Master Thesis Proposal

AI-Assisted Formalization of Finite Learning in Games in Lean 4

Duration: 6 months **Type:** library-building thesis with a fully verified core theorem

Project idea. This thesis builds a reusable Lean 4 library for the finite-game foundations of learning in games. The mathematical objects are standard but easy to mis-state in informal writing: finite action profiles, mixed strategies, expected costs, unilateral deviations, regret, empirical play, Nash equilibrium, and coarse correlated equilibrium. The goal is not only to prove one theorem, but to create a well-documented formal layer that future projects can extend toward richer learning dynamics.

The project is designed for intensive AI-assisted proof engineering. AI tools may be used for theorem search, tactic suggestions, refactoring, examples, and documentation. Correctness is nevertheless supplied only by Lean 4: every definition and theorem in the submitted library must compile without hidden axioms.

Central mathematical target. The main verified theorem is the classical link between no-regret learning and coarse correlated equilibrium, formalized in a precise finite setting:

> For a finite normal-form game, if every player has average external regret at most ε_i along a finite play history of length T, then the empirical distribution of play is an ε-coarse correlated equilibrium, with $\varepsilon = \max_i \varepsilon_i$; a second version should keep the player-dependent error bounds.

This theorem is intentionally finite and explicit. That is not a downgrade: it is the right foundation for a formal library. It avoids measure theory and stochastic-process overhead while still proving the central learning-in-games result that makes regret meaningful as an equilibrium concept.

Mathematical scope. The thesis works with finite normal-form games. Let I be a finite set of players, A_i a finite action set, and $c_i : \prod_i A_i \to \mathbb{R}$ a cost function. The library should formalize pure profiles, unilateral deviations, mixed strategies as finite probability distributions, expected costs, best responses, Nash equilibrium, distributions over profiles, empirical distributions of play, external regret, and coarse correlated equilibrium. Continuous actions, Bayesian games, measurable strategies, stochastic approximation, and infinite-dimensional dynamics are deliberately left for future work.

Core deliverables.

1. **Finite-game API.** Define players, actions, pure profiles, costs/payoffs, profile replacement for unilateral deviations, and finite sums over action profiles.
2. **Mixed strategies and expected costs.** Define finite mixed strategies, product mixed profiles, expected cost, and the linearity/affinity lemmas needed for deviations.
3. **Equilibrium definitions.** Formalize best response, ε-best response, Nash equilibrium, ε-Nash equilibrium, profile distributions, coarse correlated equilibrium, and ε-coarse correlated equilibrium.
4. **Regret and empirical play.** Define play histories, empirical distributions, player-wise external regret, average regret, and the finite-sum identities connecting regret inequalities to deviation inequalities.
5. **Main theorem.** Prove in Lean 4 that player-wise external regret implies approximate coarse correlated equilibrium of the empirical distribution.
6. **Examples and tests.** Include verified examples such as matching pennies, rock-paper-scissors, and a coordination game, mainly to test that the API is usable.
7. **Stretch goal.** Add swap regret and approximate correlated equilibrium, or develop a richer finite-game equilibrium API if the core theorem is completed early.

Expected repository structure.

```
LearningInGames/FiniteGame/Basic.lean
LearningInGames/FiniteGame/MixedStrategy.lean
LearningInGames/FiniteGame/ExpectedCost.lean
LearningInGames/FiniteGame/Nash.lean
LearningInGames/FiniteGame/Regret.lean
LearningInGames/FiniteGame/CoarseCorrelated.lean
LearningInGames/Examples/MatrixGames.lean
```

The repository should use lake, continuous integration, namespaces, comments for important definitions, and examples that compile from a clean checkout.

Six-month plan.

1. **Weeks 1–2: Lean and mathlib sprint.** Set up the project. Learn finite types, finite sums, real inequalities, finite probability distributions, notation, and existing `mathlib` conventions.
2. **Weeks 3–5: finite-game foundation.** Implement pure profiles, cost evaluation, unilateral replacement, deviation notation, and small matrix-game examples.
3. **Weeks 6–8: mixed strategies and expectations.** Implement finite mixed strategies, product profiles, expected costs, and the basic linearity lemmas.
4. **Weeks 9–11: equilibrium layer.** Implement best responses, Nash equilibrium, approximate Nash equilibrium, distributions over profiles, CCE, and approximate CCE.
5. **Weeks 12–15: regret layer.** Implement histories, empirical distributions, external regret, average regret, and the finite-sum identities needed for the main theorem.
6. **Weeks 16–19: main theorem.** Prove no-external-regret implies approximate CCE. Produce both a compact theorem statement and convenient corollaries.
7. **Weeks 20–24: ambition phase.** Add the stretch goal if feasible; otherwise improve documentation, simplify APIs, add examples, remove fragile tactic scripts, and write the thesis.

AI-assisted workflow. AI tools are part of the intended working method. They should accelerate theorem-name discovery, tactic exploration, boilerplate generation, and proof refactoring. The thesis should include a short reproducibility appendix describing how AI was used. No central result may rely on an unverified AI claim; the final authority is the Lean 4 kernel.

Success criteria. A strong thesis delivers a compiling library with clear definitions of finite games, mixed strategies, regret, Nash equilibrium, and coarse correlated equilibrium; a complete proof that low external regret implies approximate CCE; several verified examples; and documentation good enough for another student to continue the project. An excellent thesis additionally provides swap-regret or correlated-equilibrium infrastructure.

Ambition and feasibility. The thesis is ambitious because it builds reusable infrastructure, not because it tries to formalize the most general possible theorem. The core theorem is finite, explicit, and central. This makes the project feasible in six months while still producing a serious foundation for verified learning in games.

Prerequisites. The student should know basic game theory, finite probability, elementary real analysis, and be willing to work intensively in Lean 4. Prior proof-assistant experience is helpful but not required for a mathematically strong student using AI tools responsibly.