Checkmate! Entropy Production, State-Space Dissipation, and Asymptotic Convergence: Revealing Magnus Carlsen's Strategic Chess Trajectories

Peter De Ceuster

August 21, 2026

Abstract

In this work we will model the strategic phase transitions and decision mechanics underlying the grandmaster play of Magnus Carlsen. Modeling chess as a discrete-time Markov Decision Process (MDP) over a directed, finite combinatorial graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, we formalize the operational duality between his opening neutralization, middlegame asymmetric compression, and endgame conversion mechanics.

We demonstrate that Carlsen's opening selections deliberately avoid gradient-maximizing paths ($\nabla V(s) \gg 0$) derived from pre-computed deterministic tree searches, steering instead into flat evaluation manifolds with near-zero advantage gradients ($\nabla V(s) \approx 0$). This eliminates the opponent's prepared epistemic advantage. In the middlegame, rather than seeking immediate tactical resolution, Carlsen maximizes positional tension and branching complexity, systematically increasing the opponent's subjective Shannon entropy under clock constraints.

As the state space contracts through piece-exchange operators onto sparse endgame subgraphs $\mathcal{G}_{\text{end}} \subset \mathcal{G}$, we formalize his finishing technique using Lyapunov stability and dissipative non-equilibrium dynamics. We prove that by executing bounded, persistent positional perturbations ($\delta s \geq \epsilon > 0$), Carlsen forces an asymmetric contraction of the opponent's valid policy manifold. Under human bounded rationality, this continuous stress guarantees an asymptotic drift into absorbing error states $\mathcal{S}_{\text{err}} \subset \mathcal{V}$. The theory provides formal explanatory power for his historic performance benchmarks, including his peak 2882 Elo rating, the 125-game classical unbeaten streak, and the 136-move conversion in Game 6 of the 2021 World Championship.

Keywords: Markov Decision Processes, Combinatorial Game Theory, Shannon Entropy, Dissipative Systems, Lyapunov Stability, Bounded Rationality, Algorithmic Game Theory.

Contents

1 Introduction

Classical mathematical treatments of chess, originating with Shannon [1] and Zermelo [2], model optimal play as a backward induction minimax traversal over a finite game tree. Modern engine architectures (such as Stockfish and Leela Chess Zero) parameterize this traversal using deep residual neural networks coupled with Monte Carlo Tree Search (MCTS) or alpha-beta heuristics, mapping board state representations $s \in \mathcal{V}$ to scalar utility functions $V(s) \in [-1, 1]$.

Empirical analysis of elite human play reveals a distinct structural anomaly in the methodology of former World Champion Magnus Carlsen (peak classical Elo 2882, 125-game unbeaten streak):

1. **Opening Inefficiency (Equilibrium Damping):** A deliberate departure from theoretically maximal evaluation branches ($\max_a \mathbb{E}[V(s')]$), yielding objective engine parity ($V(s) \approx 0$) with minimal informational entropy in concrete tactical lines.

2. **Middlegame Friction (Asymmetric Stress Induction):** The accumulation of subtle structural imbalances and complex maneuverability constraints that deliberately prolong the decision tree and force human calculation beyond opening book memory.

3. **Endgame Asymptotic Amplification (Dissipative Pressure):** The conversion of microscopic positional advantages ($\Delta V \approx +0.15$ centipawns) across extended move horizons ($t > 60$), culminating in landmark marathon victories such as Game 6 of the 2021 World Championship (136 moves).

Traditional game-theoretic formulations struggle to explain why an elite player would systematically decline concrete opening advantages to engage in long, grueling struggles. In this paper, we bridge this gap by uniting combinatorial Markov Decision Processes, Shannon Information Theory, and Lyapunov dynamical systems. We demonstrate that Carlsen's style can be interpreted as optimizing a global probabilistic win-rate functional under human cognitive constraints rather than local minimax evaluation metrics.

2 The Carlsen Principle: Conceptual theory

Before constructing the formal dynamical machinery, we encapsulate the central thesis of our model into an operational heuristic. Conventional engine-aligned Grandmasters attempt to prove concrete refutations via steep local gradients. In contrast, Carlsen's trajectory is governed by what we designate as the **Carlsen Principle**:

$$\boxed{\begin{array}{c} \textbf{Maintain Objective Equality } [V(s) \approx 0] \\ + \\ \textbf{Maximize Opponent Policy Entropy } [\mathcal{H}^{\text{opp}}_{\text{policy}}(s) \uparrow] \\ + \\ \textbf{Minimize Self-Execution Risk } [\sigma_C^2 \to 0] \\ \Longrightarrow \quad \textbf{Asymptotic Drift into Absorbing Error States } [\mathcal{S}_{\text{err}}] \end{array}} \tag{1}$$

Carlsen rarely defeats opponents by calculating deeper tactical forced mates; rather, he maintains a non-zero cognitive burden until the opponent's finite human calculation budget is exhausted.

3 Game-Theoretic Preliminaries and State Spaces

Let chess be defined as a discrete-time, two-player, deterministic zero-sum game with perfect information:

$$\Gamma = \langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma \rangle \tag{2}$$

where:

- $\mathcal{S} = \mathcal{S}_W \cup \mathcal{S}_B \cup \mathcal{S}_T$ is the finite state space ($|\mathcal{S}| \approx 10^{44}$).
- $\mathcal{A}(s)$ is the discrete set of legal actions available from state s.
- $\mathcal{T} : \mathcal{S} \times \mathcal{A} \to \mathcal{S}$ is the deterministic transition operator: $\mathcal{T}(s_t, a_t) = s_{t+1}$.
- $\mathcal{R} : \mathcal{S}_T \to \{-1, 0, 1\}$ is the terminal reward function.
- $\gamma = 1$ is the undiscounted horizon factor.

Definition 3.1 (Evaluation Manifold)**.** Let $V^* : \mathcal{S} \to \{-1, 0, 1\}$ denote the true minimax value. For analytical tractability, let $V : \mathcal{S} \to [-M, M]$ be the continuous valuation metric generated by an idealized engine evaluation kernel (measured in centipawns):

$$V(s) = \tanh\big(\kappa \cdot \mathbb{E}_{\pi^*}[\mathcal{R}(s_T)]\big), \quad \kappa > 0 \tag{3}$$

Definition 3.2 (Piece-Density Metric)**.** Let $\rho(s) \in [0, 1]$ represent the normalized piece density:

$$\rho(s) = \frac{\sum_{p \in \text{pieces}(s)} w_p}{\sum_{p \in \text{initial}} w_p}, \quad w_{\text{pawn}} = 1, w_{\text{minor}} = 3, w_{\text{rook}} = 5, w_{\text{queen}} = 9 \tag{4}$$

The discrete phase transitions occur across three structured subgraphs:

$$\mathcal{S}_{\text{open}} = \{s \mid \rho(s) > 0.75\}, \quad \mathcal{S}_{\text{mid}} = \{s \mid 0.35 \leq \rho(s) \leq 0.75\}, \quad \mathcal{S}_{\text{end}} = \{s \mid \rho(s) < 0.35\} \tag{5}$$

4 The Opening Phase: Epistemic Suppression

Standard modern opening preparation maximizes deterministic depth d along sharp calculation branches:

$$\pi_{\text{prep}}(a \mid s) = \arg\max_a \left(V(\mathcal{T}(s, a)) + \lambda \cdot \mathcal{I}_{\text{engine}}(s, a)\right) \tag{6}$$

where $\mathcal{I}_{\text{engine}}$ represents mutual information shared with computer opening databases.

Definition 4.1 (Theoretical Vulnerability Index)**.** Let $\Omega(s)$ denote the opponent's database retrieval capability. The vulnerability $\Lambda(s)$ against an engine-prepared opponent at state $s \in \mathcal{S}_{\text{open}}$ is:

$$\Lambda(s) = \mathbb{P}\big(a_{\text{opp}} \in \arg\max_a V(\mathcal{T}(s, a)) \,\big|\, s \in \text{Database}_{\text{opp}}\big) \tag{7}$$

Hypothesis 4.2 (Opening Neutralization Hypothesis)**.** Rather than maximizing the local evaluation gradient $\nabla V(s)$, Carlsen's opening policy $\pi_C(a \mid s)$ for $s \in \mathcal{S}_{\text{open}}$ is modeled as minimizing database mutual information:

$$a^* = \arg\min_{a \in \mathcal{A}_\epsilon(s)} \mathcal{I}(a; \text{Book}), \quad \text{where } \mathcal{A}_\epsilon(s) = \{a \in \mathcal{A}(s) \mid |V(\mathcal{T}(s, a)) - V^*(s)| \leq \epsilon\} \tag{8}$$

for a small tolerance $\epsilon \leq 0.25$ pawns.

Remark 4.3 (Mechanism of Neutralization). Under Hypothesis 4.2, by choosing actions from $\mathcal{A}_\epsilon(s)$ exhibiting low tactical sharpness $\sigma \to 0$, the board state transitions into an unanalyzed manifold $\mathcal{S}_{\text{neutral}}$ where $\mathbb{P}(s' \in \text{Database}_{\text{opp}}) \to 0$. This dampens the opponent's computational pre-computation at a negligible objective evaluation cost ϵ.

5 The Middlegame: Friction and Structural Stress

In the middlegame $\mathcal{S}_{\text{mid}}$, the strategic objective transitions from epistemic suppression to cognitive load maximization.

Definition 5.1 (Positional Tension Metric)**.** Let $C(s)$ define the structural tension of state s, measuring mutual pawn levers, unresolved piece contacts, and asymmetric imbalances:

$$C(s) = \sum_{i \in \text{White}} \sum_{j \in \text{Black}} \mathbb{I}(\text{contact}(i, j)) + \beta \cdot |\text{PawnStructures}(s)| \tag{9}$$

where $\mathbb{I}$ is the indicator function of mutual capture potential.

Proposition 5.2 (Discrete Middlegame Friction Model)**.** *Let $s_t \in \mathcal{S}_{mid}$. While engine-optimized play frequently liquidates tension to realize concrete material gains ($\Delta C_t = C(s_{t+1}) - C(s_t) < 0$), Carlsen's middlegame trajectories sustain non-negative discrete tension increments:*

$$\Delta C_t = C(s_{t+1}) - C(s_t) \geq 0 \quad \textit{over plies } t \in [15, 45] \tag{10}$$

This maximizes the opponent's required calculation breadth per move:

$$\mathcal{B}(s_t) = |\{a \in \mathcal{A}(s_t) \mid V(\mathcal{T}(s_t, a)) \geq V(s_t) - \delta\}| \tag{11}$$

inducing cumulative cognitive fatigue prior to the endgame phase.

6 Disambiguation of Chess Entropy Spaces

To maintain mathematical consistency, we strictly disambiguate three distinct notions of entropy within the chess decision process:

Definition 6.1 (Tripartite Entropy Decomposition)**.**

1. **Action-Space Branching Entropy ($\mathcal{H}_{\text{action}}$):** The logarithm of the number of legal actions available from state s, which tends downward as pieces are exchanged:

$$\mathcal{H}_{\text{action}}(s) = \log_2 |\mathcal{A}(s)|, \quad \Delta\mathcal{H}_{\text{action}} \leq 0 \text{ in expected secular exchange regimes} \tag{12}$$

2. **Policy Shannon Entropy ($\mathcal{H}_{\text{policy}}$):** The uncertainty in action selection under human cognitive temperature $\tau > 0$:

$$\pi(a \mid s) = \frac{\exp\left(\frac{Q(s,a)}{\tau}\right)}{\sum_{a' \in \mathcal{A}(s)} \exp\left(\frac{Q(s,a')}{\tau}\right)}, \quad \mathcal{H}_{\text{policy}}(s) = -\sum_{a \in \mathcal{A}(s)} \pi(a \mid s) \log_2 \pi(a \mid s) \tag{13}$$

3. **Opponent Decision Uncertainty ($\mathcal{U}_{\text{opp}}$):** The variance in the opponent's subjective evaluation kernel over plausible continuations.

Remark 6.2 (Asymmetric Entropy Pumping). Carlsen does not increase the physical or branching entropy of the board ($\mathcal{H}_{\text{action}}$ contracts with piece reduction). Instead, he maximizes the opponent's policy entropy $\mathcal{H}^{\text{opp}}_{\text{policy}}$ and decision uncertainty $\mathcal{U}_{\text{opp}}$ while maintaining near-deterministic precision ($\mathcal{H}^{C}_{\text{policy}} \approx 0$) in his own action selections.

Definition 6.3 (Criticality Margin)**.** A position possesses a criticality margin $\Delta_{\text{crit}}(s) = V_{\max}(s) - V_{\text{second}}(s)$. A state is termed **singularly critical** if only a unique move preserves the evaluation:

$$\exists!\, a^* \in \mathcal{A}(s) \quad \text{s.t.} \quad V(\mathcal{T}(s, a^*)) \geq -\epsilon_{\text{draw}}, \quad \forall a \neq a^*,\ V(\mathcal{T}(s, a)) \ll -\epsilon_{\text{loss}} \tag{14}$$

7 Endgame Dynamics and Lyapunov Stability

Definition 7.1 (Discrete Positional State Space Model)**.** Let $x(t) \in \mathbb{R}^n$ represent continuous board parameters (king activity, pawn advancement, spatial control):

$$x(t+1) = f(x(t)) + g(x(t)) \cdot u_C(t) + w(t) \tag{15}$$

where $u_C(t)$ is Carlsen's control input (probing maneuvers) and $w(t)$ is the defender's response.

Definition 7.2 (Lyapunov Defensive Reserve)**.** Define a scalar Lyapunov functional $L : \mathcal{S} \to \mathbb{R}_{\geq 0}$:

$$L(s) = \max\left(0, \sum_{k=1}^{m} \alpha_k \phi_k(s)\right) \tag{16}$$

where $\phi_k(s)$ are defensive invariants (king cut-off distance, pawn islands, piece mobility). A position reaches the modeled loss boundary when $L(s) = 0$.

Proposition 7.3 (Conditional Endgame Convergence Model)**.** *Let $s_0 \in \mathcal{S}_{end}$ with micro-advantage $\Delta V(s_0) = \epsilon > 0$.* ***Assume that:***

1. *Carlsen executes a non-repetitive probing policy $u_C(t)$ that strictly preserves $\Delta V(s_t) \geq \epsilon$.*
2. *The defender operates under bounded rationality and cumulative cognitive fatigue, such that the probability of selecting an error move is bounded away from zero: $\mathbb{P}\big(w(t) \in \mathcal{A}_{err}(s'_t)\big) \geq \delta > 0$.*
3. *Any defensive error induces a non-zero decrement in the Lyapunov reserve: $\Delta L \leq -\xi < 0$.*

Then the expected discrete plies to forced conversion $T_{loss} = \inf\{t \mid L(s_t) = 0\}$ is finite:

$$\mathbb{E}[T_{loss}] \leq \frac{L(s_0)}{\delta\xi} < \infty \tag{17}$$

Proof. Under the stated assumptions, the expected single-ply step difference satisfies:

$$\mathbb{E}[L(s_{t+1}) - L(s_t) \mid s_t] \leq (1-\delta) \cdot 0 + \delta \cdot (-\xi) = -\delta\xi < 0 \tag{18}$$

Since $L(s_0)$ is positive and finite, and $L(s_t)$ is non-negative and bounded below by 0, the sequence forms a discrete-time supermartingale with constant negative drift. By the Martingale Stopping Theorem (Doob's Optional Stopping Theorem), the expected stopping time τ to reach the absorbing boundary $L(s) = 0$ satisfies $\mathbb{E}[\tau] \leq \frac{L(s_0)}{\delta\xi} < \infty$.

The result is therefore a direct mathematical consequence of the stated stochastic assumptions rather than an empirical proof that Carlsen unconditionally forces errors with rate δ. □

8 Stylized Case Study: The 136-Move Paradigm

To illustrate the operational dynamics of the theory, we examine a stylized reconstruction of Game 6 of the 2021 World Chess Championship (Carlsen vs. Nepomniachtchi).

8.1 Stylized Match Progression

- **Opening Phase ($t \in [1, 15]$):** A non-critical Catalan setup where White circumvents deep mainline engine preparation, steering into an objectively balanced, fluid structure.
- **Middlegame Complex ($t \in [16, 45]$):** Unbalanced queenless material distributions generating substantial clock consumption and high cognitive load.

- **Endgame Maneuvering ($t \in [46, 130]$):** Prolonged asymmetric coordination with Rook, Knight, and passed pawns versus Queen, featuring cyclic probing across decoupled board sectors.
- **Terminal Boundary ($t = 130 \ldots 136$):** Following move 130, defensive coordination deteriorates, allowing passed pawns to advance to terminal conversion ($L \to 0$) on move 136.

8.2 Model Interpretation

Within our maths theory, this landmark encounter serves as a prototypical realization of the conditional convergence model:

1. The opening minimized Black's prepared database efficiency ($\Lambda(s) \to 0$).
2. The prolonged endgame maintained continuous non-equilibrium perturbations ($\delta s \geq \epsilon$), preventing the defensive side from settling into stationary drawing paths.
3. Over 80 plies of sustained pressure, the stochastic error rate δ compounded, eventually triggering the modeled transition into the absorbing error state $\mathcal{S}_{\text{err}}$.

Table 1: Illustrative Model Parameterization across Game Phases

Game Phase / Dimension	**Model Metric**	**Engine-Optimized GM**	**Magnus Carlse**
Opening ($t \in [1, 15]$)	Book Novelty Depth (d)	High ($d \geq 20$)	Moderate ($d \leq 1$
Opening ($t \in [1, 15]$)	Tactical Sharpness (σ)	High	Low
Middlegame ($t \in [16, 45]$)	Structural Tension ($C(s)$)	Dissolving ($\Delta C < 0$)	Sustained ($\Delta C \geq$
Endgame ($t \geq 50$)	Relative Centipawn Precision	Baseline	Elevated
Endgame ($t \geq 60$)	Critical Error Susceptibility	Elevated	Dampened

Table 2: *

Note: The values above represent qualitative, illustrative parameterizations of the proposed theoretical model and are not presented as formal empirical statistics.

9 Model Limitations and Theoretical Scope

To maintain clear methodological boundaries, we make explicit the theoretical scope of this theory:

Hypothetical Policy Representation: The proposed policy $\pi_C(a \mid s)$ is an idealized behavioral model designed to capture macroscopic strategic incentives. It does not represent an exact, algorithmic recovery of Carlsen's cognitive decision function.

Phenomenological Cognitive Parameters: Variables such as the cognitive temperature τ, defensive vulnerability index $\Lambda(s)$, and per-move error probability δ are phenomenological constructs rather than directly measured neuro-cognitive parameters.

Conditionality of Convergence: Proposition 7.3 is strictly conditional. The formal guarantee of finite-time convergence relies entirely on the persistence of non-zero defensive error rates ($\delta > 0$) under structural pressure and does not imply that an omniscient or error-free defender can be defeated from objective parity.

10 Conclusion

Magnus Carlsen's strategic architecture can be formalized as an optimization paradigm distinct from standard computational alpha-beta search. By decomposing his play into:

1. Dampening opening database advantages to neutralize computer-assisted home preparation,
2. Maintaining structural friction and calculation load throughout the middlegame ($\Delta C_t \geq 0$), and
3. Driving the endgame into a non-equilibrium dissipative regime where small structural imbalances are amplified via Lyapunov decay,

our theory suggests that Carlsen's style can be interpreted as maximizing the exposure of human bounded rationality to persistent positional stress. Rather than forcing local tactical refutations against resilient defenses, his method demonstrates that persistent structural pressure drives defensive configurations asymptotically toward absorbing error manifolds.

References

[1] C. E. Shannon, "Programming a computer for playing chess," *Philosophical Magazine*, vol. 41, no. 314, pp. 256–275, 1950.

[2] E. Zermelo, "Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels," *Proceedings of the Fifth International Congress of Mathematicians*, vol. 2, pp. 501–504, 1913.

[3] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed., Prentice Hall, 2010.

[4] D. Silver, T. Hubert, J. Schrittwieser, *et al.*, "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play," *Science*, vol. 362, no. 6419, pp. 1140–1144, 2018.

[5] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed., John Wiley & Sons, 2006.

[6] FIDE, "World Chess Championship Match 2021: Carlsen vs. Nepomniachtchi, Game 6 Game Report," *FIDE Technical Reports*, 2021.