values represent risk-tolerant, neutral, and risk-averse baselines relative to the model's default minimum separation (1.0). The primary objective is to test for convergence: specifically, whether adaptive agents self-organize to a common safety equilibrium regardless of their starting conditions, which is a key standard for robust decentralized systems.We hypothesized that the adaptive mode produces emergent heterogeneity (threshold SD > 0), unlike fixed mode.

4.3 Experiment 2: Parameter Sensitivity

To understand the impact of learning dynamics, we varied the adaptation rate $\alpha \in \{0.02, 0.05, 0.10\}$ and the decay rate $\beta \in \{0.005, 0.01, 0.02\}$ in a full factorial design. This explores three different asymmetry ratios: $\alpha/\beta \approx 1$ –2 (near-symmetric that test necessity of asymmetry), $\alpha/\beta \approx 5$ (moderate asymmetry which is the baseline and match loss-aversion literature), and $\alpha/\beta \approx$ 10–20 (strong asymmetry, extreme conservatism). This design allows us to isolate the interaction effects between learning and decay rates on system stability.

4.4 Experiment 3: Density Robustness

To assess scalability, we tested population levels of $N \in \{100, 200, 400\}$ on a fixed-size world, while keeping the adaptation parameters constant ($\alpha = 0.05$, $\beta = 0.01$). Because the world dimensions are held constant, these levels correspond to increasing density: a sparse regime in which flocks are loosely coupled and collisions are rare, a medium regime approximating natural flock density [1], and a dense, crowded regime [17]. This four-fold increase in population (from $N = 100$ to $N = 400$) tests whether the adaptive mechanism remains robust as the system transitions from isolated movement to continuous interaction.

4.5 Control Parameters and Analysis Methods

All other parameters were held constant: vision = 5 patches, minimum separation = 1 patch, max-align-turn rate = 5°/step, max-cohere-turn rate = 3°/step, max-separate-turn rate = 1.5°/step, $T_{\min} = 0.5$, $T_{\max} = 2.0$.

Every agent starts each run at the same threshold T_{init} (not drawn from a distribution), so any heterogeneity reported below is emergent rather than seeded. Agents begin at random positions and headings with an independent random seed per replication; the world size is fixed across all runs, and Experiments 1–2 used $N = 300$.

We computed descriptive statistics (mean, SD) by condition. Statistical tests included independent samples t-tests (Exp 1), two-way ANOVA with interaction (Exp 2), and one-way ANOVA with Tukey HSD post-hoc tests (Exp 3). Effect sizes (Cohen's d, partial η^2) assessed practical significance. Because the parametric tests above assume normally distributed data, we assessed this assumption using the Shapiro–Wilk test. Most conditions were consistent with normality. Deviations arose only in a few isolated conditions in Experiment 1 and in the bounded outcomes of Experiment 2, where threshold values accumulate near $T_{\min}$ and $T_{\max}$. For these cases we re-ran the key comparisons using non-parametric alternatives, namely the Mann–Whitney U test for Experiment 1 and the Kruskal–Wallis test for Experiment 3, and both confirmed the same conclusions with strong statistical significance.

5 Results

5.1 Experiment 1: Adaptive Agents Develop Heterogeneity

We first examined whether identical agents could develop differentiated behaviors solely through interaction. Table 2 summarizes the comparison between fixed-parameter baselines and the adaptive mode,

In fixed mode, the standard deviation (SD) of thresholds was zero by definition. In adaptive mode, threshold SD ranged from 0.164 to 0.360 across initial conditions, indicating significant heterogeneity. A one-sample t-test confirmed that adaptive threshold SD was significantly greater than zero: $t(89) = 27.13$, $p = 8.18 \times 10^{-45}$ (Figure 1). This demonstrates that identical agents following

Table 2. Experiment 1 Results: Fixed vs Adaptive Comparison

Mode	Init T	Near-Miss (%)	Mean T	SD T	n
Fixed	0.75	22.8 ± 4.0	0.75	0.000	30
Fixed	1.00	41.3 ± 3.4	1.00	0.000	30
Fixed	1.50	87.4 ± 2.3	1.50	0.000	30
Adaptive	0.75	91.1 ± 2.2	1.89	0.360	30
Adaptive	1.00	93.9 ± 1.6	1.94	0.263	30
Adaptive	1.50	95.9 ± 1.3	1.98	0.164	30

identical rules develop differentiated risk tolerances purely through experiential feedback.

Independent samples t-tests comparing fixed vs adaptive modes at each initial threshold level revealed extremely large effect sizes: $T_{\text{init}} = 0.75$: $t = 81.74$, $p = 1.34 \times 10^{-61}$, Cohen's $d = 21.11$; $T_{\text{init}} = 1.0$: $t = 76.51$, $p = 6.01 \times 10^{-60}$, $d = 19.76$; $T_{\text{init}} = 1.5$: $t = 17.85$, $p = 3.07 \times 10^{-25}$, $d = 4.61$. These effect sizes confirm that the adaptive mechanism fundamentally transforms system behavior. Regardless of the initial threshold (0.75, 1.0, or 1.5), the adaptive agents

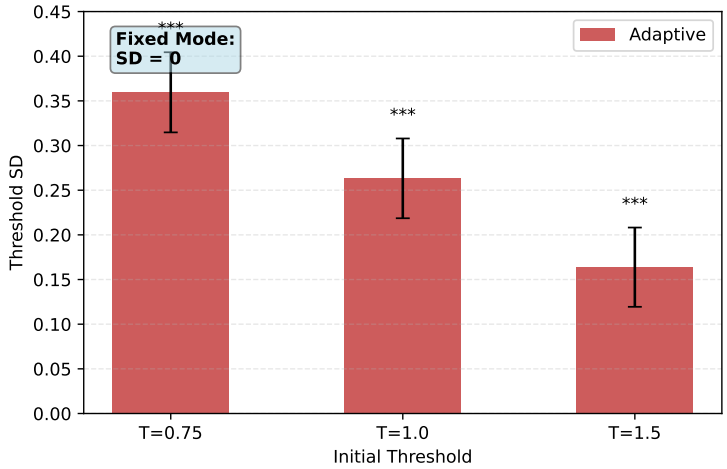

Fig. 1. Exp 1 Emergent Heterogeneity

converged to high mean thresholds (1.89–1.98), approaching $T_{\max} = 2.0$. This shows that even agents initialized with risk-tolerant thresholds ($T_{\text{init}} = 0.75$, well below minimum-separation) rapidly adapted to safe levels. Notably, the final heterogeneity depended on initial conditions: SD decreased monotonically from 0.360 at $T_{\text{init}} = 0.75$ to 0.263 at $T_{\text{init}} = 1.0$ and 0.164 at $T_{\text{init}} = 1.5$. This reveals that agents starting from lower initial thresholds experience longer learning

paths and accumulate more diverse near-miss histories, producing greater emergent variation in final risk tolerances. Agents that begin closer to $T_{\max}$ saturate quickly, leaving less room for differentiation.

Interestingly, the adaptive mode exhibited significantly higher near-miss rates (91.1–95.9%) compared to fixed modes (22.8–87.4%).This apparent paradox requires careful interpretation. In our model, a "near-miss" is defined as the detection of a neighbor within an agent's personal threshold $T_i(t)$, which triggers the separation behavior and threshold adjustment. Because adaptive agents converge to expanded safety zones (mean $T_i(t) \approx 1.89$–1.98 patches), nearly double the minimum-separation distance (1.0 patch), they detect and respond to neighbors at much greater distances than agents with smaller thresholds.

The high near-miss rates (91–96%) therefore reflect proactive monitoring activity rather than actual collision risk. Adaptive agents become alert and initiate avoidance maneuvers earlier than agents with smaller thresholds. While this produces frequent threshold violations (the technical definition of "near-miss" in our implementation), these are early warnings occurring at safe distances, not critical proximity events. Importantly, this metric conflates alertness (threshold violations at safe distances) with danger (actual close approaches below minimum-separation). Future work will implement dual metrics distinguishing safety zone activations from critical proximity events to separately quantify monitoring activity and collision risk (see Section 6).

5.2 Experiment 2: The α/β Ratio Matters

To understand the boundaries of this behavior, we conducted a full factorial sweep of the adaptation rate (α) and decay rate (β). Table 3 presents mean thresholds and heterogeneity (SD) across parameter combinations, visualized in Figures 2A and 2B. Notably, the effects of larger decay rates ($\beta > 0.01$) on heterogeneity patterns require further investigation.

Table 3. Experiment 2: Mean threshold and heterogeneity by parameters. Each cell reports the mean threshold with the threshold SD in parentheses.

	Decay Rate (β)		
α	0.005	0.010	0.020
0.02	1.75 (0.54)	0.67 (0.45)	0.51 (0.05)
0.05	1.99 (0.08)	1.94 (0.26)	1.26 (0.71)
0.10	1.99 (0.08)	1.99 (0.11)	1.96 (0.23)

The Two-way ANOVA revealed significant main effects of α ($F(2, 261) = 16312.69$, $p < .001$, $\eta_p^2 = 0.992$) and β ($F(2, 261) = 6750.73$, $p < .001$, $\eta_p^2 = 0.981$), and a significant interaction ($F(4, 261) = 2836.04$, $p < .001$, $\eta_p^2 = 0.978$). The large F-statistics and effect sizes indicate that the parameter choices dramatically influence the system outcomes.

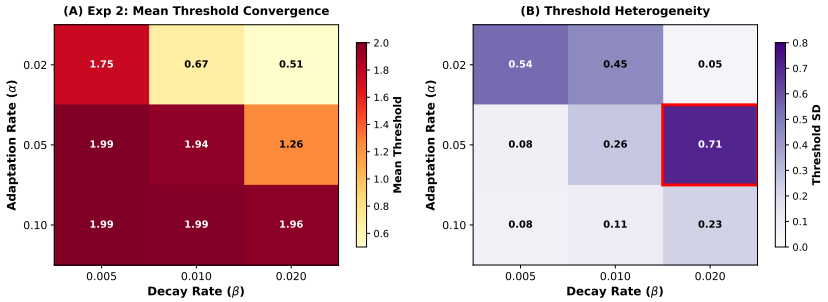

Fig. 2. Exp 2: Parameter Sensitivity Analysis

The magnitude of these effects suggests that the ratio α/β dictates the system's phase. We identify three distinct behavioral regimes as Fig 2A:

Regime 1: Strong Asymmetry ($\alpha \gg \beta$). When $\alpha/\beta \geq 10$ (e.g., $\alpha = 0.10$, $\beta = 0.005$ or 0.01; $\alpha = 0.05$, $\beta = 0.005$), thresholds converge rapidly to $T_{\max} = 2.0$ (mean = 1.99) with minimal heterogeneity (SD = 0.08–0.09). All agents become cautious, producing homogeneous risk-averse populations. This represents that fear dominates.

Regime 2: Moderate Asymmetry ($\alpha/\beta \approx 5$). When $\alpha/\beta = 5$ ($\alpha = 0.05$, $\beta = 0.01$), the baseline condition, thresholds converge to high values (mean = 1.94) with moderate heterogeneity (SD = 0.26). This balances collective safety with individual variation.

Regime 3: Weak/ Reverse Asymmetry ($\alpha \leq \beta$). When $\alpha/\beta \leq 2.5$ (e.g. $\alpha = 0.02$ with $\beta = 0.01$ or 0.02), the thresholds decay to $T_{\min}$ (mean = 0.51–0.67), which represents dangerous under sensitivity. We treat this collapse toward $T_{\min}$, i.e. the loss of the protective high-threshold equilibrium, as the failure mode of the mechanism.

Figure 2B reveals that maximum heterogeneity (SD = 0.71, highlighted by red box) emerges when $\alpha = 0.05$ and $\beta = 0.02$ ($\alpha/\beta = 2.5$). This condition produces a threshold distribution that some agents become cautious (high T_i), others remain risk-tolerant (low T_i), depending on their early near-miss experiences. This shows that initial random fluctuations amplify into lasting behavioral differences. The key finding from this experiment is that the α/β ratio, not absolute values, determines emergent behavior. We report the absolute SD of T_i rather than a coefficient of variation, since thresholds share a common bounded scale and the SD measures behavioral spread directly.

5.3 Experiment 3: Robust Across Densities

Finally, we evaluated the system's performance across varying population densities ($N \in \{100, 200, 400\}$). Results are summarized in Table 4, visualized as Figure 3.

Table 4. Experiment 3: Density Robustness Results

Population	Near-Miss (%)	Mean T	SD T
100	83.7 ± 4.8	1.76	0.51
200	91.2 ± 2.7	1.89	0.36
400	95.8 ± 1.0	1.97	0.19

One-way ANOVA showed a significant effect of population density on near-miss rates: $F(2, 87) = 107.62$, $p = 2.97 \times 10^{-24}$, $\eta^2 = 0.712$ (large effect). Post-hoc Tukey HSD tests confirmed that all pairwise comparisons were significant ($p < .001$): 100 vs 200 (mean difference = 7.45%), 100 vs 400 (12.07%), 200 vs 400 (4.62%). Although the adaptive mechanism functioned robustly throughout

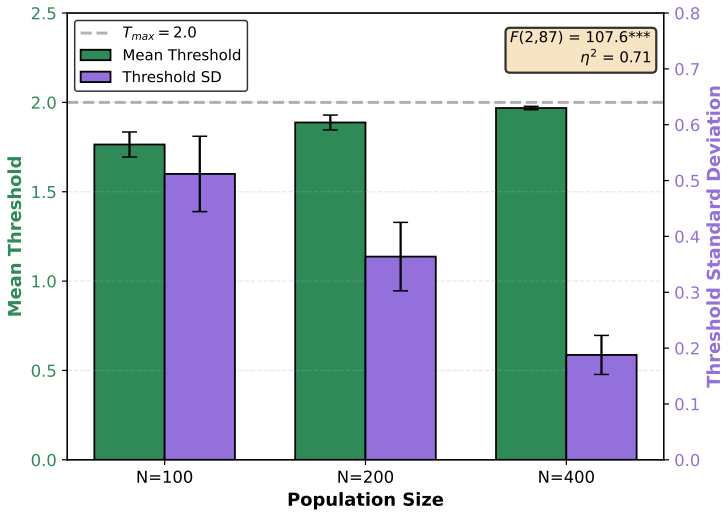

Fig. 3. Exp 3: Density Robustness

the 4-fold density range, the specific outcomes exhibited density dependence. Furthermore, the system shows the adaptive constraint. As shown in Fig.3, when density increases from 100 to 400 agents, the mean threshold automatically rises ($1.76 \rightarrow 1.97$), while heterogeneity declines ($SD : 0.51 \rightarrow 0.19$).

This indicates a functional trade-off: in sparse environments ($N = 100$), lower interaction frequencies allow agents to maintain diverse, relaxed thresholds. In dense environments ($N = 400$), continuous near-miss pressure forces the entire population towards a strict and uniform high-safety standard. This adaptive mechanism enables the group to have scalability and can automatically tighten safety measures in congested situations.

6 Discussion

All agents begin with identical rules and parameters, yet develop heterogeneous risk tolerances through experience. This shows how individual variation can arise without predisposition, which is a key principle in complexity science. Differentiation emerges from the interaction between identical agents and their stochastic environments. Our findings align with recent observations that behavioral heterogeneity can emerge even among genetically identical individuals [11] and that individual differences shape collective outcomes [15, 23]. However, previous work typically assumes pre-existing heterogeneity, whereas our model demonstrates how it can emerge dynamically from experience.

The key to this emergent stability is the asymmetry between the adaptation and decay rates ($\alpha \gg \beta$). This models the biological observation that negative experiences (near-misses) have lasting impact, while positive experiences (safety) are quickly forgotten. The "once bitten, twice shy" principle appears to be a fundamental mechanism underlying self-preservation. The key role of the α/β ratio aligns with the findings of Farkas & Wang [7] that flocking systems exhibit nonlinear sensitivity to parameter combinations. Their work on fixed parameters complements our analysis of adaptive thresholds, together showing a comprehensive study of how sensitivity, whether static or dynamic, determines collective behavior.

Beyond these theoretical insights, the choice of a simple heuristic over more complex methods warrants justification from an engineering perspective. One might ask why we employ a simple heuristic rather than a sophisticated Reinforcement Learning (RL) approach. Although RL is powerful, heuristic adaptation offers distinct advantages for low-level self-preservation: it requires no offline training, offers deterministic interpretability, and achieves constant-time per-agent threshold updates (one addition and two comparisons per step, executed after neighbor detection). This computational efficiency mirrors the fast and frugal stimulus-response patterns observed in animal learning, making it highly operable on resource-constrained hardware.

Given this computational efficiency, our findings have implications for designing decentralized systems. In swarm robotics, designers can benefit from adaptive collision thresholds. Rather than pre-programming heterogeneity, designers can allow it to emerge through experience. The α/β ratio becomes a tunable safety parameter: higher ratios for high-risk missions, lower ratios for efficiency-focused operations. Similarly, for autonomous vehicles, the model allowing vehicles to adapt following distances based on near-miss history could

mitigate traffic jams, with our results suggesting that a moderate asymmetry ($\alpha/\beta \approx 5$) optimally balances safety and flow.

While our findings demonstrate emergent heterogeneity and suggest its potential functional value, several methodological limitations should be noted. Our near-miss metric conflates monitoring activity with collision risk, as discussed in Section 5.1; future work will implement dual metrics to separately quantify these phenomena.

Additionally, we lack formal convergence validation (stationarity tests, sensitivity to simulation length), baseline comparisons to optimally-tuned fixed thresholds or density-aware heuristics, and empirical calibration of our parameters ($\alpha = 0.05$, $\beta = 0.01$) to biological data. While our α/β ratio is grounded in loss aversion theory, direct validation would require fitting model outputs to longitudinal vigilance data from field studies. The current model's homogeneous 2D environment limits generalizability. Real flocking systems navigate complex landscapes with obstacles, predators, and resource gradients. Extensions to heterogeneous environments and three-dimensional movement would test whether the adaptive mechanism remains robust under more realistic conditions. Agents currently lack communication; while sharing threat information could accelerate collective adaptation, it also risks information cascades where false alarms propagate system-wide.

Finally, while we demonstrate that emergent heterogeneity exists (SD $\approx$ 0.16–0.36 depending on initial conditions), we have not rigorously tested whether this diversity provides functional advantages over enforced homogeneity or random heterogeneity. Comparative experiments would clarify whether the specific pattern of emergent diversity contributes to collective robustness.

Regarding computational complexity, while threshold updates require only constant-time operations per agent (one addition, two comparisons), the overall system remains $O(N^2)$ per timestep due to neighbor detection inherent to spatial models. This is not a limitation of our adaptive mechanism specifically but a general characteristic of agent-based flocking simulations.

These limitations collectively suggest future work spanning empirical validation (biological calibration, dual metrics), methodological (refinement (convergence analysis, baseline comparisons), theoretical development (analytical equilibrium characterization), and domain extension (3D environments, communication, heterogeneous landscapes). Addressing these gaps would transform our proof-of-concept demonstration into a validated framework for adaptive collective behavior in real-world applications.

7 Conclusion

This study addressed the internal limitation of traditional flocking models: the reliance on static behavioral parameters that do not adapt to dynamic environmental risks. We introduced a memoryless, asymmetric self-preservation mechanism that allows agents to autonomously regulate their collision thresholds based on local interaction history.

Our experiments yield three key insights. First, we demonstrated that heterogeneous risk tolerances need not be hard-coded; they emerge spontaneously from identical agents interacting with their environment, converging to a physics-driven equilibrium. Second, the asymmetry of adaptation (α/β) serves as the primary control lever, where a "fast-panic, slow-relaxation" strategy is important for system stability. Third, the mechanism demonstrates robustness across the conditions we tested. It is effective to extend the sudden safety behaviors observed in a dispersed crowd to a dense and highly pressured crowd, without the need to adjust the parameters.

Theoretically, this work bridges the gap between rigid fixed-parameter models and computationally expensive reinforcement learning systems. By mimicking biological heuristics, specifically the principle of computational parsimony, we achieve robust safety behaviors using only local binary stimuli and constant-time per agent decision rules. Within the single 2D flocking model studied here, this suggests that complex collective coordination need not always require complex individual cognition; whether the result generalizes beyond this setting remains to be tested.

Future research should extend this work to 3D environments and specific real-world domains like pedestrian systems, while also exploring limited inter-agent communication and heterogeneous initialization. Additionally, we can compare this heuristic approach with reinforcement learning models that would provide valuable insights into the trade-offs between computational simplicity and optimized performance.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Ballerini, M., Cabibbo, N., Candelier, R., et al.: Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study of starling flocks. PNAS **105**(4), 1232–1237 (2008)
2. Brown, J.S., Laundré, J.W., Gurung, M.: The ecology of fear: optimal foraging, game theory, and trophic interactions. Journal of Mammalogy **80**(2), 385–399 (1999)
3. Couzin, I.D., Krause, J., Franks, N.R., Levin, S.A.: Effective leadership and decision-making in animal groups on the move. Nature **433**(7025), 513–516 (2005)
4. Couzin, I.D., Krause, J., James, R., Ruxton, G.D., Franks, N.R.: Collective memory and spatial sorting in animal groups. Journal of Theoretical Biology **218**(1), 1–11 (2002)
5. Dill, L.M.: Distance-to-cover and the escape decisions of an african cichlid fish. Environmental Biology of Fishes **27**(2), 147–152 (1990)
6. Ebert, J.T., Gauci, M., Mallmann-Trenn, F., Nagpal, R.: Bayes bots: Collective bayesian decision-making in decentralized robot swarms. In: ICRA 2020. pp. 7186–7192 (2020)
7. Farkas, I.J., Wang, S.: Spatial flocking: Control by speed, distance, noise and delay. PLOS ONE **13**(5), e0191745 (2018)

8. Gescheider, G.A.: Psychophysics: The Fundamentals. Lawrence Erlbaum Associates, Mahwah, NJ, 3rd ed. edn. (1997)
9. Gigerenzer, G., Gaissmaier, W.: Heuristic decision making. Annual Review of Psychology **62**, 451–482 (2011)
10. Griffin, A.S., Evans, C.S., Blumstein, D.T.: Learning specificity in acquired predator recognition. Animal Behaviour **62**(3), 577–589 (2001)
11. Harpaz, R., Tkačik, G., Schneidman, E.: Discrete modes of social information processing predict individual behavior of fish in a group. PNAS **114**(38), 10149–10154 (2017)
12. Hüttenrauch, M., Adrian, S., Neumann, G., Hoogendoorn, S.P.: Deep reinforcement learning for swarm systems. Journal of Machine Learning Research **20**(54), 1–31 (2019)
13. Kahneman, D., Tversky, A.: Prospect theory: An analysis of decision under risk. Econometrica **47**(2), 263–291 (1979)
14. Lima, S.L., Dill, L.M.: Behavioral decisions made under the risk of predation: A review and prospectus. Canadian Journal of Zoology **68**(4), 619–640 (1990)
15. Ling, H., Mclvor, G.E., van der Vaart, K., Vaughan, R.T., Thornton, A., Ouellette, N.T.: Costs and benefits of social relationships in the collective motion of bird flocks. Nature Ecology & Evolution **3**(6), 943–948 (2019)
16. Martinez-Gil, F., Lozano, M., Fernández, F.: Emergent behaviors and scalability for multi-agent reinforcement learning-based pedestrian models. Simulation Modelling Practice and Theory **106**, 102147 (2021)
17. Moussaïd, M., Helbing, D., Theraulaz, G.: How simple rules determine pedestrian behavior and crowd disasters. PNAS **108**(17), 6884–6888 (2011)
18. Papadopoulou, M., Hildenbrandt, H., Hemelrijk, C.K.: Diffusion during collective turns in bird flocks under predation. Frontiers in Ecology and Evolution **11**, 1198248 (2023)
19. Reynolds, C.W.: Flocks, herds and schools: A distributed behavioral model. In: Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques. pp. 25–34 (1987)
20. Schelling, T.C.: Micromotives and Macrobehavior. W. W. Norton & Company, New York, revised edition edn. (2006)
21. Schwarting, W., Alonso-Mora, J., Rus, D.: Planning and decision-making for autonomous vehicles. Annual Review of Control, Robotics, and Autonomous Systems **1**, 187–210 (2018)
22. Seitz, M.J., Bode, N.W., Köster, G.: How cognitive heuristics can explain social interactions in spatial movement. Journal of the Royal Society Interface **13**(121), 20160439 (2016)
23. Sridhar, V.H., Li, L., Gorbonos, D., Nagy, M., Schell, B.R., Sorochkin, T., Couzin, I.D.: The geometry of decision-making in individuals and collectives. PNAS **118**(50), e2102157118 (2021)
24. Sumpter, D.J.T.: The principles of collective animal behaviour. Philosophical Transactions of the Royal Society B: Biological Sciences **361**(1465), 5–22 (2006)
25. Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., Shochet, O.: Novel type of phase transition in a system of self-driven particles. Physical Review Letters **75**(6), 1226 (1995)
26. Wilensky, U.: Netlogo flocking model. Center for Connected Learning and Computer-Based Modeling, Northwestern University (1998), http://ccl.northwestern.edu/netlogo/models/Flocking
27. Wilensky, U.: Netlogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University (1999), http://ccl.northwestern.edu/netlogo/