

MAXWELL FISHELSON

maxkfish.com

maxfish@ias.edu

POSITIONS

Postdoctoral Associate

2026-present

Institute for Advanced Study (IAS)

Computer Science and Discrete Mathematics

EDUCATION

Ph.D. in Computer Science

2020-2026

Massachusetts Institute of Technology (MIT)

Supervisor: Prof. Constantinos Daskalakis

B.S. in Mathematics

2016-2020

Massachusetts Institute of Technology (MIT)

GPA: 4.9/5

RESEARCH

Swap Regret Minimization Through Response-Based Approachability

COLT 2026

[\[Paper\]](#)

Ioannis Anagnostides, Gabriele Farina, Maxwell Fishelson, Haipeng Luo, Jon Schneider

- Developed a simple, computationally efficient online learning algorithm that achieves $O(d\sqrt{T})$ linear swap regret over general convex sets
- Improved best known regret bound of $O(d^4\sqrt{T})$ (STOC'25) and eliminated computationally intensive per-iteration ellipsoid calls
- Extended the algorithmic framework to profile swap regret minimization, improving best known bound of $O(d^4\sqrt{T})$ (EC'25 best paper)
- Proved a matching information-theoretic lower bound of $\Omega(d\sqrt{T})$ for linear swap regret

High-Dimensional Calibration from Swap Regret

NeurIPS 2025

[\[Paper\]](#) [\[NeurIPS Talk\]](#)

Oral Presentation (77 of 25000 submissions)

Maxwell Fishelson, Noah Golowich, Mehryar Mohri, Jon Schneider

- Proposed *Tree-Cal*: a simple, efficient forecaster for calibration over high-dimensional action sets. Calibration in this regime was previously not known to be possible
- Established universal analysis that demonstrates that single algorithm simultaneously minimizes calibration with respect to all norm distances with optimal rate simultaneously
- Linked calibration rates to Online Linear Optimization rates: whenever OLO over $(P, \|\cdot\|_*)$ admits regret $O(\sqrt{\rho T})$, Tree-Cal attains the corresponding ϵ -calibration rate.
- Proved strong lower bounds: any algorithm achieving ℓ_1 -calibration error $\leq \epsilon T$ on the simplex requires exponential rounds in $1/\epsilon$

Breaking the $T^{2/3}$ Barrier for Sequential Calibration

STOC 2025

[\[Paper\]](#) [\[STOC Talk\]](#)

Invited to SICOMP Special Issue

Yuval Dagan, Constantinos Daskalakis, Maxwell Fishelson

Invited to HALG 2026

Noah Golowich, Robert Kleinberg, Princewill Okoroafor

- Established an algorithm for sequential prediction achieving $O(T^{2/3-\epsilon})$ calibration error, improving on the state of the art $O(T^{2/3})$ of Foster and Vohra, 1998
- Improved lower bound $\Omega(T^{0.54})$ on calibration error achievable in sequential prediction
- Established first $\omega(T^{1/2})$ calibration lower bound for an oblivious adversary

- New upper bound and tighter lower bound reductions from calibration to the sign preservation game of Qiao and Valiant
- Devised new recursive approach to establish upper bounds in the sign preservation game

Efficient Learning and Computation of Linear Correlated Equilibrium in General Convex Games

STOC 2025

[\[Paper\]](#) [\[STOC Talk\]](#)

Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Charilaos Pipis, Jon Schneider

- Established a general approach for minimizing linear swap regret with arbitrary action sets
- Established analogous ellipsoid based approaches for the computation of linear correlated equilibrium in games with arbitrary action sets
- Introduced the technique of regret minimization via semi-separation oracle (as opposed to the classic separation oracle) to handle intractable deviation sets
- Semi separation has already found broader applications (e.g. establishing the first poly-time algorithms for profile swap regret minimization)

From External to Swap Regret 2.0: An Efficient Reduction for Large Action Spaces

STOC 2024

Invited to SICOMP Special Issue

[\[Paper\]](#) [\[STOC Talk\]](#)

Yuval Dagan, Constantinos Daskalakis, Maxwell Fishelson, Noah Golowich

- Devised *Tree-Swap*: the first swap regret minimization algorithm with rate depending logarithmically on the number of actions N , resolving an open problem of Blum and Mansour from 2007
- Proved an equivalent swap regret rate only dependent on the Littlestone dimension of a concept class, establishing the first no-swap-regret algorithm for potentially infinite concept classes
- Established matching lower bounds, showing even oblivious, ℓ_1 -constrained adversaries can force a swap regret learner to expend exponentially many rounds
- Improved the best known dependence on N for swap regret learning in bandit learning
- Resolved open problems about the query and communication complexity of computing ϵ -approximate correlated equilibrium in games
- Proved ϵ -approximate correlated equilibria in extensive-form games can be computed efficiently for constant ϵ , advancing a long-standing open problem for extensive-form games

Full Swap Regret and Discretized Calibration

ALT 2025

[\[Paper\]](#) [\[ALT Talk\]](#)

Maxwell Fishelson, Robert Kleinberg, Princewill Okoroafor, Renato Paes Leme, Jon Schneider, Yifeng Teng

- Establish an efficient learning algorithm for swap regret minimization in large, structured games
- Introduce the problem of full swap regret minimization, a metric that encompasses all previously studied regret measures
- Established a family of algorithms attaining various full swap regret minimization rates for various loss settings
- Provided an algorithm for L_2 calibration matching the best-known rate, but with substantial algorithmic simplification over the state-of-the-art
- Introduced the problem of discretized calibration and established non-trivial rates using a novel regret minimization technique for nearly-strongly convex loss functions, which may be of independent interest

Online Learning and Solving Infinite Games with an ERM Oracle

COLT 2023

[\[Paper\]](#) [\[COLT Talk\]](#)

Angelos Assos, Idan Attias, Yuval Dagan, Constantinos Daskalakis, Maxwell Fishelson

- Devised the first algorithm for online learning of arbitrary potentially-infinite hypothesis classes using only ERM oracle calls
- Provided a theoretical backing to the prevalent “double-oracle” algorithms that are the standard for solving large games in practice, unlike previous online learning algorithms that relied on the SOA oracle and had no practical utility
- Obtained similar results for non-parametric games, providing learning algorithms that only rely on best response oracles and converge to approximate-minimax equilibria in two-player zero-sum games and approximate coarse correlated equilibria in multi-player general-sum games

Near-Optimal No-Regret Learning for Correlated Equilibria in Multi-Player General-Sum Games

STOC 2022

[\[Paper\]](#) [\[STOC Talk\]](#)

Ioannis Anagnostides, Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Noah Golowich, Tuomas Sandholm

- Constructed algorithmic extensions of Optimistic Hedge attaining $\text{poly}(\log T)$ internal-regret and swap-regret in multi-player general-sum games respectively
- Established a means of simulating the Stoltz-Lugosi Optimistic-Hedge algorithm as an instance of Optimistic Hedge on a combinatorial space, demonstrating that the no-external-regret of Optimistic Hedge implies no-internal-regret for Stoltz-Lugosi
- Introduced novel techniques for the analysis of Taylor expansions of multinomial functions arising from no-regret algorithms, enabling a proof that Blum-Mansour Optimistic Hedge achieves no-swap-regret
- A corollary of our bound is that both Stoltz-Lugosi and Blum-Mansour Optimistic Hedge converge to correlated equilibrium in general games at a rate of $\tilde{O}(1/T)$.

Near-Optimal No-Regret Learning in General Games

NeurIPS 2021

[\[Paper\]](#)

Oral Presentation (55 of 9000 submissions)

Constantinos Daskalakis, Maxwell Fishelson, Noah Golowich

- Established that Optimistic Hedge – a common variant of multiplicative-weights-updates with recency bias – attains $\text{poly}(\log T)$ regret in multi-player general-sum games
- Exponentially improved on the best known regret attainable by no-regret learners in general games
- Introduced many novel techniques for the analysis of the performance of regret minimization algorithms, including Fourier analysis
- A corollary of our bound is that Optimistic Hedge converges to coarse correlated equilibrium in general games at a rate of $\tilde{O}(1/T)$.

Multi-item Non-truthful Auctions Achieve Good Revenue

SICOMP 2022

(a.k.a. Simple, Credible, and Approximately Optimal Multi-item Auctions) EC 2020

[\[Paper\]](#) [\[EC Talk\]](#)

Constantinos Daskalakis, Maxwell Fishelson, Brendan Lucier, Vasilis Syrgkanis, Santhoshini Velusamy

- Established that first-price-type auctions can achieve a constant factor of the optimal revenue in the multi-item auction setting, resolving an open question
- Established the first credible and static multi-item auction that is approximately revenue optimal

- Obtained approximately-revenue-optimal multi-item mechanisms with fixed entry fees that are amenable to tuning via online learning techniques
- Proved a geometric lemma that enabled analysis of the utility of a first price auction, showing that welfare loss in a first price auction is at most 4 times the revenue of the posted price mechanism

Pattern Avoidance Over a Hypergraph

Electronic Journal of Combinatorics 2021

[\[Paper\]](#)

Maxwell Fishelson, Benjamin Gunby

- Achieved a generalization of the Stanley-Wilf theorem, bounding the number of n -permutations avoiding a fixed sub-permutation at indices corresponding to the edges of a hypergraph
- Achieved bounds for both random and deterministic avoidance hypergraphs
- In deterministic case, devised a hypergraph formulation of pattern-avoidance, enabling the use of the hypergraph containers method

WORK EXPERIENCE

Google Research

Student Researcher, 2024-2025

- Established Tree-Cal algorithm for calibration over high-dimensional sets
- Devised the first algorithm for linear swap regret minimization over arbitrary convex action sets
- Established algorithms for full swap regret minimization over convex action sets
- Studied generalizations of high-dimensional swap regret minimization to Blackwell approachability
- Worked to close the upper and lower bound on the optimal rate of sequential calibration

Microsoft Research

EconCS Research Extern, 2020

- Established that first-price multi-item auctions achieve a constant factor of the optimal revenue
- Devised ways of compressing the information in matching problems via agent classifications while maintaining welfare guarantees for the derived matchings
- Position intended for graduate students; was employed during undergrad; only 1 opening

Optiver US LLC

Quantitative Research and Trading Intern, 2018

- Created a machine learning model to predict volatility of the S&P500 following a day with abnormally high realized volatility
- Devised a strategy to adjust predicted volatility in response to a market data
- Traded S&P futures and options in a simulated environment using real market data
- Coded automated trading algorithms to compete against fellow interns in market making games

TEACHING EXPERIENCE

MIT Primes

Mentor, 2025

- Mentored two high school researchers in advanced theoretical and experimental projects
- With Michael Han, supervised a theoretical project investigating the existence of combinatorial graph structures known as “novelty games”, awarded as a **2026 Regeneron STS Scholar** and a **National Finalist in the S.-T. Yau High School Science Award**
- With Ashley Yu, guided an experimental project utilizing simulations to test a conjecture on the existence of a fast per-iterate swap regret minimization algorithm

- Provided weekly personalized mentorship, teaching known mathematical tools and fostering research curiosity

Proofs and Goofs

Educational Video Creator, 2024-

- Produced a series of [video lectures](#) introducing fundamental concepts in the theory of computer science in elementary terms
- Covered topics including linear programming duality and gradient descent
- Accompanied each lecture with an original math song as a fun reward for sticking through the lecture

MIT Math Learning Center

Tutor, 2019

- Helped undergraduates with coursework from classes spanning MIT's mathematics curriculum
- Provided individual attention to students struggling on a specific topic, helping them gain intuition
- Lectured groups of students on foundational topics

Awesome Math Summer Program

Teaching Assistant, 2015-16

- 4 three-week sessions at Cornell U. (twice), UC Berkeley, and Univ. of Puget Sound, WA
- Worked alongside instructor explaining high-level olympiad math contest techniques from geometry, combinatorics, and number theory
- Worked with 50 students each camp (ages 12-16) helping them work through problems and gain problem solving intuition during classes and office hours

SERVICE

Program Committee ALT 2027, Easy Peasy at EC 2026, COLT 2025, EC 2025, ALT 2025

Reviewer SODA 2027, NeurIPS 2026, NeurIPS 2025 (**Top Reviewer**), ITCS 2026, SODA 2026, FOCS 2025, COLT 2025, Journal of Mathematics of Operations Research 2025, EC 2025, ALT 2025, ITCS 2025, SODA 2025, FOCS 2024, FOCS 2023, Journal of Mathematics of Operations Research 2023, SICOMP 2022, FOCS 2020

AWARDS AND HONORS

Akamai Presidential Graduate Fellowship	2020
Honorable Mention, USA Junior Math Olympiad	(#12 nationally) 2014
Honorable Mention, USA Junior Math Olympiad	(#15 nationally) 2013
Harvard-MIT Math Tournament individual round	(#19 internationally) 2015
#1 Individual Scorer, NY State Math Tournament	(Curt Boddie Award) 2014
#1 Individual Scorer, NYC Math Tournament (NYCIML)	(3 years in a row) 2014-16