

Benjamin Laufer

Website: bendlaufer.github.io

Email: bdl56@cornell.edu

Phone: +1-917-628-8880

EMPLOYMENT

- **University of Washington Information School** Seattle, WA
Incoming Assistant Professor *September 2026 – Present*
- **Microsoft** San Francisco, CA
Senior Applied Scientist *July 2026 – Present*

EDUCATION

- **Cornell Tech** New York, NY
Ph.D. Candidate, Bowers School of Computing and Information Sciences *2021 – 2026*
 - **Advisors:** Jon Kleinberg, Helen Nissenbaum*M.S., 2024*
- **Princeton University** Princeton, NJ
B.S. in Engineering, Operations Research and Financial Engineering, Cum Laude *2015 – 2019*

HONORS

- **Rising Star in Data Science:** University of Chicago, UCSD, and Stanford program for ‘rising stars’ in Data Science (2025)
- **Seth Bonder Foundation Scholar:** INFORMS Doctoral Student Colloquium award (2025)
- **Cosmos Institute × FIRE quick grant:** Funding for work on the “AI Ecosystem” (2025)
- **Georgia Tech ISyE-MS&E-IOE Rising Star:** Georgia Tech program for ‘rising stars’ in Industrial and Systems Engineering, Management Science and Engineering, and Industrial and Organizational Engineering (2025)
- **Rising Star in MS&E:** Stanford program for ‘rising stars’ in Management Science & Engineering (2024)
- **LinkedIn Bowers CIS PhD Fellowship:** Full support plus discretionary funding (2023-2024) [[link1](#), [link2](#)]
- **Meta PhD Fellowship Finalist:** Runner-up for a two-year, fully-funded fellowship in Computational Social Science. Over 3200 applicants internationally. (2023) [[link](#)]
- **Digital Life Initiative Doctoral Fellowship:** Yearlong discretionary funding for PhD students interested in ethics, politics, and quality of life in digital societies. (2022)
- **Fulbright Fellowship Alternate:** MSc in computing at Imperial College London, UK (2021)
- **Lime Hackathon Winner:** 1st place prize in Lime’s annual company-wide hackathon. Project entitled “Head First: Real-Time Helmet Detector using Computer Vision” (2020)
- **Best Senior Thesis in Urban Studies:** Princeton Urban Studies department thesis prize (2019) [[link](#)]
- **Kenneth H. Condit Graduation Prize:** Awarded to a graduating senior at Princeton who has exhibited excellence in leadership, academic achievement and community service (2019) [[link](#)]
- **Departmental Distinction:** Princeton ORFE Honors (2019)
- **Sigma Xi:** Nominated to the Scientific Research Honors Society upon graduating (2019)
- **Fulbright Fellowship, Laos:** 1 year full support (2019; declined)
- **Princeton SEAS Research Fund:** Funding for senior thesis research (2018-2019)
- **Parkhurst Research Fund:** Funding for senior thesis research (2018-2019)
- **Yale Law School Arthur Liman Fellowship:** Award supporting a summer internship in the public interest and colloquium attendance (2018) [[link](#)]
- **Princeton SEAS Research Fund:** Funding for junior independent research (2018)
- **Princeton Environmental Institute Fellowship:** Award for environmental research (2017) [[link](#)]
- **Princeton International Internship Program:** Award for international summer research (2016)

PAPERS

- **Chris Hays, Benjamin Laufer, Solon Barocas, Manish Raghavan:** “Statistical Guarantees in the Search for Less Discriminatory Algorithms,” International Conference on Learning Representations (ICLR) 2026, also appeared at *CS & Law 2026* (non-archival), and on [ArXiv](#), 2025.
- **Benjamin Laufer, Hamidah Oderinwale, Jon Kleinberg:** “Anatomy of a Machine Learning Ecosystem: 2 Million Models on Hugging Face,” *Revision at Proceedings of National Academy of Sciences (PNAS)*, and on [ArXiv](#), 2025.
- **Benjamin Laufer, Jon Kleinberg:** “Measuring Rule-Following in Language Models” *ICML 2025 Workshop on Assessing World Models*, 2025.
- **Tori Qiu, Benjamin Laufer, Jon Kleinberg, Hoda Heidari:** “On the Economic Impacts of AI Openness Regulation: A Game-theoretic model of Performance and Openness Implications,” *Advances in Neural Information Processing Systems (NeurIPS)*, and on [ArXiv](#), 2025.
- **Benjamin Laufer, Hoda Heidari, Jon Kleinberg:** “The Backfiring Effect of Weak AI Safety Regulation” *Revision at Proceedings of National Academy of Sciences (PNAS)*, previously appeared at *ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)* (non-archival), and [ArXiv](#), 2025.
- **Benjamin Laufer, Manish Raghavan, Solon Barocas:** “What Constitutes a Less Discriminatory Algorithm?” *ACM CS & Law ’25 [presentation]*, earlier versions appeared at *NeurIPS RegML and NeurIPS AFME workshops [oral spotlight]*, and on [ArXiv](#), 2024.
- **Benjamin Laufer, Jon Kleinberg, Hoda Heidari:** “Fine-Tuning Games: Bargaining and Adaptation for General-Purpose Models,” *The Web Conference 2024 (WWW) [oral spotlight]*, and on [ArXiv](#).
- **Benjamin Laufer, Helen Nissenbaum:** “Algorithmic Displacement of Social Trust,” *Knight Institute, Symposium and Essay Series on Algorithmic Amplification and Society*, 2023.
- **Benjamin Laufer, Jon Kleinberg, Karen Levy, Helen Nissenbaum:** “Strategic Evaluation: Subjects, Evaluators, and Society,” *ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*, 2023.
- **Benjamin Laufer, Thomas Gilbert, Helen Nissenbaum:** “Optimization’s Neglected Normative Commitments.” *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2023.
- **Benjamin Laufer, Manish Raghavan, Solon Barocas:** “Are computational interventions to advance fair lending robust to different modeling choices about the nature of lending?” *NeurIPS AFT workshop*, 2023.
- **Benjamin Laufer, Emma Pierson, Nikhil Garg:** “End-to-end Auditing for Decision Pipelines.” *ICML Workshop on Responsible Decision Making in Dynamic Environments (RDMDE)*, 2022.
- **Benjamin Laufer, Niko Grupen:** “Collective Obfuscation and Crowdsourcing.” *KDD MIS2-TrueFact and ICML DisCoML workshops*, 2022.
- **Benjamin Laufer, Sameer Jain, A. Feder Cooper, Jon Kleinberg, Hoda Heidari:** “Four Years of FAccT: A Reflexive, Mixed-Methods Analysis of Research Contributions, Shortcomings, and Future Prospects.” *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2022.
- **A. Feder Cooper, Emanuel Moss, Benjamin Laufer, Helen Nissenbaum:** “Accountability in an Algorithmic Society.” *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2022.
- **Benjamin Laufer:** “Abandoning Criminal Risk and Recidivism: On Dangerous Goals in ML Scoring-Decision Systems.” *NeurIPS Resistance AI Workshop*, 2020.
- **Benjamin Laufer:** “Compounding Injustice: History and Prediction in Carceral Decision-Making.” Undergraduate Thesis, [ArXiv](#), [Princeton Mudd Manuscript Library](#), April 2019.
- **Benjamin Laufer:** “Data Analytics and Machine Learning for Environmental Protection: Targeting Air Inspections.” Junior Research Paper, *Princeton ORFE*, January 2018.

TALKS AND EVENTS

- **Johns Hopkins University Computer Science Seminar:** “AI Ecosystems: Structure, Strategy, Risk and Regulation,” Baltimore, MD, January 2026.
- **University of Chicago Data Science Institute and Computer Science Joint Colloquium:** “AI Ecosystems: Structure, Strategy, Risk and Regulation,” Chicago, IL, January 2026.
- **University of Washington iSchool:** “AI Ecosystems: Structure, Strategy, Risk and Regulation,” Seattle, WA, January 2026.
- **Columbia University AI/ML Business Analytics Conference:** “Anatomy of a Machine Learning Ecosystem,” New York, NY, December 2025.
- **University of North Carolina Chapel Hill School of Information and Library Sciences:** “AI Ecosystems: Structure, Strategy, Risk and Regulation,” Chapel Hill, NC, November 2025.

- **Rising Stars in Data Science:** Invited workshop and talk, “AI Ecosystems: Structure, Strategy, and Regulation,” Stanford, CA, November 2025.
- **EAAMO '25:** Proceedings paper and talk, “The Backfiring Effect of Weak AI Safety Regulation,” Pittsburgh, PA, November 2025.
- **NYU Stern Technology, Operations and Statistics group:** “AI Ecosystems: Structure, Strategy, Risk and Regulation,” New York, NY, October 2025.
- **Carnegie Mellon Workshop on Human-AI Complementarity for Decision-Making:** NSF AI Institute for Societal Decision Making invited workshop, and poster presentation, “Game-theoretic model of AI Safety, Openness and Regulation,” Pittsburgh, PA, September 2025.
- **Toyota Technical Institute Chicago workshop on “Incentives for Collaborative Learning and Data Sharing”:** Invited workshop, and poster presentation, “Anatomy of an AI Ecosystem,” Chicago, IL, August 2025.
- **Northeastern AIDE summer school in Data Ethics:** Eight week program, Boston, MA, June-July, 2025.
- **Princeton University Fairness in AI Seminar Series:** Invited talk, “What Constitutes a Less Discriminatory Algorithm?” Princeton, NJ, April, 2025.
- **IMT Atlantique Data Science Research Seminar:** Invited talk, “Searching for accurate and fair classifiers,” Brest, France, April, 2025.
- **Université de Bretagne Occidentale Summer School:** Invited speaker, “Social Networks and the Integrity of Democratic Processes,” Brest, France, April, 2025.
- **MIT Human and Machine Decisions Group:** Visiting Speaker, “Regulation along the AI Development Pipeline for Fairness, Safety and Related Goals” Boston, MA, April, 2025.
- **Simons Institute Theory of Computing for Fairness (TOC4Fairness) Seminar:** Talk, “Regulation along the AI Development Pipeline for Fairness, Safety and Related Goals” Virtual, April, 2025.
- **Northeastern Network Science Institute:** Visiting Speaker, “Regulation along the AI Development Pipeline for Fairness, Safety and Related Goals” Boston, MA, April, 2025.
- **CSLaw '25:** Proceedings paper and talk, “What Constitutes a Less Discriminatory Alternative?” Munich, Germany, March, 2025.
- **Technical University Munich (TUM) School of Social Sciences and Technology:** Invited talk, “Regulation along the AI Development Pipeline for Fairness, Safety and other Desirables” Munich, Germany, March, 2025.
- **NYU Responsible AI Course:** Invited lecture, “What Constitutes a Less Discriminatory Algorithm?” New York, NY, February, 2025.
- **Simons Collaboration on the Theory of Algorithmic Fairness:** Invited participant, New York, NY, February, 2025.
- **Inaugural Conference of the International Association for Safe and Ethical AI:** Invited participant, ‘A game-theoretic model of fine-tuning and its implications for regulating AI’ Paris, France, February, 2025.
- **Workshop on Defining Holistic Private Data Science for Practice:** Invited participant and lightning talk, “Fine-Tuning Games,” San Diego, California, January, 2025.
- **Meta AI and Society Seminar:** Invited talk, “Modeling the Development of ML Applications and their Ethical Implications,” Virtual, December, 2024.
- **NeurIPS '24:** AFME Workshop spotlight oral presentation, “Fundamental Limits in the Search for Less Discriminatory Algorithms—and How to Avoid Them,” Vancouver, CA (given virtually), December, 2024.
- **Meta Responsible AI Group Visit:** Invited talk, “Fine-Tuning Games: Modeling the Ecosystem of AI Application Development,” New York, NY, November, 2024.
- **Johns Hopkins Berman Institute of Bioethics:** Invited participant and travel grant awardee, “AI Ethics and Governance Symposium,” Washington, D.C., October, 2024.
- **MLESI '24:** Invited participant, Machine Learning and Economics Summer Institute, Chicago, IL, August, 2024.
- **Cornell Milstein Fellowship Program:** Lecture, “Can We Find a Less Discriminatory Algorithm?” July, 2024.
- **WALE '24:** Invited participant and poster presentation, “Fine-Tuning Games,” Workshop on Algorithms, Learning, and Economics, Kefalonia, Greece, June, 2024.
- **BIRS Workshop:** Invited participant and presenter, “Bridging Prediction and Intervention Problems in the Social Sciences,” Banff, British Columbia, Canada, June, 2024.
- **WWW '24:** Proceedings paper and oral presentation, “Fine-tuning Games: Bargaining and Adaptation for General-Purpose Models,” Singapore, Singapore, May, 2024.
- **Stanford Management Science & Engineering:** Rising stars event, poster and lightning talk, “Fine-tuning Games: Bargaining and Adaptation for General-Purpose Models,” Palo Alto, California, April, 2024.

- **Simons Collaboration on the Theory of Algorithmic Fairness:** Invited participant, New York, NY, February, 2024.
- **NYU Center for Data Science (CDS) Seminar:** Visit and talk, “Fine-tuning Games: Bargaining and Adaptation for General-Purpose Models,” New York, NY, January, 2024.
- **Harvard EconCS Seminar:** Visit and talk, “Fine-tuning Games: Bargaining and Adaptation for General-Purpose Models,” Cambridge, MA, November, 2023. [\[link\]](#)
- **MIT AI and Society Group:** Visit and talk, “Fine-tuning Games: Bargaining and Adaptation for General-Purpose Models,” Cambridge, MA, November, 2023.
- **EAAMO '23:** Proceedings paper and poster presentation, “Strategic Evaluation,” Boston, MA, October, 2023.
- **Symposium on Applications of Contextual Integrity:** Proceedings talk, “Party’s Over: First-Party Surveillance Through the Lens of Contextual Integrity,” Toronto, Canada, September, 2023.
- **Industria 5.0:** Presentation, “Algorithmic Amplification,” New York, NY, September, 2023.
- **BIRS Workshop:** Invited participant, “Contextual Integrity for Differential Privacy” Kelowna, British Columbia, Canada, August, 2023.
- **Cornell Milstein Fellowship Program:** Lecture, “When is Algorithmic Amplification Problematic?” July, 2023.
- **Cornell AI, Policy and Practice:** “First-party Surveillance” July, 2023.
- **FACCT '23:** Proceedings talk, “Optimization’s Neglected Normative Commitments,” Chicago, IL, June, 2023.
- **PLSC:** Workshop participant as paper-writer and discussant, “First-Party Surveillance,” Boulder, CO, June, 2023.
- **Columbia Knight Institute:** Symposium entitled ‘Optimizing for What? Algorithmic Amplification and Society.’ Presenting work, “When is Algorithmic Amplification Wrongful?” New York, NY, April, 2023. [\[event\]](#), [\[video\]](#)
- **Cornell Digital Life Initiative Seminar:** “When is Algorithmic Amplification Wrongful?” New York, NY, March, 2023. [\[event\]](#), [\[video\]](#)
- **Cornell AI, Policy and Practice:** “When is Algorithmic Amplification Wrongful?” February, 2023.
- **c3.ai:** Workshop on Data, Learning and Markets Workshop, “Disparate Impact over a Sequence of Capacity-Constrained Subset Selections,” Poster Session, Champaign, IL, October, 2022. [\[event\]](#)
- **EAAMO '22:** Queer in AI panel, October, 2022. [\[event\]](#)
- **Cornell AI, Policy and Practice:** “Accounting for Unforeseeable Machine Learning Harms via Economic Models of Product Liability,” October, 2022.
- **KDD '22:** MIS2-TrueFact Workshop talk, “Collective Obfuscation and Crowdsourcing,” August, 2022. [\[event\]](#)
- **ICML '22:** Workshop on Disinformation Countermeasures and Machine Learning, “Collective Obfuscation and Crowdsourcing,” July, 2022. [\[event\]](#)
- **ICML '22:** Workshop on Responsible Decision-Making in Dynamic Environments, “End-to-end Auditing of Decision Pipelines,” July, 2022. [\[video\]](#)
- **Economics & Computing:** Poster Session, “End-to-end Auditing of Decision Pipelines,” July, 2022. [\[link\]](#)
- **FACCT '22:** Proceedings talk, “Four Years of FACCT: A Reflexive, Mixed-Methods Analysis of Research Contributions, Shortcomings, and Future Prospects,” June, 2022. [\[video\]](#)
- **Cornell Information Science Colloquium:** “Neglected Normative Assumptions in Applied Optimization,” April, 2022.
- **Cornell Tech Tree Day:** “Auditing Efficiency and Equity throughout an Operational Workflow,” April, 2022.
- **Cornell AI, Policy and Practice:** “Crowdsourcing, Allocation, and Equity in NYC Parks Department Policies,” April, 2022.
- **Cornell Tech Course Lecture:** Data Science in the Wild Guest Lecture, “Data Visualization,” February, 2022.
- **CUNY Experimental Philosophy Group:** “On Accountability,” November, 2021.
- **Cornell AI, Policy and Practice:** “Accountable Metrics,” November, 2021.
- **NeurIPS '20:** Workshop on Resistance AI, “Dangerous Goals in ML Scoring-Decision Systems,” December 2020. [\[event\]](#)
- **Princeton ORFE Department:** Invited lecture on algorithmic fairness in ORF473-FinTech Lending (prof: Margaret Holen), May 2020.
- **Princeton Reunions:** “Princeton in Prison: The Petey Greene Tutors” Panel, June 2019.
- **Princeton Urban Studies Department:** Research Symposium, May 2019.
- **High Meadows Environmental Institute:** Research Symposium, April 2019.
- **Yale Law Public Interest Colloquium:** Invited as a fellow, April 2018. [\[event\]](#)

WORK EXPERIENCE

- **Northeastern Ethics Institute** Boston, MA
Research Intern Jun 2025 - Aug 2025 (3 months)
- **Microsoft Research (MSR)** New York, NY
Research Intern Jun 2024 - Sep 2024 (4 months)
- **Lime** San Francisco, CA
Data Scientist II (prev. Data Scientist, Data Analyst Intern) 2019 - 2021 (2 years)
- **National Center for Access to Justice** New York, NY
Arthur Liman Research Fellow, Yale Law School Jun 2018 - Aug 2018 (3 months)
- **U.S. Environmental Protection Agency (EPA)** New York, NY
Data Science Intern Jun 2017 - Jan 2018 (8 months)
- **Princeton Plasma Physics Laboratory (PPPL)** Princeton, NJ
Research Assistant Oct 2016 - May 2017 (8 months)
- **International Water Management Institute** Accra, Ghana
Intern Jun 2016 - Aug 2016 (3 months)
- **Princeton High School Debate Team** Princeton, New Jersey
Private Debate Coach September 2015 - June 2016 (10 months)
- **National Symposium for Debate** Minneapolis, Minnesota
Debate Teacher Jul 2015 - Aug 2015 (2 months)

SERVICE AND VOLUNTEERING

- **Academic Conferences and Journals**
Program Committee Reviewer
 - **FAccT**: 2023, 2024, 2025, 2026
 - **AIES**: 2023, 2024, 2025
 - **NeurIPS**: 2025
 - **CHI**: 2023
 - **WWW**: 2023, 2024, 2025
 - **EC**: 2026
 - **Journal of Responsible Computing**: Reviewer (x3)
- **Cornell University** New York, NY
Information Science Graduate Student Association (ISGSA) and PhDs at Cornell Tech (PACT) 2023-2024
 - **Faculty Hiring Representative**: Elected position; responsible for attending faculty hiring meetings and representing student interests.
- **Princeton University** Princeton, NJ
Alumni volunteer 2020-present
 - **Volunteer alumni interviewer**: Serve as an interviewer for high school students applying to college.
- **UCLA COVID Behind Bars Research Group** Los Angeles, CA
Data Science Researcher 2020
 - **Data reporting for COVID-19 cases in ICE detention facilities**: Creating a single location for COVID-19 case and death counts in ICE facilities.
- **Petey Greene Program** Princeton, NJ
Tutor and Undergraduate President 2016 - 2019 (4 years)
 - **Bringing Education into Prisons**: Tutored inside NJ correctional facilities weekly for four years.
 - **Undergraduate President**: Led the team of 100+ tutors and organized educational programs and trainings for tutors.

PROGRAMMING SKILLS

- **Languages:** Python, SQL, MATLAB, R, Java, C, html, CSS, JS, STATA
- **Technologies:** AWS, Sagemaker, React, Tableau

REFERENCES

- **Helen Nissenbaum, Andrew H. & Ann R. Tisch Professor, Information Science, Cornell Tech:** hn288 [at] cornell.edu
- **Jon Kleinberg, Tisch University Professor, Information Science and Computer Science, Cornell University:** kleinberg [at] cornell.edu
- **Solon Barocas, Principal Researcher, Microsoft Research:** solon.barocas [at] microsoft.com
- **Hoda Heidari, K&L Gates Career Development Assistant Professor in Ethics and Computational Technologies, Caregie Mellon University:** hheidari [at] cmu.edu
- **Manish Raghavan, Drew Houston (2005) Career Development Professor, MIT:** mragh [at] mit.edu