DANIEL HALPERN

300 Buchanan St, San Francisco, CA | +1 (607) 227–4045 | dhalpern@google.com | https://daniel-halpern.com

PROFESSIONAL APPOINTMENTS

Google Research — **Mountain View, CA**
Research Scientist — *Jul. 2025–*

EDUCATION

Harvard University — **Cambridge, MA**
Ph.D. in Computer Science — *Aug. 2020–Jun. 2025*
- Advisor: Ariel D. Procaccia

University of Toronto — **Toronto, ON**
B.Sc. in Computer Science with High Distinction — *Sep. 2016–Jun. 2020*

JOURNAL ARTICLES

Note on author order for all publications: (α) denotes alphabetical, (r) denotes random.

J4. AI Alignment from Social Choice Perspectives.
(α) D. Halpern, E. Micha, A. D. Procaccia, B. Schiffer, I. Shapira, and S. Zhang.
In *SIGecom Exchanges*, 2026.

J3. Representation with Incomplete Votes.
(α) D. Halpern, G. Kehne, A. D. Procaccia, J. Tucker-Foltz, and M. Wüthrich.
In *Theory and Decision*, 2026.

J2. Dynamic Fair Division with Partial Information.
(α) G. Benadè, D. Halpern, and A. Psomas.
In *Operations Research*, 2025.

J1. Tracking Truth with Liquid Democracy.
(α) A. Berinsky, D. Halpern, J. Y. Halpern, A. Jadbabaie, E. Mossel, A. D. Procaccia, and M. Revel.
In *Management Science*, 2025.

CONFERENCE PUBLICATIONS

C25. Truthful Fair Division under Stochastic Valuations.
(α) D. Halpern, A. Psomas, and S. Zhang.
In *Proceedings of the 26th ACM Conference on Economics and Computation (**EC**)*, 2026.

C24. Robust AI Evaluation through Maximal Lotteries.
H. Khalaf, S. Wang, D. Halpern, I. Shapira, F. P. Calmon, and A. D. Procaccia.
In *Proceedings of the 43rd International Conference on Machine Learning (**ICML**)*, 2026.

C23. Pairwise Calibrated Rewards for Pluralistic Alignment.
(α) D. Halpern, E. Micha, I. Shapira, and A. D. Procaccia.
In *Proceedings of the 39th Conference on Neural Information Processing Systems (**NeurIPS**)*, 2025.

C22. Online Envy Minimization and Multicolor Discrepancy: Equivalences and Separations.
(α) D. Halpern, A. Psomas, P. Verma, and D. Xie.
In *Proceedings of the 26th ACM Conference on Economics and Computation (**EC**)*, 2025.

C21. The Proportional Veto Principle for Approval Ballots.
(α) D. Halpern, A. D. Procaccia, and W. Suksompong.
In *Proceedings of the 34th International Joint Conference on Artificial Intelligence (**IJCAI**)*, 2025.

C20. Federated Assemblies.
(α) D. Halpern, A. D. Procaccia, E. Shapiro, and N. Talmon.
In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (**AAAI**)*, 2025.
★ Oral presentation (4.6% of submissions)

C19. Axioms for AI Alignment from Human Feedback.
(α) L. Ge, D. Halpern, E. Micha, A. D. Procaccia, I. Shapira, Y. Vorobeychik, and J. Wu.
In *Proceedings of the 38th Conference on Neural Information Processing Systems (**NeurIPS**)*, 2024.
★ Spotlight presentation (2.1% of submissions)

C18. Computing Voting Rules with Elicited Incomplete Votes.
(α) D. Halpern, S. Hossain, and J. Tucker-Foltz.
In *Proceedings of the 25th ACM Conference on Economics and Computation (**EC**)*, 2024.

C17. On the Existence of Envy-Free Allocations Beyond Additive Valuations.

(α) G. Benadè, D. Halpern, A. Psomas, and P. Verma.
In *Proceedings of the 25th ACM Conference on Economics and Computation (**EC**)*, 2024.

C16. Metric Distortion with Elicited Pairwise Comparisons.
(α) S. Ebadian, D. Halpern, and E. Micha.
In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (**IJCAI**)*, 2024.

C15. Optimal Engagement-Diversity Tradeoffs in Social Media.
(α) F. Baumman, D. Halpern, I. Rahwan, I. Shapira, A. D. Procaccia, and M. Wüthrich.
In *Proceedings of the 33rd ACM Web Conference (**WWW**)*, 2024.

C14. Strategyproof Voting under Correlated Beliefs.
(α) D. Halpern, R. Li, and A. D. Procaccia.
In *Proceedings of the 37th Conference on Neural Information Processing Systems (**NeurIPS**)*, 2023.

C13. Smoothed Analysis of Social Choice Revisited.
(α) B. Flanigan, D. Halpern, and A. Psomas.
In *Proceedings of the 19th Conference on Web and Internet Economics (**WINE**)*, 2023.

C12. In Defense of Liquid Democracy.
(α) D. Halpern, J. Y. Halpern, A. Jadbabaie, E. Mossel, A. D. Procaccia, and M. Revel.
In *Proceedings of the 24th ACM Conference on Economics and Computation (**EC**)*, 2023.

C11. Representation with Incomplete Votes.
(α) D. Halpern, G. Kehne, A. D. Procaccia, J. Tucker-Foltz, and M. Wüthrich.
In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (**AAAI**)*, 2023.

C10. Dynamic Fair Division with Partial Information.
(α) G. Benadè, D. Halpern, and A. Psomas.
In *Proceedings of the 36th Conference on Neural Information Processing Systems (**NeurIPS**)*, 2022.

C9. Liquid Democracy in Practice: An Empirical Analysis of its Epistemic Performance.
M. Revel, D. Halpern, A. Berinsky, and A. Jadbabaie.
In *Proceedings of the 2nd ACM conference on Equity and Access in Algorithms, Mechanisms, Optimization (**EAAMO**)*, 2022.

C8. Distortion in Voting with Top-t Preferences.
(α) A. Borodin, D. Halpern, M. Latifian, and N. Shah.
In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (**IJCAI**)*, 2022.

C7. Can Buyers Reveal for a Better Deal?.
(α) D. Halpern, G. Kehne, and J. Tucker-Foltz.
In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (**IJCAI**)*, 2022.

C6. How Many Representatives Do We Need? The Optimal Size of an Epistemic Congress.
(r) M. Revel, T. Lin, and D. Halpern.
In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (**AAAI**)*, 2022.

C5. Fair and Efficient Resource Allocation with Partial Information.
(α) D. Halpern and N. Shah.
In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (**IJCAI**)*, 2021.

C4. Aggregating Binary Judgments Ranked By Accuracy.
(α) D. Halpern, G. Kehne, D. Peters, A. D. Procaccia, N. Shah, and P. Skowron.
In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (**AAAI**)*, 2021.

C3. Fair Division with Binary Valuations: One Rule to Rule Them All.
(α) D. Halpern, A. D. Procaccia, A. Psomas, and N. Shah.
In *Proceedings of the 16th Conference on Web and Internet Economics (**WINE**)*, 2020.

C2. Resolving the Optimal Metric Distortion Conjecture.
(α) V. Gkatzelis, D. Halpern, and N. Shah.
In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (**FOCS**)*, 2020.
★ Invited to the EC 2021 plenary session: Highlights Beyond EC

C1. Fair Division with Subsidy.
(α) D. Halpern and N. Shah.
In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (**SAGT**)*, 2019.

WORKING PAPERS

W3. Proportionality in Ranking Compression.
(α) D. Halpern, X. Liu, E. Micha, Y. P. Ng, and W. Suksompong.

W2. Linear Social Choice with Few Queries: A Moment-Based Approach.
(α) L. Ge, D. Halpern, G. Kehne, and Y. Vorobeychik.

W1. Fair Division Under Inaccurate Preferences.
(α) T. Dang, D. Halpern, A. Makur, A. Psomas, J. Singh, and P. Verma.

SELECTED HONORS AND AWARDS

- *Siebel Scholarship* 2024
- *NSF Graduate Research Fellowship* 2021
- *University of Toronto Computer Science Undergraduate Research Award* 2020
- *Harold Willet Stewart Memorial Scholarship* 2020
- *Anna And Alex Beverly Memorial Fellowship* 2020
- *Samuel Beatty In Course Scholarship* 2019
- *C. L. Burton Scholarship For Mathematics and Physical Sciences* 2019
- *Dr. James A. & Connie P. Dickson Scholarship in Science & Mathematics* 2018
- *Alan Milne McCombie Scholarship* 2017
- *University of Toronto President's Scholars of Excellence Program* 2016

TEACHING EXPERIENCE

GEC Academy **Online**
Teaching Fellow *Summer 2024*
- Mathematics for Economics

Harvard University **Cambridge, MA**
Teaching Fellow *Spring 2022*
- Optimized Democracy (CS238)

University of Toronto **Toronto, ON**
Undergraduate Teaching Assistant *Spring 2020*
- Data Structures and Analysis (CSC263)
- Algorithm Design, Analysis & Complexity (CSC373)

SERVICE

PC Member: AAAI ('23, '24, '25), IJCAI ('23, '24), SAGT ('23), NeurIPS ('24)
Journal Reviewer: ARTINT ('21, '22, '24, '25), JAAMAS ('21, '21, '21, '22), MOR ('22, '23), MSS ('21, '22, '23), TEAC ('25)
Subreviewer: AAMAS ('25), EAAMO ('22), SAGT ('21), SODA ('24), STOC ('25)

INVITED TALKS

Social Choice and Welfare Seminar September, 2025
A Social Choice Perspective on AI Alignment

National University of Singapore Workshop on Algorithmics of Fair Division and Social Choice December, 2024
Aggregating Preferences with Limited Queries

Cornell Theory Seminar November, 2024
Aggregating Preferences with Limited Queries

FOCS Workshop on Distortion in Social Choice October, 2024
Optimal Randomized Utilitarian Distortion

INFORMS Annual Meeting October, 2024
Tracking Truth with Liquid Democracy

University of Chicago Computer Science Colloquium October, 2024
Aggregating Preferences with Limited Queries

Carnegie Mellon Formal Epistemology Lecture Series September, 2024
Aggregating Preferences with Limited Queries

Oxford Algorithmic Game Theory Seminar June, 2024
Computing Voting Rules with Elicited Incomplete Votes

MSRI/SLMath Social Choice Seminar November, 2023
Resolving the Optimal Metric Distortion Conjecture

INFORMS Annual Meeting October, 2023
Representation with Incomplete Votes

HalpernFest at Cornell University June, 2023
In Defense of Liquid Democracy

McGill Bellairs Workshop on Multi-Agent Systems March, 2023
Representation with Incomplete Votes

COMSOC Video Seminar February, 2023
Representation with Incomplete Votes

LAMSADE Mini-Workshop on Cooperative Games, Social Choice, and Fair Division September, 2022
In Defense of Liquid Democracy

Highlights Beyond EC July, 2021
Resolving the Optimal Metric Distortion Conjecture

Drexel Theory Seminar May, 2021
Fair and Efficient Resource Allocation with Partial Information

Cornell Theory Seminar November, 2020
Resolving the Optimal Metric Distortion Conjecture

Harvard EconCS Seminar September, 2020
Resolving the Optimal Metric Distortion Conjecture