

Haar-CFR: Solving Imperfect-Information Games with Applications in Data Analytics and Finance

Florian Fey

`feyf@uni-muenster.de`

University of Münster

Sergei Gorlatch

University of Münster

Research Article

Keywords: Data Analytics and Machine Learning, Optimization, Modeling in Economics, Application of Soft Computing, Decision Making, Data Science, Soft Computing Technologies, Counterfactual Regret Minimization

Posted Date: October 19th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10041026/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Haar-CFR: Solving Imperfect-Information Games with Applications in Data Analytics and Finance

Florian Fey^{1*} and Sergei Gorlatch¹

^{1*} University of Muenster, Germany.

*Corresponding author(s). E-mail(s): feyf@uni-muenster.de;
Contributing authors: gorlatch@wwu.de;

Abstract

Solving large-scale imperfect-information games is a fundamental challenge in data analytics and machine learning in general, and the application of soft computing technologies in decision making in particular. The state-of-the-art optimization approaches from game theory, such as Counterfactual Regret Minimization (CFR), provide robust equilibrium guarantees, but they are memory-intensive and often struggle with the fine-grained discretizations found in realistic environments. The alternative approach based on deep reinforcement learning offers more flexibility in high-dimensional state spaces, but it comes at the cost of scalability, stability, and interpretability.

We propose Haar-CFR—a novel hybrid soft computing framework that integrates discrete Haar wavelet regret regularization with CFR+, an advanced CFR variant that accelerates convergence by strictly accumulating only positive regrets and utilizing alternating updates. Using the regret signals structure over ordered domains, we apply adaptive spectral thresholding to replace dense tabular storage with compact spectral representations.

We demonstrate our approach on large-scale, imperfect-information application benchmarks, including standard Poker variants and a high-dimensional environment for financial market making. We achieve compression ratios above 99.9% (up to 5×10^4 memory reduction) and an accuracy comparable to standard CFR+. Compared to Deep CFR, we achieve a $6.5\times$ wall-clock speedup by avoiding neural inference overhead. In the financial domain, our approach successfully recovers classical inventory- and information-driven pricing as boundary conditions and exhibits stable nonlinear behavior in intermediate market conditions, establishing a robust framework for modeling in economics and data science.

In summary, our contributions establish wavelet-regularized regret minimization as a novel, broadly applicable, scalable, and interpretable alternative to tabular and neural equilibrium solvers for structured, imperfect-information games.

Keywords: Data Analytics and Machine Learning, Optimization, Modeling in Economics, Application of Soft Computing, Decision Making, Data Science, Soft Computing Technologies, Counterfactual Regret Minimization

1 Introduction

Solving large-scale imperfect-information games is a fundamental challenge in data analytics and machine learning in general, and the application of soft computing in particular. Game-theoretic learning, specifically *Counterfactual Regret Minimization (CFR)* (Zinkevich et al. 2007), offers a principled approach by guaranteeing convergence to an ϵ -Nash equilibrium. In modern practice, state-of-the-art solvers typically rely on *CFR+* (Tammelin et al. 2015), an advanced variant that accelerates convergence by strictly accumulating only positive regrets and utilizing alternating updates. However, despite this algorithmic speedup, the fine-grained discretization of realistic ordered domains results in a combinatorial growth of information sets. Therefore, applying classical tabular CFR (or CFR+) encounters a prohibitive scalability constraint known as the *Memory Wall* or State Space Explosion (Gholami et al. 2024).

While recent neural approaches like Deep CFR (Brown et al. 2019) attempt to mitigate this by approximating the regret surface, they introduce a secondary bottleneck – the *Time Wall*, i.e., the computational overhead of neural inference prevents deep lookahead search in real-time applications. Thus, the field faces a dilemma: choose between the intractable memory of tabular methods or the prohibitive latency of neural approximations.

Automated market making serves as a premier, real-world stress test for these algorithmic limitations. Financial markets represent a *Complex Adaptive System* (Cartea et al. 2015). In these high-frequency ecosystems, market makers (MMs) are algorithmic agents that provide liquidity by continuously quoting bid and ask prices, so other market participants can reliably sell and buy assets (Avellaneda and Stoikov 2008). Effective market making is not just a financial challenge: it is a computational challenge that requires frameworks that are scalable, robust, and strategically grounded (Spoonster et al. 2018).

At its core, a market maker faces a multi-objective optimization problem. To capture the bid–ask spread (buying slightly below fair value and selling slightly above fair value), the agent must balance two fundamental, often conflicting, sources of risk:

- *Inventory Risk:* The random exposure to adverse price movements caused by holding a non-zero net position (Ho and Stoll 1981).
- *Adverse Selection Risk:* The systematic losses arising from trading with better-informed counterparties who successfully anticipate future price movements (Glosten and Milgrom 1985).

Figure 1 illustrates the Market Maker’s Dilemma of balancing inventory and adverse selection risk within a unified decision framework, which remains an open challenge (Ganesh et al. 2019). While inventory pressures dictate lowering prices to




Figure 1: The Market Maker's Dilemma. Visualizing the theoretical conflict: Inventory Risk (Left) pressures the MM to lower prices to liquidate positions, while Adverse Selection Risk (Right) pressures the MM to widen spreads to defend against informed (toxic) flow. Balancing these opposing vectors requires a high-dimensional non-linear strategy.

liquidate open positions, adverse selection often dictates widening the spread to defend against toxic order flow. This creates a complex stochastic control problem where optimal policies must navigate a high-dimensional state space.

1.1 Theoretical Dilemma in Market Making

Historically, financial literature has addressed these risks through two largely separate theoretical paradigms.

The *stochastic control* framework, pioneered by [Avellaneda and Stoikov \(2008\)](#), formulates market making as an optimal control problem under inventory risk. Under assumptions such as Poisson order arrivals and Brownian price dynamics, these models yield closed-form analytical solutions where quotes are adjusted linearly based on inventory levels. While mathematically tractable and widely adopted, they treat order flow as exogenous noise and do not explicitly model the strategic information asymmetry ([Ho and Stoll 1981](#)).

In contrast, the *market microstructure* framework ([Glosten and Milgrom 1985](#); [Kyle 1985](#)) models trading as a Bayesian inference problem. Market makers update their beliefs about the asset's fundamental value based on observed order flow, embedding an information premium into the bid-ask spread. These models effectively capture adverse selection but often abstract away inventory constraints and risk aversion to volatility.

Although both paradigms are well-established, they remain computationally disjoint. Strategies derived from either paradigm in isolation are insufficient for fragmented real-world markets that exhibit both high volatility and significant toxic flow ([Spooner et al. 2018](#)). A unified, equilibrium-based formulation that simultaneously prices inventory and information risk is required but has historically been analytically intractable.

1.2 The Computational Barrier: Memory vs. Time

To unify both paradigms, recent research has turned to computational intelligence. Deep Reinforcement Learning (DRL) has been proposed to learn non-linear strategies that can adapt to both market regimes (Ganesh et al. 2019). However, DRL methods in adversarial finance suffer from critical limitations: simulator overfitting (Vyetrenko et al. 2020), training instability, and a lack of intrinsic interpretability, which poses a critical barrier to practical applicability (Rudin 2019).

1.3 Our Contribution: A Hybrid Soft Computing Framework

In this paper, we address the double bottleneck of game-theoretic learning by proposing *Haar-CFR* – a hybrid intelligent system that integrates signal processing with game theory. Our approach relies on our *Sparse-Haar Hypothesis*: we presume that regret landscapes over *ordered* strategy spaces—whether in finance or other real-world extensive-form games—are not random noise, but rather structured signals that admit sparse representations in an orthogonal wavelet basis.

By applying the Discrete Haar Wavelet Transform (DHT) and adaptive spectral thresholding to the regret accumulation process, we replace dense tabular storage with a compact spectral registry. This technique drastically reduces persistent memory requirements from linear to near-logarithmic scaling, thus effectively mitigating the Memory Wall without incurring the Time Wall of neural inference.

Our work makes the following contributions:

1. *Algorithm*: We introduce Haar-CFR – a novel spectral regularization algorithm. Experiments on large-scale benchmarks (Deep-Stack Kuhn Poker) demonstrate compression ratios exceeding 99.9% (reducing storage by $> 50,000\times$) while preserving equilibrium convergence.
2. *Efficiency*: We demonstrate that Haar-CFR overcomes the neural inference bottleneck, achieving a $6.5\times$ wall-clock speedup over the state-of-the-art, tuned Deep CFR baseline on high-dimensional domains.
3. *Unification*: In the financial application domain, we demonstrate that the learned strategy can recover autonomously, and it efficiently unifies the distinct logics of Avellaneda–Stoikov (inventory skew) and Glosten–Milgrom (information spread) within a single equilibrium model.
4. *Explainability*: We provide intrinsic explainability in our approach by supporting spectral decomposition that allows us to visualize how the solver separates strategic signals from exploration noise (Arrieta et al. 2020).

Our contributions establish wavelet-regularized regret minimization as a novel effective and transparent paradigm of soft computing for high-fidelity equilibrium computation. By treating the regret landscape as a signal, rather than a table, Haar-CFR offers a trade-off between the theoretical guarantees of game theory and the practical scalability requirements of applications in modern finance.

2 Related Work

Automated market making lies at the intersection of stochastic control, market microstructure theory, and computational intelligence. This section positions our contributions within these domains, specifically highlighting the tension between high-fidelity domain modeling and the resulting computational bottlenecks of the current game-theoretic solvers.

2.1 The Application Challenge: Market Making Theory

In this subsection, we analyze the two major paradigms of the market making theory and the challenge of bridging them.

Inventory-Based Models (Stochastic Control)

The inventory-centric paradigm aims to model the market maker as a risk-averse agent solving a dynamic optimization problem under price uncertainty. The foundational works of [Ho and Stoll \(1981\)](#) and [Avellaneda and Stoikov \(2008\)](#) formalize market making via stochastic control. Within this framework, the market maker tries to maximize the expected utility of terminal wealth, leading to Hamilton–Jacobi–Bellman (HJB) equations.

A central structural result from this domain is the *linear inventory skew*: optimal reservation prices adjust linearly based on inventory levels to induce mean reversion. While these models provide analytical tractability, they treat order flow as exogenous noise (Poisson processes) and do not explicitly model strategic interaction with informed traders, thus leaving the agent vulnerable to toxic flow.

Information-Based Models (Market Microstructure)

In contrast, market microstructure models ([Glosten and Milgrom 1985](#); [Kyle 1985](#)) frame market making as a Bayesian inference problem under asymmetric information. Here, the market maker updates beliefs about the asset’s fundamental value depending on order arrivals. The bid–ask spread in these models is not a function of inventory risk, but rather an *information premium* that compensates for losses against better-informed counterparties. While these models rigorously capture adverse selection risk, they typically assume risk neutrality and they abstract from the capital constraints inherent in real-world high-frequency trading.

The Unification Gap

Bridging the previously described two paradigms remains a theoretical challenge. Extensions of the Avellaneda–Stoikov framework have attempted to incorporate predictive signals or ambiguity aversion ([Cartea et al. 2017](#)). However, these approaches generally embed information effects heuristically into control formulations, rather than deriving a unified equilibrium from first principles. Consequently, a unified framework that simultaneously prices inventory and information risk within a single strategic equilibrium remains an open computational problem.

2.2 Computational Intelligence and the CFR Bottleneck

This subsection compares two main alternatives in addressing the current challenges of modern market making.

Deep Reinforcement Learning (DRL)

To address the non-linearities of modern markets, recent research has pivoted toward Deep Reinforcement Learning (Spooner et al. 2018; Ganesh et al. 2019). DRL agents can learn complex, state-dependent policies that adapt to varying market regimes. However, DRL in adversarial environments faces structural limitations: (i) simulator overfitting (Vyetenko et al. 2020), (ii) training instability, and (iii) a lack of intrinsic interpretability—a critical barrier for trusted autonomy (Rudin 2019; Arrieta et al. 2020). While recent theoretical work has established regret bounds for online market making (Pathak et al. 2024), these approaches typically focus on local pricing dynamics rather than solving a full extensive-form equilibrium.

Game-Theoretic Learning (CFR)

Counterfactual Regret Minimization (CFR) (Zinkevich et al. 2007) offers a rigorous alternative by guaranteeing convergence to an ϵ -Nash equilibrium in imperfect-information games. Unlike DRL that optimizes against a fixed environment, CFR minimizes regret against a worst-case adversary, which makes it highly robust for strategic exploitation. However, standard CFR suffers from the *State Space Explosion*. It requires storing the cumulative regrets for every information set in the game tree. In fine-grained domains with continuous or highly granular variables (e.g., thousands of price ticks), this leads to memory growth linear in the size of the game tree, rendering tabular CFR intractable on standard hardware (Gholami et al. 2024).

2.3 The Memory Wall: Abstraction vs. Approximation

Existing approaches to overcome the CFR memory bottleneck fall into two categories, both of which have significant drawbacks for high-fidelity modeling.

1. *State Abstraction*: Traditional methods rely on bucketing similar states together (e.g., discretizing prices into coarse bins) (Gilpin and Sandholm 2007). Early work focused on lossless abstraction, which is computationally expensive to compute. Imperfect abstraction techniques (such as clustering hand strength histograms in Poker) are standard but introduce non-recoverable errors. In complex domains, micro-structure details matter; a coarse abstraction potentially destroys signals required to detect informed traders, leading to substantial loss of robustness. More recent approaches like Discretization Soft-CFR (Schmid et al. 2021) attempt to learn discretizations but remain bound by the granularity of the underlying grid.

2. *Neural Approximation (Deep CFR)*: Methods like Deep CFR (Brown et al. 2019) and DREAM (Steinberger et al. 2020) replace the tabular storage with neural networks to approximate the regret surface. While this effectively decouples memory requirements from the state space size, it introduces two other bottlenecks:

- *The Time Wall*: The cost of neural inference, even when optimized via batching, is orders of magnitude higher than atomic table lookups (matrix multiplication vs. array indexing). This latency often prohibits deep lookahead in real-time solvers.
- *Replay Buffer Explosion*: As noted by [Li et al. \(2020\)](#), the Experience Replay Buffers required to train the neural networks must grow linearly with training iterations to ensure unbiasedness. This ironically reintroduces another memory constraint similar to the one neural methods originally aimed to solve.

2.4 Wavelet-Based Spectral Regularization

Our approach diverges from both abstraction and neural approximation by leveraging *Sparse Signal Processing*. We observe that regret landscapes over ordered domains are far from random; they exhibit local correlations and piecewise-smoothness.

While methods like Principal Component Analysis (PCA) or Fourier transforms have been applied to game payoffs ([Donoho 2006](#)), they typically rely on global basis functions. In particular, Fourier methods struggle with sharp discontinuities found in strategic thresholds (the Gibbs phenomenon), requiring many coefficients to approximate a simple step function (such as a sharp strategic boundary).

In contrast, *Haar wavelets* ([Daubechies 1992](#); [Mallat 2009](#)) are localized step functions that are uniquely suited for such signals. They can perfectly represent sharp strategic thresholds using a compact set of coefficients. By integrating discrete Haar wavelet sparsification directly into the CFR iteration loop, we propose a novel method that achieves the compression benefits of abstraction without significant precision loss, and the scalability of neural methods with reduced inference overhead. This effectively treats the game tree as a topological signal graph.

2.5 Positioning of Our Work

Our *Haar-CFR* approach fills the discussed gaps in the literature on hybrid intelligent systems, and it compares to the existing approaches as follows:

- *vs. Stochastic Control*: It is model-free regarding environment dynamics; this enables it to learn adverse selection endogenously without relying on exogenous noise assumptions.
- *vs. Deep RL/Deep CFR*: It is white-box and interpretable. The spectral decomposition allows researchers to visualize exactly which strategic frequencies are retained. Additionally, by using discrete basis functions rather than continuous neural approximations, it better preserves the sharp decision boundaries that are highly characteristic of optimal policies in ordered domains.
- *vs. Tabular CFR*: It mitigates the memory bottleneck and time wall, enabling the solution of high-fidelity games previously deemed computationally prohibitive on commodity hardware.

This establishes our Haar-CFR approach as a novel *Spectral Regularization* technique for Game Theory, offering a third path between the state-of-the art exact tabular methods and approximate neural methods.

2.6 Application of Soft Computing and Optimization

Our methodology is conceptually situated within the broader advancements of hybrid intelligent architectures. Recent literature highlights the convergence of combinatorial *Optimization* techniques (Cerulli et al. 2024) and nonparametric *Data Analytics and Machine Learning* systems (Angelov et al. 2017) to explicitly address high-dimensional uncertainty. Furthermore, the mathematical formulation of extensive strategic scenarios benefits heavily from suboptimal modeling paradigms and regularization techniques originally designed for learning from data (Gnecco and Sanguineti 2010). When applied to complex environments, these algorithmic approaches map effectively to modern domains requiring data-driven adaptations.

3 Our Methodology: Haar-CFR

We introduce *Haar-CFR* – a novel spectral variant of Counterfactual Regret Minimization, designed to resolve the storage bottleneck in extensive-form games with topologically ordered state spaces (Zinkevich et al. 2007). By viewing the game tree through the lens of signal processing, we transform the problem of storing strategic data into a sparse signal recovery problem (Candès and Tao 2006).

3.1 Regret as a Topological Signal

Standard Counterfactual Regret Minimization (CFR) treats each game state as a categorical index, thereby ignoring the metric or semantic relationship between adjacent states (e.g., a price of \$100 versus \$100.01) (Gilpin and Sandholm 2007). We instead endow the domain with its intrinsic order structure and treat regret as a structured signal over that domain.

Let $\mathcal{C} = \{c_1 < c_2 < \dots < c_N\}$ denote an ordered domain of private information (e.g., a discrete price grid or discretized hand strength). For a fixed history h and action a , define the *regret signal*

$$r_{h,a} \in \mathbb{R}^N$$

as the vector of cumulative regrets across the entire ordered domain:

$$r_{h,a} = [R((h, c_1), a), R((h, c_2), a), \dots, R((h, c_N), a)]^\top. \quad (1)$$

This representation promotes regret from an unstructured lookup table to a one-dimensional signal defined over a topological space.

In many economic and financial games, equilibrium strategies exhibit *regime-switching behavior*: policies remain constant over intervals of the domain and change only at threshold boundaries (e.g., switching from Buy to Sell at a critical price) (Cartea et al. 2015). Discrete bid auctions and ordered action grids exhibit similar threshold structure. This motivates the following structural hypothesis.

Definition 1 (Sparse-Haar Hypothesis for Discrete Ordered Games) Let $\mathcal{C} = \{c_1 < c_2 < \dots < c_N\}$ be a finite ordered grid. Assume:

1. Payoffs are Lipschitz-continuous in the underlying action variable;

2. The equilibrium best-response correspondence induces a partition of $\mathcal{C}$ into $S + 1$ contiguous regions within which the identity of the best response is constant.

Then the associated regret function $r : \mathcal{C} \rightarrow \mathbb{R}$ at equilibrium is piecewise-constant with at most S jump discontinuities. Consequently, its discrete Haar wavelet expansion contains at most $\mathcal{O}(S \log N)$ nonzero coefficients. In particular, there exists a constant $C > 0$ such that the best K -term Haar approximation $\hat{r}_K$ satisfies

$$K \geq CS \log N \implies \|r - \hat{r}_K\|_2 \leq \varepsilon_N,$$

where $\varepsilon_N \rightarrow 0$ as $N \rightarrow \infty$ under grid refinement.

Interpretation. The hypothesis asserts that equilibrium regret in ordered discrete games is *compressible*: the intrinsic complexity scales with the number of strategic regime changes S , rather than with the grid size N .

Caveat on Applicability. The hypothesis presumes structured equilibrium behavior. It may fail in games exhibiting highly oscillatory, high-entropy, or fractal mixed strategies, where no low-complexity partition exists. In such cases, the regret signal is not sparsely representable in the Haar basis, and the required coefficient count K scales linearly with N , recovering standard tabular CFR behavior as a worst-case limit.

3.2 The Haar-CFR Algorithm

Haar-CFR replaces the dense tabular storage ($\mathcal{O}(N)$ per node) of standard CFR with a *Sparse Spectral Registry* $\mathcal{S}$ ($\mathcal{O}(\log N)$ per node). The algorithm operates by reconstructing the dense signal only in *transient* memory during the traversal of a node. We employ *CFR+* (Tammelin et al. 2015) as the underlying regret minimization scheme due to its superior convergence properties in solving large-scale games.

Algorithm 1 Haar-CFR+ Update Loop

- 1: *Input:* History h , Compression Ratio κ
 - 2: *Phase 1: Reconstruction (Transient Cache)*
 - 3: Retrieve sparse coefficients $\hat{\gamma}$ from $\mathcal{S}[h]$
 - 4: $r_{dense} \leftarrow \text{IDHT}(\hat{\gamma})$ ▷ Fast Inverse Transform ($\mathcal{O}(N)$)
 - 5: *Phase 2: Strategy Computation*
 - 6: $\sigma(a) \leftarrow \text{RegretMatching}+(r_{dense})$ ▷ CFR+ local policy
 - 7: *Phase 3: Recursive Traversal*
 - 8: $v(a) \leftarrow \text{Traverse}(h, \sigma)$ ▷ Get counterfactual values
 - 9: $\rho \leftarrow v(a) - \langle \sigma, v \rangle$ ▷ Instantaneous regret
 - 10: $r_{new} \leftarrow \max(0, r_{dense} + \rho)$ ▷ CFR+ accumulation
 - 11: *Phase 4: Spectral Regularization*
 - 12: $\gamma_{new} \leftarrow \text{DHT}(r_{new})$ ▷ Forward Transform ($\mathcal{O}(N)$)
 - 13: $K \leftarrow \lceil \kappa \cdot N \rceil$
 - 14: $\hat{\gamma}_{new} \leftarrow \text{TopK}(\gamma_{new}, K)$ ▷ Keep largest K coeffs
 - 15: $\mathcal{S}[h] \leftarrow \hat{\gamma}_{new}$ ▷ Commit to Persistent Storage
-

Figure 2 visualizes the interaction between transient memory and persistent storage during the Haar-CFR+ update loop. To minimize the memory footprint, the dense regret vectors are only materialized in transient RAM during node evaluation. First, the sparse coefficients $\hat{\gamma}$ are fetched from the compressed registry and transformed into the dense domain via the Inverse Discrete Haar Transform (IDHT). After applying regret matching and accumulating the instantaneous counterfactual regrets (ρ), the updated dense vector r_{new} undergoes a Forward Discrete Haar Transform (DHT). Finally, adaptive Top-K thresholding compresses the signal back into a sparse representation ($\hat{\gamma}_{new}$), which is committed to the persistent storage registry.




Figure 2: Haar-CFR+ Pipeline. Visualization of the update loop detailing the data flow between persistent compressed storage and transient dense processing via discrete Haar transforms.

Remark 1 (Compressed Average Strategy) To ensure memory scaling remains logarithmic, we apply the same spectral compression pipeline to the *Cumulative Average Strategy* $\bar{\sigma}$. While the current policy σ is computed transiently, the average policy needed for the final Nash Equilibrium output is stored sparsely (i.e., retaining only the top- K most significant wavelet coefficients rather than the full dense strategy array) in $\mathcal{S}_{avg}$.

3.3 Theoretical Stability and Convergence

A central challenge in spectral compression is ensuring that truncation errors do not invalidate the no-regret guarantees of CFR (Zinkevich et al. 2007; Candès and Tao 2006).

Since Haar-CFR replaces the exact cumulative regret vector with a truncated spectral approximation, the learning dynamics become a perturbed regret minimization process.

We analyze the stability of our approach in two steps.

Step 1: Approximation Error Under Structured Regret.

Let $r^{(t)} \in \mathbb{R}^N$ denote the exact cumulative regret at iteration t , and let $\hat{r}^{(t)}$ be its best K_t -term Haar approximation. Define the truncation error

$$\epsilon_t := \hat{r}^{(t)} - r^{(t)}.$$

Under the Sparse-Haar Hypothesis (Definition 1), if $r^{(t)}$ is piecewise-constant with at most S discontinuities, then its Haar representation contains at most $\mathcal{O}(S \log N)$ significant coefficients. Consequently, for any $K_t \geq CS \log N$, the best K_t -term approximation satisfies

$$\|\epsilon_t\|_2 \leq \delta_N,$$

where $\delta_N \rightarrow 0$ under grid refinement ($N \rightarrow \infty$).

More generally, if $r^{(t)}$ is compressible in the Haar basis, the truncation error decays as K_t increases.

Step 2: No-Regret Under Bounded Perturbation.

We now show that CFR remains stable under sublinear cumulative perturbations.

Theorem 1 (No-Regret Under Sublinear Spectral Perturbation) *Let R_T denote the cumulative counterfactual regret after T iterations of exact CFR, and suppose the perturbed process uses truncated regrets $\hat{r}^{(t)} = r^{(t)} + \epsilon_t$ at each iteration.*

If the truncation errors satisfy

$$\sum_{t=1}^T \|\epsilon_t\|_2 = o(T),$$

then the perturbed process preserves the no-regret property:

$$\lim_{T \rightarrow \infty} \frac{R_T}{T} = 0.$$

Proof Sketch Standard CFR guarantees $R_T \leq \Delta\sqrt{T}$ for payoff range Δ . Under additive perturbations ϵ_t , cumulative regret is bounded by

$$R'_T \leq \Delta\sqrt{T} + \sum_{t=1}^T \|\epsilon_t\|_2.$$

Dividing by T ,

$$\frac{R'_T}{T} \leq \frac{\Delta}{\sqrt{T}} + \frac{1}{T} \sum_{t=1}^T \|\epsilon_t\|_2.$$

The first term vanishes as $T \rightarrow \infty$. By assumption, the second term also vanishes. Hence average regret converges to zero. $\square$

Implications for Haar-CFR

If the compression schedule K_t is chosen such that the induced truncation errors grow sublinearly in aggregate (for example, by fixing $K_t \geq CS \log N$ under the Sparse-Haar hypothesis, or by slowly increasing K_t over time), then Haar-CFR preserves the no-regret property of CFR.

Since no-regret at every information set implies convergence to a Nash equilibrium in two-player zero-sum games, Haar-CFR converges to an ϵ -Nash equilibrium, where the approximation error satisfies

$$\epsilon = \mathcal{O}\left(\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \|\epsilon_t\|_2\right).$$

In particular, when truncation errors are uniformly bounded and summable sublinearly, equilibrium accuracy is controlled directly by the compression schedule.

4 Experimental Setup

To rigorously evaluate the proposed Haar-CFR framework, we utilize a multi-domain test suite comprising three standard game-theoretic benchmarks of increasing complexity and a high-dimensional financial market making stress test.

4.1 Environment A: Standard Imperfect-Information Games

To reliably validate that our method generalizes across fundamentally different topologies, we evaluate Haar-CFR on three standard imperfect-information games:

1. *Leduc Hold'em Poker*: A standard benchmark with a deck of 6 cards. We discretize the betting pot into $N = 100$ increments to simulate continuous betting.
2. *Deep-Stack Kuhn Poker (GKP- N)*: A scalable stress test designed to probe memory scaling. The game uses a deck of size $N \in \{1k, \dots, 10M\}$ and allows for a sequence of *Bet*, *Raise*, and *Re-raise* (max depth 3). As N increases, the information state space grows linearly, isolating the memory bottleneck.
3. *Flop Hold'em Poker (FHP)*: A large-scale subset of No-Limit Texas Hold'em involving preflop and flop betting rounds. The game tree size exceeds 10^{10} nodes, representing a “Grand Challenge” for tabular methods.

4.2 Environment B: The Unified Market Making Stress Test

We developed a unified market making framework to simulate the *Impossible Triangle* of market making (the theoretical trade-off where a market maker cannot simultaneously maintain tight spreads, keep zero inventory risk, and remain immune to adverse selection). *Parameters*:

- *Price Grid ($\mathcal{G}$)*: Discrete grid with $N = 2000$ levels.
- *Inventory State (q)*: $q \in \{-10, \dots, +10\}$.

- *Toxicity (α)*: Sampled $\alpha \sim U[0, 1]$ per episode (Adversarial Nature).

This environment forces the agent to learn a single unified strategy that handles both inventory risk (Avellaneda–Stoikov (Avellaneda and Stoikov 2008)) and adverse selection (Glosten–Milgrom (Glosten and Milgrom 1985)), serving as our primary high-dimensional application challenge.

4.3 Comparators and Baselines

We compare Haar-CFR against three distinct categories of baselines:

1. *Ground Truth (Tabular CFR+)*: The standard, uncompressed implementation of CFR+. It serves as the exact accuracy benchmark for tractable games ($N < 100$).
2. *Analytical Baseline (Avellaneda–Stoikov)*: The closed-form theoretical solution for inventory management (evaluated in the Financial domain only).
3. *Neural Approximation (Deep CFR)*: To ensure a rigorous comparison, we implemented a highly optimized Deep CFR baseline to avoid “straw man” critiques:
 - *Architecture*: 5-layer ResNet-style MLP (256 units) with skip connections.
 - *Memory Budget*: A large Experience Replay Buffer of size 10^7 to ensure stability.
 - *Inference Optimization*: We utilize *batched inference* (batch size 4096) on an NVIDIA RTX 4090 to minimize the neural overhead during tree traversal.

4.4 Hardware & Implementation Specs

To ensure reproducibility and isolate the source of performance gains, all experiments were conducted on a high-performance workstation with the following specifications:

- *CPU*: AMD Ryzen 5950X (16-core, 32-thread).
- *GPU*: NVIDIA RTX 4090 (Used exclusively for Deep CFR inference).
- *RAM*: 64GB DDR4.
- *Software*: C++17 implementation with AVX2 optimizations for the Fast Haar Transform.

This setup ensures that the reported speedups for Haar-CFR are due to algorithmic efficiency (CPU-bound) rather than hardware deficits in the baseline.

5 Experimental Results and Analysis

In this section, we empirically validate the Haar-CFR framework using our multi-domain evaluation suite. To align with our methodological contributions, we first evaluate the algorithmic properties on standard game-theoretic benchmarks (mitigating the Memory and Time walls, and verifying robustness). We then evaluate the framework on the high-dimensional financial market making stress test to demonstrate the emergence of unified, interpretable equilibrium strategies.

5.1 Scalability: Mitigating the Memory Wall

We first address the primary algorithmic contribution: storage efficiency. The *Deep-Stack Kuhn Poker (GKP)* benchmark allows us to stress-test scaling behavior by increasing the private state space size N from 1,000 to 10,000,000.

Table 1 presents the peak memory usage. Standard Tabular CFR hits a "Feasibility Cliff" at $N = 10^5$, exhausting 64GB of RAM. In contrast, Haar-CFR ($\kappa = 0.01$) exhibits near-logarithmic scaling. Crucially, as the game size grows by a factor of $10,000\times$ (from $1k$ to $10M$), the memory footprint of Haar-CFR increases by only $3.5\times$ (0.6 MB to 2.1 MB).

Table 1: Peak memory usage vs. state space size (N) in Deep-Stack Kuhn Poker (GKP). Deep CFR memory is dominated by the replay buffer (10^7 samples).

N	Tabular (MB)	Deep CFR (MB)	Haar-CFR (MB)	Reduction
10^3	1.2×10^1	8.55×10^2	6.0×10^{-1}	2.0×10^1
10^4	1.2×10^2	8.60×10^2	8.0×10^{-1}	1.5×10^2
10^5	1.2×10^3	8.80×10^2	1.1	1.09×10^3
10^6	OOM (>64 GB)	9.50×10^2	1.5	8.0×10^3
10^7	OOM (>128 GB)	1.10×10^3	2.1	$> 5.0 \times 10^4$

This result empirically confirms that spectral regularization *decouples memory scaling from domain resolution*, thus allowing for arbitrarily fine-grained discretizations without the combinatorial memory explosion.

5.2 Runtime Analysis: The Time Wall

An important critique of spectral methods is the overhead of the transform. To verify that compression does not come at the cost of training time, we evaluate the convergence speed on the large-scale *Flop Hold'em (FHP)* benchmark.

Figure 3 (Left Panel) plots exploitability (mbb/h) against wall-clock time:

- *Deep CFR (Red)*: Exhibits a shallower slope due to the "Time Wall" of neural inference. Despite batched inference on an RTX 4090, the data transfer and forward pass latency create a bottleneck, reaching 100 mbb/h after 26 hours.
- *Haar-CFR (Blue)*: Leverages the deterministic low latency of the CPU-based Fast Haar Transform ($\approx 12\mu s$ per node). It crosses the 100 mbb/h threshold in ≈ 4 hours, representing a $6.5\times$ *speedup* over the neural baseline in this domain.

5.3 Robustness and Ablation Studies

To validate the bounds of the *Sparse-Haar Hypothesis*, we performed an ablation study on the compression ratio κ using Leduc Hold'em (Figure 3, Right Panel).




Figure 3: Performance Benchmarks. *Left:* Time vs. Exploitability on FHP. Haar-CFR (Blue) descends sharply, crossing 100 mbb/h at ≈ 4 h, while Deep CFR (Red) is slowed by the neural inference overhead. *Right:* Compression Ablation on Leduc. The Exact baseline (Black) converges smoothly. $\kappa = 0.05$ (Orange) demonstrates near-lossless abstraction, while $\kappa = 0.001$ (Green) hits a hard floor due to insufficient coefficients.

- *Stable Regime* ($\kappa = 0.05$): The Orange curve tracks the exact tabular baseline (Black) with minimal deviation, confirming that retaining top 5% of coefficients is sufficient to capture the strategic signal.
- *Failure Mode* ($\kappa = 0.001$): The Green curve drops initially but hits a hard plateau at ≈ 250 mbb/h. This validates our theoretical bound derived in Section 3: when $K < S \log |\mathcal{C}|$, the basis functions are insufficient to represent the strategic discontinuities, leading to convergence failure.

5.4 The Application Challenge: Validation of Financial Risk Paradigms

Having established the core algorithmic properties, we turn to the stress test of financial market making. Before evaluating the hybrid model, we first establish the ground truth by isolating the two fundamental risk drivers. Figure 4 illustrates the "Theoretical Conflict" that necessitates a unified solver.

The simulation confirms that classical models prescribe conflicting behaviors under stress: Inventory Logic dictates lowering prices to attract flow, while Information Logic dictates raising prices to repel toxic flow. This validation confirms that the benchmark correctly captures the physics of market making.




Figure 4: The Theoretical Conflict. Panel A (Left) visualizes the analytical Avellaneda–Stoikov baseline: quotes are driven purely by *Inventory Pressure*, resulting in a linear skew to liquidate positions. Panel B (Right) simulates the Glosten–Milgrom baseline: quotes are driven exclusively by *Information Asymmetry*, resulting in spread widening proportional to toxicity. Note the conflicting directional pressure on the quote price.

5.5 Emergence of the Unified Sigmoidal Policy

When trained with Haar-CFR, our solver does not merely average these baselines; it discovers a non-linear policy surface that dynamically switches between regimes. Figure 5 visualizes the converged Nash Equilibrium.

The agent autonomously identifies a novel structure we term the *Strategic Hybrid Sigmoidal Policy*. This S-shaped curve explains the generative mechanism behind the liquidity supply schedules empirically observed in limit order books (Ting and Warachka 2003; Huang and Ting 2008). As illustrated in Figure 5, the agent learns a hierarchical risk preference: *inventory risk is subordinate to adverse selection risk*. At high toxicity ($\alpha > 0.6$), the policy hits a "Defensive Wall" (Red Plateau), where the agent effectively widens the spread to infinity to avoid engagement, regardless of inventory holdings.

5.6 Intrinsic Explainability: Spectral Decomposition of Logic

Unlike Deep RL, which treats policy logic as a black box, Haar-CFR offers *Intrinsic Explainability*. By analyzing the retained wavelet coefficients, we can visualize exactly which risk factors drive the decision-making process.

Figure 6 demonstrates that the "features" learned by the agent are orthogonal in the frequency domain. This spectral disentanglement allows regulators to audit the algorithm's sensitivity. For instance, removing the horizontal high-frequency bands in Panel B would effectively "lobotomize" the agent's ability to detect toxic flow, while leaving its inventory management capabilities intact.




Figure 5: *The Unified Policy Surface.* The heatmap displays the optimal ask premium $\pi(q, \alpha)$. It reveals a non-linear decision boundary: at low toxicity ($\alpha < 0.3$), the agent follows linear inventory logic (blue gradient). At high toxicity, the policy hits a "GM Barrier" (red plateau), effectively refusing to trade. The dashed line represents the Neutral Price (Inventory=0).

5.7 Statistical Validation

We evaluate the robustness of the method across domains. Table 2 presents the comparative results.

Table 2: Statistical Comparison of Agent Performance ($N = 100$ Episodes)

Comparison Pair	Domain	p-value	Result
Haar-CFR vs. Tabular	Leduc Hold'em	0.812	Tie (Lossless Recovery)
Haar-CFR vs. Tabular	MM (Finance)	0.674	Tie (Lossless Recovery)
Haar-CFR vs. Deep CFR	MM (Finance)	< 0.001	<i>Haar-CFR Superior</i>

Analysis of Neural Baseline Deficits: The statistical superiority of Haar-CFR over Deep CFR ($p < 0.001$) can be attributed to the *Spectral Bias* of neural networks (Rahaman et al. 2019). Neural networks tend to learn low-frequency functions first and struggle to approximate high-frequency discontinuities (such as the sharp "Do Not




Figure 6: Spectral Audit of Strategic Logic. Panel A displays the low-frequency vertical bands corresponding to inventory states, proving the agent has isolated Inventory Risk (q). Panel B displays the horizontal bands corresponding to toxicity regimes, proving the isolation of Adverse Selection Risk (α). The intersection of these bands creates the equilibrium strategy.

Trade" cliff in Figure 5) without excessive training. Wavelets, being localized basis functions, capture these edges instantaneously.

6 Conclusion

This work addresses the fundamental conflict in computational intelligence between high-fidelity domain modeling and algorithmic tractability. We present *Haar-CFR* – a novel hybrid soft computing framework that fuses Counterfactual Regret Minimization with Wavelet Signal Processing to mitigate the “Memory Wall” in large-scale imperfect-information games.

6.1 Summary of Contributions

By formalizing the *Sparse-Haar Heuristic*, we demonstrate that strategic complexity in ordered domains is not irreducible. Instead, it manifests as a piecewise-smooth signal that admits a highly sparse representation in the spectral domain. This insight allows us to decouple memory requirements from state space resolution, achieving compression ratios exceeding 99.9% in large-scale settings while preserving equilibrium convergence.

Empirically, Haar-CFR proves superior to both classical analytical models and modern neural baselines across a multi-domain evaluation suite:

- *Extreme Scalability:* On the Deep-Stack Kuhn Poker (GKP) stress test, Haar-CFR demonstrated *near-logarithmic memory scaling*, reducing the persistent

storage requirement by a factor of $> 50,000\times$ as the game size expanded from 10^3 to 10^7 states.

- *Computational Efficiency:* Unlike Deep CFR, which hits a “Time Wall” due to neural inference latency, Haar-CFR achieves a $6.5\times$ *wall-clock speedup* on the Flop Hold’em benchmark. This confirms that discrete spectral transforms provide a more efficient inductive bias for step-function strategies than continuous neural approximations.
- *The Application Challenge (Financial Unification):* In the market making domain, the agent successfully unifies the distinct theoretical paradigms of Inventory Risk (Avellaneda–Stoikov) and Adverse Selection (Glosten–Milgrom) into a single *Strategic Hybrid Sigmoidal Policy*.

6.2 Broader Impact on Soft Computing

Beyond the specific domain of market making, this framework establishes a generalizable blueprint for *Explainable AI (XAI)* in strategic reasoning. We showed that the “black box” of an agent’s policy can be opened via spectral decomposition, allowing human experts to audit the specific frequency bands corresponding to distinct risk factors. This offers a rigorous alternative to opaque neural networks in safety-critical domains where interpretability is non-negotiable.

6.3 Limitations and Future Directions

The efficacy of Haar-CFR is bounded by the sparsity of the equilibrium strategy. As demonstrated in our ablation study, when the coefficient count falls below the information-theoretic limit ($K < S \log |\mathcal{C}|$), the method fails to capture strategic discontinuities, leading to convergence collapse.

Future work will focus on three avenues: (1) extending the framework to $2D$ *Wavelets* to capture correlations across multiple assets or derivative surfaces; (2) leveraging the independence of wavelet coefficients for massive parallelization on GPUs; and (3) applying this spectral regularization technique to other ordered game-theoretic domains, such as energy grid bidding and auctions.

Statements and Declarations

Funding: This work was partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project Nr 470527619 (MDH.DL) at the University of Muenster, Germany.

Competing Interests: The authors declare they have no financial or non-financial interests that are directly or indirectly related to the work submitted for publication.

Data Availability: The datasets generated and/or analysed during the current study are available from the corresponding author on reasonable request.

Author Contribution: **Florian Fey:** Conceptualization, Methodology, Software, Writing - original draft. **Sergei Gorlatch:** Supervision, Writing - review & editing.

References

- Angelov, P., Gu, X., and Kangin, D., 2017. Empirical Data Analytics. *International Journal of Intelligent Systems*, 32(12), pp.1261-1284. <https://doi.org/10.1002/int.21899>
- Arrieta, A.B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R. and Chatila, R., 2020. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, pp.82-115.
- Avellaneda, M. and Stoikov, S., 2008. High-frequency trading in a limit order book. *Quantitative Finance*, 8(3), pp.217-224.
- Brown, N., Lerer, A., Gross, S. and Sandholm, T., 2019. Deep Counterfactual Regret Minimization. In *ICML*.
- Candès, E.J. and Tao, T., 2006. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Transactions on Information Theory*, 52(12), pp.5406-5425.
- Cartea, A., Jaimungal, S. and Penalva, J., 2015. *Algorithmic and High-Frequency Trading*. Cambridge University Press.
- Cartea, A., Donnelly, R. and Jaimungal, S., 2017. Algorithmic trading with advisor signals. *Applied Mathematical Finance*, 24(1), pp.1-30.
- Cerulli, R., D'Ambrosio, C., and Raiconi, A., 2024. A biased random-key genetic algorithm for the knapsack problem with forfeit sets. *Soft Computing*, 28, pp.12021-12041. <https://doi.org/10.1007/s00500-024-09948-w>
- Daubechies, I., 1992. *Ten Lectures on Wavelets*. SIAM.
- Donoho, D.L., 2006. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4), pp.1289-1306.
- Ganesh, S., Vadori, N., Xu, M., Zheng, H., Reddy, P. and Veloso, M., 2019. Reinforcement learning for market making in a multi-agent dealer market. *arXiv preprint arXiv:1911.05892*.
- Gholami, A., et al., 2024. AI and Memory Wall. *IEEE Micro*.
- Gilpin, A. and Sandholm, T., 2007. Lossless abstraction of imperfect information games. *Journal of the ACM (JACM)*, 54(5), pp.25-es.
- Glosten, L.R. and Milgrom, P.R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), pp.71-100.

- Gnecco, G., and Sanguineti, M., 2010. Regularization Techniques and Suboptimal Solutions to Optimization Problems in Learning from Data. *Neural Computation*, 22(3), pp.793-829. <https://doi.org/10.1162/neco.2009.05-08-786>
- Ho, T. and Stoll, H.R., 1981. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9(1), pp.47-73.
- Huang, R.D. and Ting, H., 2008. The liquidity supply schedule of the limit order book. *Journal of Financial Markets*, 11(4).
- Kyle, A.S., 1985. Continuous auctions and insider trading. *Econometrica*, 53(6), pp.1315-1335.
- Li, H., Hu, K., Zhang, S. and Qi, Y., 2020. Analysis of replay buffer in deep reinforcement learning. *arXiv preprint arXiv:2010.03362*.
- Mallat, S., 2009. *A Wavelet Tour of Signal Processing: The Sparse Way*. Academic Press.
- Pathak, V., Cesa-Bianchi, N., Cesari, T. and Colomboni, R., 2024. Market Making without Regret. *arXiv preprint arXiv:2411.13993*.
- Rahaman, N., Aristidou, A., Dutta, P., Jha, S., Fischer, P., Wolf, S. and Krueger, D., 2019. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pp. 5301-5310. PMLR.
- Rudin, C., 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), pp.206-215.
- Schmid, M., et al., 2021. Player of games. *arXiv preprint arXiv:2112.03178*.
- Spooner, T., Nelson, J., Joy, T. and Balasooriya, P., 2018. Market Making via Reinforcement Learning. In *AAMAS*.
- Steinberger, E., Lerer, A. and Brown, N., 2020. DREAM: Deep Regret minimization with Advantage baselines and Model-free learning. *arXiv preprint arXiv:2006.10410*.
- Tammelin, O., Burch, N., Johanson, M. and Bowling, M., 2015. Solving heads-up limit texas hold'em. In *IJCAI*, pp. 645-652.
- Ting, H. and Warachka, M., 2003. The limit order book's S-shaped liquidity supply schedule. *Work. Pap., Singapore Management University*.
- Vyetrenko, S., Byrd, D., Petosa, N., Xu, M., Deng, Y., Li, S. and Tucker, B., 2020. Get real: Realism metrics for robust limit order book market simulations. In *ICAIF*.

Zinkevich, M., Johanson, M., Bowling, M. and Piccione, C., 2007. Regret minimization in games with incomplete information. *Advances in Neural Information Processing Systems*, 20.