

Vincent Conitzer

Office: GHC 7113

Computer Science Department

Carnegie Mellon University

5000 Forbes Avenue

Pittsburgh, PA 15213, USA

Office phone: (412) 268-9612

conitzer@cs.cmu.edu

<https://www.cs.cmu.edu/~conitzer/>

Academic Employment (since Ph.D.)

- **Carnegie Mellon University:**
 - **Director of the Foundations of Cooperative AI Lab (FOCAL).** 2022-present.
 - **Professor of Computer Science.** 2022-present.
 - ... with affiliate/courtesy appointments in **Machine Learning, Philosophy, and the Tepper School of Business.** 2022-present.
- **University of Oxford:**
 - **Head of Technical AI Engagement, Institute for Ethics in AI.**
 - **Professor of Computer Science and Philosophy.**
 - **Visiting Fellow at Pembroke College.**
- **Duke University:**
 - **Kimberly J. Jenkins Distinguished University Professor of New Technologies.** 2015-2022.
 - **Professor of Computer Science.** (Prior to the Jenkins chair, was also Arts and Sciences Professor of CS and Sally Dalton Robinson Professor of CS.) 2011-2022.¹
 - **Professor of Economics.** 2011-2022.
 - **Professor of Philosophy.** 2015-2022.
 - **Assistant Professor of Computer Science.** 2006-2011.
 - **Assistant Professor of Economics.** 2006-2011.

Extended Visits

- **UC Berkeley Simons Institute for the Theory of Computing** program on Economics and Computation. Fall 2015.
- **Centrum Wiskunde & Informatica (CWI) Amsterdam** (the national research institute for mathematics and computer science in the Netherlands), Algorithms, Combinatorics and Optimization group. Spring 2011.

Education

¹Double promotion from Assistant Professor to full Professor, skipping the rank of Associate Professor. Youngest full Professor at Duke (as recently as June 19, 2014).

- **Ph.D., Computer Science, Carnegie Mellon University, 2006.**
Dissertation title: *Computational Aspects of Preference Aggregation*. Committee: Tuomas Sandholm (chair), Avrim Blum, Craig Boutilier, Tom Mitchell, Christos Papadimitriou.
- **M.S., Computer Science, Carnegie Mellon University, 2003.**
- **A.B., Applied Mathematics, Harvard University, 2001.**
Subconcentrations in Computer Science and Economics.

Awards and Distinctions

Other group members' awards such as student dissertation awards are not listed here, but are sometimes listed with the general description of the group member given further below.

- Best Talk Award, COMSOC 2025.
Paper: Berker, Casacuberta, Robinson, Ong, Conitzer, Elkind, "From Independence of Clones to Composition Consistency: A Hierarchy of Barriers to Strategic Nomination." Talk by Emin Berker.
- Best Poster Award, AAAI 2025.
Paper: Tewolde, Zhang, Oesterheld, Sandholm, and Conitzer, "Computing Game Symmetries and Equilibria That Respect Them." Poster by Emanuel Tewolde.
- Distinguished Senior Program Committee Member, AAMAS 2025.
- Honorable Mention for Best Paper Award, HCOMP 2022.
Paper: Jecmen, Zhang, Liu, Fang, Conitzer, and Shah, "Near-Optimal Reviewer Splitting in Two-Phase Paper Reviewing and Conference Experiment Design."
- Outstanding Paper Award, workshop on ML Evaluation Standards at ICLR 2022.
Paper: Jecmen, Shah, Fang, and Conitzer, "Tradeoffs in Preventing Manipulation in Paper Bidding for Reviewer Assignment."
- Our paper was listed under Oxford University Press' "Best of Philosophy" for 2021.
Paper: Oesterheld and Conitzer, "Extracting Money from Causal Decision Theorists."
- IFAAMAS Influential Paper Award, 2022.
Paper: Conitzer and Sandholm, "Computing the optimal strategy to commit to."
- ACM/SIGAI Autonomous Agents Research Award, 2021.
- Top 5% of all undergraduate instructors teaching in the Natural Sciences at Duke, Fall 2020 (Computational Microeconomics).
- Best Blue Sky Idea Award, HCOMP 2020.
Paper: "Using Human Cognitive Limitations to Enable New Systems."
- Meritorious Prize Winner (among top 7 entries out of 800+ submissions) in the NSF 2026 Idea Machine competition, 2020.
- Fellow of the Association for Computing Machinery (ACM), 2019.
- Fellow of the Association for the Advancement of Artificial Intelligence (AAAI), 2019.
- AMiner Most Influential Scholar in Artificial Intelligence (ranked 10th), 2018 (announced in 2019).
- Co-author on paper receiving the Outstanding Student Paper Honorable Mention, AAAI 2018.
Paper: Freedman, Schaich Borg, Sinnott-Armstrong, Dickerson, and Conitzer, "Adapting a Kidney Exchange Algorithm to Align with Human Values."

- Top 5% of all undergraduate instructors teaching in the Natural Sciences at Duke, Fall 2017 (Artificial Intelligence).
- First place, New Yorker Cartoon Caption Contest, January 30, 2017.
- Computing Community Consortium Blue Sky Award, 2017.
Paper: Vincent Conitzer, Walter Sinnott-Armstrong, Jana Schaich Borg, Yuan Deng, and Max Kramer. Moral Decision Making Frameworks for Artificial Intelligence.
- AMiner Most Influential Scholar in Artificial Intelligence (ranked 26th), 2016
- Guggenheim Fellowship, 2015.
- Outstanding Senior Program Committee Member Award, AAAI 2015.
- Social Choice and Welfare Prize, 2014.
- Kavli Fellow, U.S. National Academy of Sciences (NAS) Kavli Frontiers of Science, 2012.
- Presidential Early Career Award for Scientists and Engineers (PECASE), 2011 (announced 2012).
- Thomas Langford Lectureship Award, 2012.
- Duke Bass Society of Fellows, 2011-present.
- The IJCAI Computers and Thought Award, 2011.
- IEEE Intelligent Systems' "AI's 10 to Watch" (Jan./Feb. 2011 issue).
- Runner-up for the Best SPC member award, AAMAS 2011.
- National Science Foundation CAREER Award, 2010.
- Runner-up for the Best SPC member award, AAMAS 2010.
- Top 5% of all undergraduate instructors at Duke, Fall 2009 (Computational Microeconomics).
- Outstanding Paper Award, AAAI 2008.
Paper: Wagman and Conitzer, "Optimal False-Name-Proof Voting Rules with Costly Voting."
- Co-author on the winner of the Pragnesh Jay Modi Best Student Paper Award, AAMAS 2008.
Paper: Ohta, Conitzer, Satoh, Iwasaki, and Yokoo, "Anonymity-Proof Shapley Value: Extending Shapley Value for Coalitional Games in Open Environments."
- Alfred P. Sloan Research Fellowship, 2008.
- Honorable Mention for the 2007 ACM Doctoral Dissertation Award.
- The 2006 IFAAMAS Victor Lesser Distinguished Dissertation Award (inaugural award, presented at AAMAS 2007).
- The Best Program Committee Member Award, AAMAS 2006.
- Runner-up for the Best Student Paper Award, AAMAS 2006.
Paper: Conitzer and Sandholm, "A Technique for Reducing Normal-Form Games to Compute a Nash Equilibrium."
- IBM Ph.D. Fellowship, 2005.
- Awarded Paper, JSAI 2005.
Paper: Yokoo, Conitzer, Sandholm, Ohta, and Iwasaki, "A New Solution Concept for Coalitional Games in Open Anonymous Environments."

- Phi Beta Kappa (Alpha Iota of Massachusetts).
- Tenth place, silver medalist, ACM Collegiate Programming Contest World Finals, 2001. (Member of Harvard team.)
- Top 100 ranking, 61st William Lowell Putnam Mathematical Competition, 2000.
- Twelfth place, Certamen Ciceronianum Arpinas (international Latin translation contest for high school students), 1995.

Publications

Dissertation

1. Vincent Conitzer. Computational Aspects of Preference Aggregation. Ph.D. Dissertation. Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, July 2006. Committee: Tuomas Sandholm (chair), Avrim Blum, Craig Boutilier, Tom Mitchell, Christos Papadimitriou. Available as technical report CMU-CS-06-145. (Winner of the 2006 IFAAMAS Victor Lesser Distinguished Dissertation Award and an Honorable Mention for the 2007 ACM Doctoral Dissertation Award.)

Books

2. Jana Schaich Borg, Walter Sinnott-Armstrong, and Vincent Conitzer. *Moral AI And How We Get There*. Pelican Books, February 2024. On Amazon, was, at various points, the #1 new release in Science & Maths Ethics, the #1 new release in Ethics & Morality, and the #1 best seller in Science & Maths Ethics. Also listed by Nature in books in brief (“five of the best science picks”).
1. Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (eds.). *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

Book chapters

10. Vincent Conitzer. What Would It Look Like to Align Humans with Ants? To appear as Chapter 17 in Nyholm, Sven, Kasirzadeh, Atoosa, & Zerilli, John (eds.) (2026): *Contemporary Debates in the Ethics of Artificial Intelligence*. Hoboken: Wiley-Blackwell. Also available as PhilSci 26351.
9. Lok Chan, Walter Sinnott-Armstrong, Jana Schaich Borg, and Vincent Conitzer. Should Responsibility Affect Who Gets a Kidney? In Ben Davies, Neil Levy, Gabriel De Marco, and Julian Savulescu (editors), *Responsibility and Healthcare*, chapter 1, pages 35-60, Oxford University Press, 2024.
8. Markus Brill and Vincent Conitzer. Strategic Voting and Strategic Candidacy. In S. Kurz, N. Maaser, and A. Mayer, editors, *Advances in Collective Decision Making: Interdisciplinary Perspectives for the 21st Century*, Studies in Choice and Welfare, pages 69-84. Springer, 2023. See also conference paper 110 below.
7. Walter Sinnott-Armstrong and Vincent Conitzer. How Much Moral Status Could AI Ever Achieve? Chapter 16 (pages 269-289) in *Rethinking Moral Status*, Clarke, S., Zohny, H., and Savulescu, J. (eds.), Oxford University Press, 2021.
6. Vincent Conitzer and Rupert Freeman. Algorithmically Driven Shared Ownership Economies. Chapter in *Future of Economic Design*, pp. 275-285, Springer, 2019.

5. Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. Introduction to Computational Social Choice. To appear as Chapter 1 in *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia (eds.), Cambridge University Press, 2016.
4. Vincent Conitzer and Toby Walsh. Barriers to Manipulation in Voting. To appear as Chapter 6 in *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia (eds.), Cambridge University Press, 2016.
3. Felix Brandt, Vincent Conitzer, and Ulle Endriss. Computational Social Choice. Chapter in G. Weiss (Ed.), *Multiagent Systems*, pp. 213-283, MIT Press, March 2013.
2. Vincent Conitzer. Metareasoning as a Formal Computational Problem. Appears as Chapter 8 in *Metareasoning: Thinking about Thinking*, Michael Cox and Anita Raja (editors), MIT Press, 2011.
1. Vincent Conitzer. Auction Protocols. Appears as Chapter 16 in the *CRC Algorithms and Theory of Computation Handbook, Second Edition, Volume 2: Special Topics and Techniques*, Mikhail Atallah and Marina Blanton (editors), 2010.

Journal papers

54. Vincent Conitzer. A Visual Proof of the Parallelogram Law. *The American Mathematical Monthly*, to appear. DOI: 10.1080/00029890.2025.2588111.
53. Ryan Liu, Steven Jecmen, Vincent Conitzer, Fei Fang, and Nihar B. Shah. Testing for reviewer anchoring in peer review: A randomized controlled trial. *PLOS ONE*, published November 18, 2024.
52. Eric Horvitz, Vincent Conitzer, Sheila McIlraith, and Peter Stone. Now, Later, and Lasting: Ten Priorities for AI Research, Policy, and Practice. *Communications of the ACM*, to appear. (Accepted December 2023.)
51. Vincent Conitzer, Gillian Hadfield, and Shannon Vallor. Technical Perspective: The Impact of Auditing for Algorithmic Bias. *Communications of the ACM*, Volume 66, Issue 1, January 2023, pp. 100.
50. Caspar Oesterheld and Vincent Conitzer. Safe Pareto improvements for delegated game playing. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, Volume 36, Article Number 46, 2022. Supersedes conference paper 155 below.
49. Edmond Awad, Sydney Levine, Michael Anderson, Susan Leigh Anderson, Vincent Conitzer, M.J. Crockett, Jim A.C. Everett, Theodoros Evgeniou, Alison Gopnik, Julian C. Jamison, Tae Wan Kim, S. Matthew Liao, Michelle N. Meyer, John Mikhail, Kweku Opoku-Agyemang, Jana Schaich Borg, Juliana Schroeder, Walter Sinnott-Armstrong, Marija Slavkovik, Josh B. Tenenbaum. Computational Ethics. *Trends in Cognitive Sciences*, Volume 26, Issue 5, May 2022, Pages 388-405.
48. Lok Chan, Jana Schaich Borg, Vincent Conitzer, Dominic Wilkinson, Julian Savulescu, Hazem Zohny, and Walter Sinnott-Armstrong. Which features of patients are morally relevant in ventilator triage? A survey of the UK public. *BMC Medical Ethics*, Volume 23, Article Number 33, 2022.
47. Masoud Afnan, Michael Anis Mihdi Afnan, Yanhe Liu, Julian Savulescu, Abhishek Mishra, Vincent Conitzer, Cynthia Rudin. Data solidarity for machine learning for embryo selection: a call for the creation of an open access repository of embryo data. *Reproductive BioMedicine Online*, DOI: 10.1016/j.rbmo.2022.03.015.

46. Vincent Conitzer, Christian Kroer, Debmalaya Panigrahi, Okke Schrijvers, Eric Sodomka, Nicolas E. Stier-Moses, and Chris Wilkens. Pacing Equilibrium in First-Price Auction Markets. To appear in *Management Science*. Supersedes conference paper 139 below.
45. Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E. Stier-Moses. Multiplicative Pacing Equilibria in Auction Markets. *Operations Research*, 70(2): 963-989 (2022). DOI: 10.1287/opre.2021.2167. Supersedes conference paper 135 below.
44. Michael Anis Mihdi Afnan, Yanhe Liu, Vincent Conitzer, Cynthia Rudin, Abhishek Mishra, Julian Savulescu, and Masoud Afnan. Interpretable, not black-box, artificial intelligence should be used for embryo selection. *Human Reproduction Open*, Volume 2021, Issue 4, 2021, hoab040.
43. Caspar Oesterheld and Vincent Conitzer. Extracting Money from Causal Decision Theorists. *Philosophical Quarterly*, Volume 71, Issue 4, October 2021, DOI: 10.1093/pq/pqaa086.
42. Vincent Conitzer. The Personalized A-Theory of Time and Perspective. *Dialectica*, Volume 74, Number 1, pages 1-29, 2020. Official version: pdf, html. Also available as arXiv:1802.2008.13207.
41. Michael Albert, Vincent Conitzer, Giuseppe Lopomo, and Peter Stone. Mechanism Design for Correlated Valuations: Efficient Methods for Revenue Maximization. *Operations Research*, 70(1): 562-584 (2022). DOI: 10.1287/opre.2020.2092. Builds on conference papers 125 and 127 below.
40. Andrew Kephart and Vincent Conitzer. The Revelation Principle for Mechanism Design with Signaling Costs. *ACM Transactions on Economics and Computation (TEAC)*, Volume 9, Issue 1, March 2021, Article Number 6, pages 1-35, DOI: 10.1145/3434408. Supersedes conference paper 123 below.
39. Aaron Kolb and Vincent Conitzer. Crying about a strategic wolf: A theory of crime and warning. *Journal of Economic Theory*, Volume 189, September 2020, 105094.
38. Rachel Freedman, Jana Schaich Borg, Walter Sinnott-Armstrong, John Dickerson, and Vincent Conitzer. Adapting a Kidney Exchange Algorithm to Align with Human Values. *Artificial Intelligence*, Volume 283, June 2020, 103261, DOI:10.1016/j.artint.2020.103261. Supersedes conference paper 131 below.
37. Joshua August Skorburg, Walter Sinnott-Armstrong, and Vincent Conitzer. AI Methods in Bioethics. *American Journal of Bioethics: Empirical Bioethics*, Volume 11, Issue 1, pages 37-39, 2020. DOI:10.1080/23294515.2019.1706206.
36. Vincent Conitzer. The Exact Computational Complexity of Evolutionarily Stable Strategies. *Mathematics of Operations Research*, 44(3): 783-792, 2019. DOI:10.1287/moor.2018.0945. Supersedes conference paper 100 below.
35. Vincent Conitzer. A Puzzle about Further Facts. *Erkenntnis*, June 2019, Volume 84, Issue 3, pp. 727-739 (DOI: 10.1007/s10670-018-9979-6). Also available as arXiv:1802.01161.
34. Suguru Ueda, Atsushi Iwasaki, Vincent Conitzer, Naoki Ohta, Yuko Sakurai, and Makoto Yokoo. Coalition structure generation in cooperative games with compact representations. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, 32(4): 503-533, 2018. Supersedes conference paper 61 below.
33. Vincent Conitzer. Technical Perspective: Designing Algorithms and the Fairness Criteria They Should Satisfy. *Communications of the ACM*, Volume 61, Issue 2, February 2018, pp. 92.
32. Yuqian Li and Vincent Conitzer. Game-Theoretic Question Selection for Tests. *Journal of Artificial Intelligence Research (JAIR)*, Volume 59, pp. 437-462, 2017. Supersedes conference paper 99 below.

31. Vincent Conitzer and Preston McAfee. Farewell Editorial: Looking Back on Our Terms Editing ACM TEAC and into the Future. *ACM Transactions on Economics and Computation (TEAC)*, Article 9e, Volume 5, Issue 2, March 2017.
30. Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified Representation in Approval-Based Committee Voting. In *Social Choice and Welfare*, Volume 48, Issue 2, pp. 461-485, February 2017. Supersedes conference paper 107 below.
29. Vincent Conitzer. On Stackelberg Mixed Strategies. In *Synthese* (special issue on Logic and the Foundations of Decision and Game Theory), Volume 193, Issue 3, pp. 689-703, March 2016.
28. Vincent Conitzer. Can rational choice guide us to correct *de se* beliefs? *Synthese*, Volume 192, Issue 12, pp. 4107-4119, December 2015.
27. Vincent Conitzer. A Dutch Book against Sleeping Beauties Who Are Evidential Decision Theorists. *Synthese*, Volume 192, Issue 9, pp. 2887-2899, October 2015.
26. Vincent Conitzer. A Devastating Example for the Halfer Rule. *Philosophical Studies*, Volume 172, Issue 8, pp. 1985-1992, August 2015.
25. Joshua Letchford, Dmytro Korzhyk, and Vincent Conitzer. On the Value of Commitment. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, Volume 28, Issue 6, pp. 986-1016, November 2014.
24. Mingyu Guo and Vincent Conitzer. Better Redistribution with Inefficient Allocation in Multi-Unit Auctions. *Artificial Intelligence (AIJ)*, Volume 216, pp. 287-308, November 2014. Supersedes conference paper 49 below.
23. Mathijs M. de Weerd, B. Paul Harrenstein, and Vincent Conitzer. Strategy-Proof Contract Auctions and the Role of Ties. *Games and Economic Behavior*, Special Issue on EC'08/'09, Volume 86, July 2014, pp. 405-420. Based on conference paper 59 below (very much rewritten).
22. Liad Wagman and Vincent Conitzer. False-Name-Proof Voting with Costs over Two Alternatives. *International Journal of Game Theory (IJGT)*, Volume 43, Issue 3, pp. 599-618, August 2014. Supersedes most of conference paper 46 below.
21. Mingyu Guo, Evangelos Markakis, Krzysztof R. Apt, and Vincent Conitzer. Undominated Groves Mechanisms. *Journal of Artificial Intelligence Research (JAIR)*, Volume 46, 2013, pp. 129-163. Supersedes conference paper 51 and most of conference paper 42 below.
20. Vincent Conitzer and Preston McAfee. The ACM Transactions on Economics and Computation: An Introduction. *ACM Transactions on Economics and Computation (TEAC)*, Article 1, Volume 1, Issue 1, January 2013.
19. Vincent Conitzer, Curtis Taylor, and Liad Wagman. Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases. *Marketing Science*, Volume 31, Number 2, 2012, pp. 277-292.
18. Vincent Conitzer. Should Social Network Structure Be Taken into Account in Elections? Short communication in *Mathematical Social Sciences (MSS)*, Special Issue on Computational Foundations of Social Choice, Volume 64, Issue 1, 2012, pp. 100-102.
17. Liad Wagman and Vincent Conitzer. Choosing Fair Lotteries to Defeat the Competition. *International Journal of Game Theory (IJGT)*, Volume 41, Issue 1, 2012, pp. 91-129. Supersedes conference paper 41 below.
16. Vincent Conitzer and Tuomas Sandholm. Computing Optimal Outcomes under an Expressive Representation of Settings with Externalities. *Journal of Computer and System Sciences (JCSS)*, Special Issue devoted to Knowledge Representation and Reasoning, Volume 78, Issue 1, January 2012, Pages 2-14. Supersedes conference paper 25 below.

15. Dmytro Korzhyk, Zhengyu Yin, Christopher Kiekintveld, Vincent Conitzer, and Milind Tambe. Stackelberg vs. Nash in Security Games: An Extended Investigation of Interchangeability, Equivalence, and Uniqueness. *Journal of Artificial Intelligence Research (JAIR)*, Volume 41, 2011, pp. 297-327. Supersedes conference paper 67 below.
14. Lirong Xia and Vincent Conitzer. Determining Possible and Necessary Winners under Common Voting Rules Given Partial Orders. *Journal of Artificial Intelligence Research (JAIR)*, Volume 41, 2011, pp. 25-67. Supersedes conference paper 45 below.
13. Vincent Conitzer and Makoto Yokoo. Using Mechanism Design to Prevent False-Name Manipulations. *AI Magazine*, Special Issue on Algorithmic Game Theory, Volume 31, Issue 4, December 2010, pp. 65-77.
12. Vincent Conitzer and Tuomas Sandholm. Expressive Markets for Donating to Charities. *Artificial Intelligence (AIJ)*, Special Issue on Representing, Processing, and Learning Preferences: Theoretical and Practical Challenges, Volume 175, Issues 7-8, May 2011, pp. 1251-1271. Supersedes conference paper 13 below.
11. Vincent Conitzer. Comparing Multiagent Systems Research in Combinatorial Auctions and Voting. *Annals of Mathematics and Artificial Intelligence (AMAI)*, Volume 58, Issue 3, 2010, pp. 239-259. Supersedes nonarchival conference paper 1 below.
10. Mingyu Guo and Vincent Conitzer. Optimal-in-Expectation Redistribution Mechanisms. *Artificial Intelligence (AIJ)*, Volume 174, Issues 5-6, April 2010, pp. 363-381. Supersedes conference paper 43 below.
9. Joseph Farfel and Vincent Conitzer. Aggregating Value Ranges: Preference Elicitation and Truthfulness. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, Special Issue on Computational Social Choice, Volume 22, Number 1, January 2011, pp. 127-150.
8. Vincent Conitzer. Making Decisions Based on the Preferences of Multiple Agents. *Communications of the ACM (CACM)*, Volume 53, Number 3, March 2010, pp. 84-94.
7. Vincent Conitzer. Eliciting Single-Peaked Preferences Using Comparison Queries. *Journal of Artificial Intelligence Research (JAIR)*, Volume 35, 2009, pp. 161-191. Supersedes conference paper 37 below.
6. Mingyu Guo and Vincent Conitzer. Worst-Case Optimal Redistribution of VCG Payments in Multi-Unit Auctions. *Games and Economic Behavior*, Special Section Dedicated to the 8th ACM Conference on Electronic Commerce, Volume 67, Issue 1, 2009, pp. 69-98. Supersedes conference paper 38 below.
5. Vincent Conitzer and Tuomas Sandholm. New Complexity Results about Nash Equilibria. *Games and Economic Behavior*, Special Issue on the Second World Congress of the Game Theory Society, Volume 63, Issue 2, 2008, pp. 621-641. Supersedes conference paper 5 below.
4. Mehmet Serkan Apaydin, Vincent Conitzer, and Bruce Randall Donald. Structure-based protein NMR assignments using native structural ensembles. *Journal of Biomolecular NMR*, 2008; 40(4):263-276. PMID: 18365752.
3. Vincent Conitzer, Tuomas Sandholm, and Jérôme Lang. When Are Elections with Few Candidates Hard to Manipulate? *Journal of the ACM (JACM)*, Volume 54, Issue 3, June 2007, Article 14 (33 pages). Supersedes conference papers 2 and 8 below.
2. Vincent Conitzer and Tuomas Sandholm. AWESOME: A General Multiagent Learning Algorithm that Converges in Self-Play and Learns a Best Response Against Stationary Opponents. *Machine Learning*, Special Issue on Learning and Computational Game Theory, Volume 67, Numbers 1-2, May 2007, pp. 23-43. Supersedes conference paper 10 below.

1. Vincent Conitzer and Tuomas Sandholm. Complexity of Constructing Solutions in the Core Based on Synergies Among Coalitions. *Artificial Intelligence (AIJ)*, volume 170, issues 6-7, May 2006, pp. 607-619. Supersedes conference paper 6 below.

Long archival conference papers

Note: the Workshop on Internet and Network Economics (WINE), which was later renamed the Conference on Web and Internet Economics (WINE), is listed as a conference. Other workshops are listed in a separate category.

195. Vijay Keswani, Cyrus Cousins, Breanna Nguyen, Vincent Conitzer, Hoda Heidari, Jana Schaich Borg, and Walter Sinnott-Armstrong. Moral Change or Noise? On Problems of Aligning AI With Temporally Unstable Human Feedback. In *Proceedings of the Fortieth AAAI Conference on Artificial Intelligence (AAAI-26) Alignment Track*, Singapore, 2026. (Acceptance rate 23% for alignment track.)
194. Ratip Emin Berker, Emanuel Tewolde, Vincent Conitzer, Mingyu Guo, Marijn Heule, and Lirong Xia. On the Edge of Core (Non-)emptiness: An Automated Reasoning Approach to Approval-based Multi-winner Voting. In *Proceedings of the Fortieth AAAI Conference on Artificial Intelligence (AAAI-26)*, Singapore, 2026. (Acceptance rate 18%.)
193. Vojtech Kovarik, Eric Olav Chen, Sami Petersen, Alexis Ghersengorin, and Vincent Conitzer. AI Testing Should Account for Sophisticated Strategic Behaviour. In *Proceedings of the Thirty-Ninth Conference on Neural Information Processing Systems (NeurIPS-25) Position Paper Track*, San Diego, CA, USA and Mexico City, Mexico, 2025.
192. Ratip Emin Berker, Silvia Casacuberta, Isaac Robinson, Christopher Ong, Vincent Conitzer, and Edith Elkind. From Independence of Clones to Composition Consistency: A Hierarchy of Barriers to Strategic Nomination. In *Proceedings of the 26th ACM Conference on Economics and Computation (EC-25)*, Stanford, CA, USA, 2025. (Acceptance rate 22%.)
191. Brian Hu Zhang, Ioannis Anagnostides, Emanuel Tewolde, Ratip Emin Berker, Gabriele Farina, Vincent Conitzer, and Tuomas Sandholm. Learning and Computation of Φ -Equilibria at the Frontier of Tractability. In *Proceedings of the 26th ACM Conference on Economics and Computation (EC-25)*, Stanford, CA, USA, 2025. (Acceptance rate 22%.)
190. Caspar Oesterheld and Vincent Conitzer. Choosing what game to play without selecting equilibria: results on inferring safe (Pareto) improvements in binary constraint structures. In *Proceedings of the 20th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-25)*, Düsseldorf, Germany, 2025.
189. Scott Emmons, Caspar Oesterheld, Vincent Conitzer, and Stuart Russell. Observation Interference in Partially Observable Assistance Games. In *Proceedings of the Forty-Second International Conference on Machine Learning (ICML-25)*, Vancouver, Canada, 2025. (Acceptance rate 27%.)
188. Brian Hu Zhang, Ioannis Anagnostides, Emanuel Tewolde, Ratip Emin Berker, Gabriele Farina, Vincent Conitzer, and Tuomas Sandholm. Expected Variational Inequalities. In *Proceedings of the Forty-Second International Conference on Machine Learning (ICML-25)*, Vancouver, Canada, 2025. (Acceptance rate 1% for oral presentations.)
187. Vijay Keswani, Vincent Conitzer, Walter Sinnott-Armstrong, Breanna K. Nguyen, Hoda Heidari, and Jana Schaich Borg. Can AI Model the Complexities of Human Moral Decision-Making? A Qualitative Study of Kidney Allocation Decisions. In *Proceedings of the ACM conference on Human Factors in Computing Systems (CHI-25)*, Yokohama, Japan, 2025.
186. Vojtech Kovarik, Nathaniel Sauerberg, Lewis Hammond, and Vincent Conitzer. Game Theory with Simulation in the Presence of Unpredictable Randomisation. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-25)*, Detroit, MI, USA, 2025. (Acceptance rate 24%.)

185. Emery Cooper, Caspar Oesterheld, and Vincent Conitzer. Characterising simulation-based program equilibria. In *Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI-25)*, Philadelphia, PA, USA, 2025. (Acceptance rate 23%.)
184. Emanuel Tewolde, Brian Hu Zhang, Caspar Oesterheld, Tuomas Sandholm, and Vincent Conitzer. Computing Game Symmetries and Equilibria That Respect Them. In *Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI-25)*, Philadelphia, PA, USA, 2025. (Acceptance rate 23%.)
183. Ratip Emin Berker, Emanuel Tewolde, Ioannis Anagnostides, Tuomas Sandholm, and Vincent Conitzer. The Value of Recall in Extensive-Form Games. In *Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI-25)*, Philadelphia, PA, USA, 2025. (Acceptance rate 23%.)
182. Yixuan Even Xu, Hanrui Zhang, Yu Cheng, and Vincent Conitzer. Aggregating Quantitative Relative Judgments: From Social Choice to Ranking Prediction. In *Proceedings of the Thirty-Eighth Conference on Neural Information Processing Systems (NeurIPS-24)*, Vancouver, Canada, 2024. (Acceptance rate 26%.)
181. Rakshit Trivedi, Akbir Khan, Jesse Clifton, Lewis Hammond, Edgar Duenez-Guzman, Dipam Chakraborty, John Agapiou, Jayd Matyas, Sasha Vezhnevets, Barna Pásztor, Yunke Ao, Omar G. Younis, Jiawei Huang, Benjamin Swain, Haoyuan Qin, Deng, Ziwei Deng, Utku Erdoğanaras, Yue Zhao, Marko Tescic, Natasha Jaques, Jakob Foerster, Vincent Conitzer, José Hernández-Orallo, Dylan Hadfield-Menell, Joel Leibo. Melting Pot Contest: Charting the Future of Generalized Cooperative Intelligence. In *Proceedings of the Thirty-Eighth Conference on Neural Information Processing Systems (NeurIPS-24)*, Datasets and Benchmarks Track, Vancouver, Canada, 2024. (Acceptance rate 25% for the track.)
180. Vijay Keswani, Vincent Conitzer, Hoda Heidari, Jana Schaich Borg, and Walter Sinnott-Armstrong. On the Pros and Cons of Active Learning for Moral Preference Elicitation. In *Proceedings of the Seventh AAAI/ACM Conference on AI, Ethics, and Society (AIES-24)*, San Jose, CA, USA, 2024. (Acceptance rate 32%.)
179. Kyle Boerstler, Vijay Keswani, Lok Chan, Jana Schaich Borg, Vincent Conitzer, Hoda Heidari, and Walter Sinnott-Armstrong. On The Stability of Moral Preferences: A Problem with Computational Elicitation Methods. In *Proceedings of the Seventh AAAI/ACM Conference on AI, Ethics, and Society (AIES-24)*, San Jose, CA, USA, 2024. (Acceptance rate 32%.)
178. Brian Hu Zhang, Gabriele Farina, Ioannis Anagnostides, Federico Cacciamani, Stephen Marcus McAleer, Andreas Alexander Haupt, Andrea Celli, Nicola Gatti, Vincent Conitzer, and Tuomas Sandholm. Steering No-Regret Learners to a Desired Equilibrium. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC-24)*, New Haven, CT, USA, 2024.
177. Vincent Conitzer, Rachel Freedman, Jobst Heitzig, Wesley H. Holliday, Bob M. Jacobs, Nathan Lambert, Milan Mossé, Eric Pacuit, Stuart Russell, Hailey Schoelkopf, Emanuel Tewolde, and William S. Zwicker. Social Choice Should Guide AI Alignment in Dealing with Diverse Human Feedback. In *Proceedings of the 41st International Conference on Machine Learning (ICML-24)*, Vienna, Austria, 2024. (Acceptance rate for position papers 26%.)
176. Ratip Emin Berker and Vincent Conitzer. Computing Optimal Equilibria in Repeated Games with Restarts. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24)*, Jeju, South Korea, 2024.
175. Emanuel Tewolde, Brian Zhang, Caspar Oesterheld, Manolis Zampetakis, Tuomas Sandholm, Paul Goldberg, and Vincent Conitzer. Imperfect-Recall Games: Equilibrium Concepts and Their Complexity. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24)*, Jeju, South Korea, 2024.

174. Emanuel Tewolde and Vincent Conitzer. Game Transformations That Preserve Nash Equilibria or Best Response Sets. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24)*, Jeju, South Korea, 2024.
173. Vincent Conitzer. The Complexity of Computing Robust Mediated Equilibria in Ordinal Games. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI-24)*, Vancouver, Canada, 2024. (Acceptance rate 24%.)
172. Yixuan Even Xu, Hanrui Zhang, and Vincent Conitzer. Non-Excludable Bilateral Trade Between Groups. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI-24)*, Vancouver, Canada, 2024. (Acceptance rate 24%.)
171. Caspar Oesterheld, Johannes Treutlein, Roger Grosse, Vincent Conitzer, and Jakob Foerster. Similarity-based cooperative equilibrium. In *Proceedings of the Thirty-Seventh Conference on Neural Information Processing Systems (NeurIPS-23)*, New Orleans, LA, USA, 2023.
170. Brian Hu Zhang, Gabriele Farina, Ioannis Anagnostides, Federico Cacciamani, Stephen Marcus McAleer, Andreas Alexander Haupt, Andrea Celli, Nicola Gatti, Vincent Conitzer, and Tuomas Sandholm. Computing Optimal Equilibria and Mechanisms via Learning in Zero-Sum Extensive-Form Games. In *Proceedings of the Thirty-Seventh Conference on Neural Information Processing Systems (NeurIPS-23)*, New Orleans, LA, USA, 2023.
169. Caspar Oesterheld, Abram Demski, and Vincent Conitzer. A theory of bounded inductive rationality. In *Proceedings of the 19th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-23)*, Oxford, UK, 2023. (Acceptance rate 50%.)
168. Hanrui Zhang, Yu Cheng, and Vincent Conitzer. Efficiently Solving Turn-Taking Stochastic Games with Extensive-Form Correlation. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC-23)*, London, UK, 2023.
167. Vojtech Kovarik, Caspar Oesterheld, and Vincent Conitzer. Game Theory with Simulation of Other Players. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI-23)*, Macao, 2023. (Acceptance rate 15%.)
166. Emanuel Tewolde, Caspar Oesterheld, Vincent Conitzer, and Paul Goldberg. The Computational Complexity of Single-Player Imperfect-Recall Games. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI-23)*, Macao, 2023. (Acceptance rate 15%.)
165. Steven Jecmen, Minji Yoon, Vincent Conitzer, Nihar Shah, and Fei Fang. A Dataset on Malicious Paper Bidding in Peer Review. *The ACM Web Conference proceedings (TheWebConf-23)*, Austin, TX, USA, 2023. (Acceptance rate 19%.)
164. Vincent Conitzer and Caspar Oesterheld. Foundations of Cooperative AI. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI-23) Senior Member / Blue Sky Track*, Washington, DC, USA, 2023.
163. Steven Jecmen, Hanrui Zhang, Ryan Liu, Fei Fang, Vincent Conitzer, and Nihar Shah. Near-Optimal Reviewer Splitting in Two-Phase Paper Reviewing and Conference Experiment Design. In *the Tenth AAAI Conference on Human Computation and Crowdsourcing (HCOMP-22)*, pages 102-113, 2022. (Acceptance rate 31%. Honorable Mention for Best Paper Award.)
162. Hanrui Zhang, Yu Cheng, and Vincent Conitzer. Efficient Algorithms for Planning with Participation Constraints. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC-22)*, pages 1121-1140, Boulder, CO, USA, 2022. (Acceptance rate 27%.)

161. Scott Emmons, Caspar Oesterheld, Andrew Critch, Vincent Conitzer, and Stuart Russell. For Learning in Symmetric Teams, Local Optima are Global Nash Equilibria. In *Proceedings of the 39th International Conference on Machine Learning (ICML-22)*, pages 5924-5943, Baltimore, MD, USA, 2022. (Acceptance rate 22%.)
160. Hanrui Zhang, Yu Cheng, and Vincent Conitzer. Planning with Participation Constraints. In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI-22)*, pages 5260-5267, 2022. (Acceptance rate 15%.)
159. Vincent Conitzer, Debmalaya Panigrahi, and Hanrui Zhang. Learning Influence Adoption in Heterogeneous Networks. In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI-22)*, pages 6411-6419, 2022. (Acceptance rate 15%.)
158. Vincent Conitzer, Zhe Feng, David Parkes, and Eric Sodomka. Welfare-Preserving ϵ -BIC to BIC Transformation with Negligible Revenue Loss. In *the 17th Conference on Web and Internet Economics (WINE-21)*, pp. 76-94, 2021. (Acceptance rate 27%.)
157. Hanrui Zhang and Vincent Conitzer. Automated Dynamic Mechanism Design. In *Proceedings of the Thirty-Fifth Conference on Neural Information Processing Systems (NeurIPS-21)*, pages 27785-27797, 2021. (Acceptance rate 26%.)
156. Michael Anis Mihdi Afnan, Cynthia Rudin, Vincent Conitzer, Julian Savulescu, Abhishek Mishra, Yanhe Liu, and Masoud Afnan. Ethical Implementation of Artificial Intelligence to Select Embryos in In Vitro Fertilization. In *Proceedings of the Fourth AAAI/ACM Conference on AI, Ethics, and Society (AIES-21)*, pp. 316-326, 2021. (Acceptance rate 37%.)
155. Caspar Oesterheld and Vincent Conitzer. Safe Pareto improvements for delegated game playing. In *Proceedings of the Twentieth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-21)*, pp. 983-991, 2021. (Acceptance rate 25%.)
154. Duncan McElfresh, Lok Chan, Kenzie Doyle, Walter Sinnott-Armstrong, Vincent Conitzer, Jana Schaich Borg, and John Dickerson. Indecision Modeling. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)*, pp. 5975-5983, 2021. (Acceptance rate 21%.)
153. Hanrui Zhang, Yu Cheng, and Vincent Conitzer. Automated Mechanism Design for Classification with Partial Verification. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)*, pp. 5789-5796, 2021. (Acceptance rate 21%.)
152. Hanrui Zhang and Vincent Conitzer. Incentive-Aware PAC Learning. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)*, pp. 5797-5804, 2021. (Acceptance rate 21%.)
151. Anilesh Krishnaswamy, Haoming Li, David Rein, Hanrui Zhang, and Vincent Conitzer. Classification with Strategically Withheld Data. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)*, pp. 5514-5522, 2021. (Acceptance rate 21%.)
150. Hanrui Zhang, Yu Cheng, and Vincent Conitzer. Classification with Few Tests through Self-Selection. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)*, pp. 5805-5812, 2021. (Acceptance rate 21%.)
149. Steven Jecmen, Hanrui Zhang, Ryan Liu, Nihar B. Shah, Vincent Conitzer, and Fei Fang. Mitigating Manipulation in Peer Review via Randomized Reviewer Assignments. In *Proceedings of the Thirty-Fourth Conference on Neural Information Processing Systems (NeurIPS-20)*, 2020. (Acceptance rate 20%.)
148. Caspar Oesterheld and Vincent Conitzer. Decision Scoring Rules. In *the 16th Conference on Web and Internet Economics (WINE-20)*, Beijing, China / virtual, 2020. *Accepted as full paper, but opted to put only abstract in proceedings.* (Acceptance rate 31%.)

147. Caspar Oesterheld and Vincent Conitzer. Minimum-regret contracts for principal-expert problems. In *the 16th Conference on Web and Internet Economics (WINE-20)*, pp. 430-443, Beijing, China / virtual, 2020. (Acceptance rate 31%.)
146. Vincent Conitzer, Yuan Deng, and Shaddin Dughmi. Bayesian Repeated Zero-Sum Games with Persistent State, with Application to Security Games. In *the 16th Conference on Web and Internet Economics (WINE-20)*, pp. 444-458, Beijing, China / virtual, 2020. (Acceptance rate 31%.)
145. Hanrui Zhang and Vincent Conitzer. Learning the Valuations of a k -demand Agent. In *Proceedings of the 37th International Conference on Machine Learning (ICML-20)*, pp. 11066-11075, 2020. (Acceptance rate 22%.)
144. Vincent Conitzer, Debmalaya Panigrahi, and Hanrui Zhang. Learning Opinions in Social Networks. In *Proceedings of the 37th International Conference on Machine Learning (ICML-20)*, pp. 2122-2132, 2020. (Acceptance rate 22%.)
143. Hanrui Zhang and Vincent Conitzer. Combinatorial Ski Rental and Online Bipartite Matching. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC-20)*, pp. 879-910, 2020. (Acceptance rate 20%.)
142. Lok Chan, Kenzie Doyle, Duncan McElfresh, Vincent Conitzer, John Dickerson, Jana Schaich Borg, and Walter Sinnott-Armstrong. Artificial Artificial Intelligence: Measuring Influence of AI “Assessments” on Moral Decision-Making. In *Proceedings of the Third AAAI/ACM Conference on AI, Ethics, and Society (AIES-20)*, pp. 214-220, New York, NY, USA, 2020. (Acceptance rate 34%.)
141. Hanrui Zhang, Yu Cheng, and Vincent Conitzer. Distinguishing Distributions When Samples Are Strategically Transformed. In *Proceedings of the Thirty-Third Conference on Neural Information Processing Systems (NeurIPS-19)*, Vancouver, Canada, 2019. (Acceptance rate 21%.)
140. Hanrui Zhang, Yu Cheng, and Vincent Conitzer. When Samples Are Strategically Selected. In *Proceedings of the 36th International Conference on Machine Learning (ICML-19)*, pp. 7345-7353, Long Beach, CA, USA, 2019. (Acceptance rate 23%.)
139. Vincent Conitzer, Christian Kroer, Debmalaya Panigrahi, Okke Schrijvers, Eric Sodomka, Nicolas E. Stier-Moses, and Chris Wilkens. Pacing Equilibrium in First-Price Auction Markets. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC-19)*, Phoenix, AZ, USA, 2019. *Accepted as full paper, but opted to put only abstract in proceedings; full version available as arXiv:1811.07166.* (Acceptance rate 28%.)
138. Hanrui Zhang and Vincent Conitzer. A PAC Framework for Aggregating Agents’ Judgments. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*, pp. 2237-2244, Honolulu, HI, USA, 2019. (Acceptance rate 16%.)
137. Hanrui Zhang, Yu Cheng, and Vincent Conitzer. A Better Algorithm for Societal Tradeoffs. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*, pp. 2229-2236, Honolulu, HI, USA, 2019. (Acceptance rate 16%.)
136. Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Wortman Vaughan. Group Fairness for the Allocation of Indivisible Goods. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*, pp. 1853-1860, Honolulu, HI, USA, 2019. (Acceptance rate 16%.)
135. Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E. Stier-Moses. Multiplicative Pacing Equilibria in Auction Markets. *Fourteenth Conference on Web and Internet Economics (WINE-18)*, Oxford, United Kingdom, 2018. *Accepted as full paper, but opted to put only abstract in proceedings; full version available as arXiv:1706.07151.*

134. Mathijs de Weerd, Michael Albert, Vincent Conitzer, and Koos van der Linden. Complexity of Scheduling Charging in the Smart Grid. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence and the 23rd European Conference on Artificial Intelligence (IJCAI-ECAI-18)*, pp. 4736-4742, Stockholm, Sweden, 2018. (Acceptance rate 20%.)
133. Max Kramer, Jana Schaich Borg, Vincent Conitzer, and Walter Sinnott-Armstrong. When Do People Want AI to Make Decisions? In *Proceedings of the First AAAI/ACM Conference on AI, Ethics, and Society (AIES-18)*, pp. 204-209, New Orleans, LA, USA, 2018. (Acceptance rate 19% for oral presentations.)
132. Rupert Freeman*, Seyed Majid Zahedi*, Vincent Conitzer, and Benjamin Lee (* co-first authors). Dynamic Proportional Sharing: A Game-Theoretic Approach. In *Proceedings of the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, POMACS 2(1): 3:1-3:36, Irvine, CA, USA, 2018. (Acceptance rate between 14 and 22% for the deadline we submitted to.)
131. Rachel Freedman, Jana Schaich Borg, Walter Sinnott-Armstrong, John Dickerson, and Vincent Conitzer. Adapting a Kidney Exchange Algorithm to Align with Human Values. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, pp. 1636-1643, New Orleans, LA, USA, 2018. (Acceptance rate <25%. The Outstanding Student Paper Honorable Mention.)
130. Yuan Deng and Vincent Conitzer. Disarmament Games with Resources. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, pp. 981-988, New Orleans, LA, USA, 2018. (Acceptance rate <25%.)
129. Rupert Freeman, Seyed Majid Zahedi, and Vincent Conitzer. Fair Social Choice in Dynamic Settings. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*, pp. 4580-4587, Melbourne, Australia, 2017. (Acceptance rate 26%.)
128. Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair Public Decision Making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC-17)*, pp. 629-646, Cambridge, MA, USA, 2017.
127. Michael Albert, Vincent Conitzer, and Peter Stone. Mechanism Design with Unknown Correlated Distributions: Can We Learn Optimal Mechanisms? In *Proceedings of the Sixteenth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-17)*, pp. 69-77, Sao Paulo, Brazil, 2017. (Acceptance rate 26%.)
126. Yuan Deng and Vincent Conitzer. Disarmament Games. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)*, pp. 473-479, San Francisco, CA, USA, 2017. (Acceptance rate 25%.)
125. Michael Albert, Vincent Conitzer, and Peter Stone. Automated Design of Robust Mechanisms. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)*, pp. 298-304, San Francisco, CA, USA, 2017. (Acceptance rate 25%.)
124. Vincent Conitzer. Computing Equilibria with Partial Commitment. In *Proceedings of the Twelfth Conference on Web and Internet Economics (WINE-16)*, pp. 1-14, Montreal, Canada, 2016. (Acceptance rate 40%.)
123. Andrew Kephart and Vincent Conitzer. The Revelation Principle for Mechanism Design with Reporting Costs. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC-16)*, pp. 85-102, Maastricht, the Netherlands, 2016. (Acceptance rate 33%.)
122. Catherine Moon and Vincent Conitzer. Role Assignment for Game-Theoretic Cooperation. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, pp. 416-423, New York City, NY, USA, 2016. (Acceptance rate < 25%.)

121. Yuqian Li, Vincent Conitzer, and Dmytro Korzhyk. Catcher-Evader Games. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, pp. 329-337, New York City, NY, USA, 2016. (Acceptance rate < 25%.)
120. Garrett Andersen and Vincent Conitzer. ATUCAPTS: Automated Tests That a User Cannot Pass Twice Simultaneously. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, pp. 3662-3669, New York City, NY, USA, 2016. (Acceptance rate < 25%.)
119. Haifeng Xu, Rupert Freeman, Vincent Conitzer, Shaddin Dughmi, and Milind Tambe. Signaling in Bayesian Stackelberg Games. In *Proceedings of the Fifteenth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-16)*, pp. 150-158, Singapore, 2016. (Acceptance rate 25%.)
118. Markus Brill, Vincent Conitzer, Rupert Freeman, and Nisarg Shah. False-Name-Proof Recommendations in Social Networks. In *Proceedings of the Fifteenth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-16)*, pp. 332-340, Singapore, 2016. (Acceptance rate 25%.)
117. Vincent Conitzer, Rupert Freeman, Markus Brill, and Yuqian Li. Rules for Choosing Societal Tradeoffs. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)*, pp. 460-467, Phoenix, AZ, USA, 2016. (Acceptance rate 26%.)
116. Markus Brill, Rupert Freeman, and Vincent Conitzer. Computing Possible and Necessary Equilibrium Actions (and Bipartisan Set Winners). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)*, pp. 369-375, Phoenix, AZ, USA, 2016. (Acceptance rate 26%.)
115. Michael Albert, Vincent Conitzer, and Giuseppe Lopomo. Maximizing Revenue with Limited Correlation: The Cost of Ex-Post Incentive Compatibility. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)*, pp. 383-389, Phoenix, AZ, USA, 2016. (Acceptance rate 26%.)
114. Sune Kristian Jakobsen, Troels Bjerre Sørensen, and Vincent Conitzer. Timeability of Extensive-Form Games. In *Proceedings of the Seventh Innovations in Theoretical Computer Science Conference (ITCS-16)*, pp. 191-199, Cambridge, MA, USA, 2016. (Acceptance rate 28%.)
113. Catherine Moon and Vincent Conitzer. Maximal Cooperation in Repeated Games on Social Networks. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-15)*, pp. 216-223, Buenos Aires, Argentina, 2015. (Acceptance rate 29%.)
112. Andrew Kephart and Vincent Conitzer. Complexity of Mechanism Design with Signaling Costs. In *Proceedings of the Fourteenth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-15)*, pp. 357-365, Istanbul, Turkey, 2015. (Acceptance rate 25%.)
111. Rupert Freeman, Markus Brill, and Vincent Conitzer. General Tiebreaking Schemes for Computational Social Choice. In *Proceedings of the Fourteenth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-15)*, pp. 1401-1409, Istanbul, Turkey, 2015. (Acceptance rate 25%.)
110. Markus Brill and Vincent Conitzer. Strategic Voting and Strategic Candidacy. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, pp. 819-826, Austin, TX, USA, 2015. (Acceptance rate 12% for oral presentations.)
109. Michael Albert, Vincent Conitzer, and Giuseppe Lopomo. Assessing the Robustness of Cremer-McLean with Automated Mechanism Design. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, pp. 763-769, Austin, TX, USA, 2015. (Acceptance rate 12% for oral presentations.)

108. Yuqian Li and Vincent Conitzer. Cooperative Game Solution Concepts that Maximize Stability under Noise. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, pp. 979-985, Austin, TX, USA, 2015. (Acceptance rate 12% for oral presentations.)
107. Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified Representation in Approval-Based Committee Voting. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, pp. 784-790, Austin, TX, USA, 2015. (Acceptance rate 12% for oral presentations.)
106. Rupert Freeman, Markus Brill, and Vincent Conitzer. On the Axiomatic Characterization of Runoff Voting Rules. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI-14)*, pp. 675-681, Quebec City, Canada, 2014. (Acceptance rate 28%.)
105. Troels Bjerre Sørensen, Melissa Dalis, Joshua Letchford, Dmytro Korzhuk, and Vincent Conitzer. Beat the Cheater: Computing Game-Theoretic Strategies for When to Kick a Gambler out of a Casino. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI-14)*, pp. 798-804, Quebec City, Canada, 2014. (Acceptance rate 28%.)
104. Vincent Conitzer and Angelina Vidali. Mechanism Design for Scheduling with Uncertain Execution Time. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI-14)*, pp. 623-629, Quebec City, Canada, 2014. (Acceptance rate 28%.)
103. Haifeng Xu, Fei Fang, Albert Jiang, Vincent Conitzer, Shaddin Dughmi, and Milind Tambe. Solving Zero-Sum Security Games in Discretized Spatio-Temporal Domains. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI-14)*, pp. 1500-1506, Quebec City, Canada, 2014. (Acceptance rate 28%.)
102. Yuqian Li and Vincent Conitzer. Complexity of Stability-based Solution Concepts in Multi-issue and MC-net Cooperative Games. In *Proceedings of the Thirteenth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-14)*, pp. 581-588, Paris, France, 2014. (Acceptance rate 24%.)
101. Vincent Conitzer. The Maximum Likelihood Approach to Voting on Social Networks. In *Proceedings of the 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton-13)*, pp. 1482-1487, Allerton Retreat Center, Monticello, IL, USA, 2013. (Invited.)
100. Vincent Conitzer. The Exact Computational Complexity of Evolutionarily Stable Strategies. In *Proceedings of the Ninth Conference on Web and Internet Economics (WINE-13)*, pp. 96-108, Cambridge, MA, USA, 2013. (Acceptance rate 24%.)
99. Yuqian Li and Vincent Conitzer. Game-Theoretic Question Selection for Tests. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI-13)*, pp. 254-262, Beijing, China, 2013. (Acceptance rate 28%.)
98. Garrett Andersen and Vincent Conitzer. Fast Equilibrium Computation for Infinitely Repeated Games. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI-13)*, pp. 53-59, Bellevue, WA, USA, 2013. (Acceptance rate 29%.)
97. Joshua Letchford and Vincent Conitzer. Solving Security Games on Graphs via Marginal Probabilities. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI-13)*, pp. 591-597, Bellevue, WA, USA, 2013. (Acceptance rate 29%.)
96. Yuqian Li and Vincent Conitzer. Optimal Internet Auctions with Costly Communication. In *Proceedings of the Twelfth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-13)*, pp. 683-690, St. Paul, MN, USA, 2013. (Acceptance rate 23%.)
95. Taiki Todo and Vincent Conitzer. False-name-proof Matching. In *Proceedings of the Twelfth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-13)*, pp. 311-318, St. Paul, MN, USA, 2013. (Acceptance rate 23%.)

94. Manish Jain, Vincent Conitzer, and Milind Tambe. Security Scheduling for Real-world Networks. In *Proceedings of the Twelfth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-13)*, pp. 215-222, St. Paul, MN, USA, 2012. (Acceptance rate 23%.)
93. Vincent Conitzer. Computing Game-Theoretic Solutions and Applications to Security. In *Proceedings of the 26th National Conference on Artificial Intelligence (AAAI-12)*, pp. 2106-2112, Toronto, ON, Canada, 2012. (Invited as a “What’s Hot” paper to the AAAI-12 Sub-Area Spotlights track.)
92. Bo Waggoner, Lirong Xia, and Vincent Conitzer. Evaluating Resistance to False-Name Manipulations in Elections. In *Proceedings of the 26th National Conference on Artificial Intelligence (AAAI-12)*, pp. 1485-1491, Toronto, ON, Canada, 2012. (Accepted for oral and poster presentation. Overall acceptance rate 26%.)
91. Joshua Letchford, Liam MacDermed, Vincent Conitzer, Ronald Parr, and Charles Isbell. Computing Optimal Strategies to Commit to in Stochastic Games. In *Proceedings of the 26th National Conference on Artificial Intelligence (AAAI-12)*, pp. 1380-1386, Toronto, ON, Canada, 2012. (Acceptance rate 26%.)
90. Vincent Conitzer. An Undergraduate Course in the Intersection of Computer Science and Economics. In *Proceedings of the Third AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-12)*, pp. 2357-2362, Toronto, ON, Canada, 2012.
89. Vincent Conitzer and Lirong Xia. Paradoxes of Multiple Elections: An Approximation Approach. In *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR-12)*, pp. 179-187, Rome, Italy, 2012. (Acceptance rate 26%.)
88. Sayan Bhattacharya, Vincent Conitzer, and Kamesh Munagala. Approximation Algorithm for Security Games with Costly Resources. In *Proceedings of the Seventh Workshop on Internet and Network Economics (WINE-11)*, pp. 13-24, Singapore, 2011. (Acceptance rate 30%.)
87. Mingyu Guo, Victor Naroditskiy, Vincent Conitzer, Amy Greenwald, and Nicholas R. Jennings. Budget-Balanced and Nearly Efficient Randomized Mechanisms: Public Goods and Beyond. In *Proceedings of the Seventh Workshop on Internet and Network Economics (WINE-11)*, pp. 158-169, Singapore, 2011. (Acceptance rate 30%.)
86. Michael Zuckerman, Piotr Faliszewski, Vincent Conitzer, and Jeffrey S. Rosenschein. An NTU Cooperative Game Theoretic View of Manipulating Elections. In *Proceedings of the Seventh Workshop on Internet and Network Economics (WINE-11)*, pp. 363-374, Singapore, 2011. (Acceptance rate 30%.)
85. Vincent Conitzer and Dmytro Korzhyk. Commitment to Correlated Strategies. In *Proceedings of the 25th National Conference on Artificial Intelligence (AAAI-11)*, pp. 632-637, San Francisco, CA, USA, 2011. (Acceptance rate 25%.)
84. Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating Manipulations in Voting with Partial Information. In *Proceedings of the 25th National Conference on Artificial Intelligence (AAAI-11)*, pp. 638-643, San Francisco, CA, USA, 2011. (Acceptance rate 25%.)
83. Dmytro Korzhyk, Vincent Conitzer, and Ronald Parr. Security Games with Multiple Attacker Resources. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, pp. 273-279, Barcelona, Catalonia, Spain, 2011. (Acceptance rate 17% for oral presentations.)
82. Lirong Xia and Vincent Conitzer. A Maximum Likelihood Approach towards Aggregating Partial Orders. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, pp. 446-451, Barcelona, Catalonia, Spain, 2011. (Acceptance rate 30%.)

81. Vincent Conitzer, Jérôme Lang, and Lirong Xia. Hypercubewise Preference Aggregation in Multi-issue Domains. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, pp. 158-163, Barcelona, Catalonia, Spain, 2011. (Acceptance rate 30%.)
80. Lirong Xia, Vincent Conitzer, and Jérôme Lang. Strategic Sequential Voting in Multi-Issue Domains and Multiple-Election Paradoxes. In *Proceedings of the Twelfth ACM Conference on Electronic Commerce (EC-11)*, pp. 179-188, San Jose, CA, USA, 2010. (Acceptance rate 26%.)
79. Dmytro Korzhyk, Vincent Conitzer, and Ronald Parr. Solving Stackelberg Games with Uncertain Observability. In *Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-11)*, pp. 1013-1020, Taipei, Taiwan, 2011. (Acceptance rate 22%.)
78. Manish Jain, Dmytro Korzhyk, Ondrej Vanek, Vincent Conitzer, Michal Pechoucek, and Milind Tambe. A Double Oracle Algorithm for Zero-Sum Security Games on Graphs. In *Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-11)*, pp. 327-334, Taipei, Taiwan, 2011. (Acceptance rate 22%.)
77. Vincent Conitzer, Nicole Immorlica, Joshua Letchford, Kamesh Munagala, and Liad Wagman. False-Name-Proofness in Social Networks. In *Proceedings of the Sixth Workshop on Internet and Network Economics (WINE-10)*, pp. 209-221, Stanford, CA, 2010.
76. Lirong Xia and Vincent Conitzer. Strategy-proof Voting Rules over Multi-issue Domains with Restricted Preferences. In *Proceedings of the Sixth Workshop on Internet and Network Economics (WINE-10)*, pp. 402-414, Stanford, CA, 2010.
75. Dmytro Korzhyk, Vincent Conitzer, and Ronald Parr. Complexity of Computing Optimal Stackelberg Strategies in Security Resource Allocation Games. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI-10)*, pp. 805-810, Atlanta, GA, USA, 2010. (Acceptance rate 27%.)
74. Lirong Xia and Vincent Conitzer. Stackelberg Voting Games: Computational Aspects and Paradoxes. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI-10)*, pp. 921-926, Atlanta, GA, USA, 2010. (Acceptance rate 27%.)
73. Lirong Xia and Vincent Conitzer. Compilation Complexity of Common Voting Rules. In *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI-10)*, pp. 915-920, Atlanta, GA, USA, 2010. (Acceptance rate 27%.)
72. Joshua Letchford and Vincent Conitzer. Computing Optimal Strategies to Commit to in Extensive-Form Games. In *Proceedings of the Eleventh ACM Conference on Electronic Commerce (EC-10)*, pp. 83-92, Cambridge, MA, USA, 2010. (Acceptance rate 33%.)
71. Lirong Xia, Vincent Conitzer, and Ariel D. Procaccia. A Scheduling Approach to Coalitional Manipulation. In *Proceedings of the Eleventh ACM Conference on Electronic Commerce (EC-10)*, pp. 275-284, Cambridge, MA, USA, 2010. (Acceptance rate 33%.)
70. Mingyu Guo and Vincent Conitzer. Strategy-proof Allocation of Multiple Items between Two Agents without Payments or Priors. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-10)*, pp. 881-888, Toronto, ON, Canada, 2010. (Acceptance rate 24%.)
69. Atsushi Iwasaki, Vincent Conitzer, Yoshifusa Omori, Yuko Sakurai, Taiki Todo, Mingyu Guo, and Makoto Yokoo. Worst-case efficiency ratio in false-name-proof combinatorial auction mechanisms. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-10)*, pp. 633-640, Toronto, ON, Canada, 2010. (Acceptance rate 24%.)

68. Lirong Xia, Vincent Conitzer, and Jérôme Lang. Aggregating Preferences in Multi-Issue Domains by Using Maximum Likelihood Estimators. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-10)*, pp. 399-406, Toronto, ON, Canada, 2010. (Acceptance rate 24%.)
67. Zhengyu Yin, Dmytro Korzhyk, Christopher Kiekintveld, Vincent Conitzer, and Milind Tambe. Stackelberg vs. Nash in Security Games: Interchangeability, Equivalence, and Uniqueness. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-10)*, pp. 1139-1146, Toronto, ON, Canada, 2010. (Acceptance rate 24%.)
66. Sayan Bhattacharya, Vincent Conitzer, Kamesh Munagala, and Lirong Xia. Incentive Compatible Budget Elicitation in Multi-unit Auctions. In the *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-10)*, pp. 554-572, Austin, TX, USA, 2010. (Acceptance rate 31%.)
65. Peng Shi, Vincent Conitzer, and Mingyu Guo. Prediction Mechanisms That Do Not Incentivize Undesirable Actions. In *Proceedings of the Fifth Workshop on Internet and Network Economics (WINE-09)*, pp. 89-100, Rome, Italy, 2009. (Acceptance rate 23% for full-length papers.)
64. Mingyu Guo, Vincent Conitzer, and Daniel Reeves. Competitive Repeated Allocation Without Payments. In *Proceedings of the Fifth Workshop on Internet and Network Economics (WINE-09)*, pp. 244-255, Rome, Italy, 2009. (Acceptance rate 23% for full-length papers.)
63. Joshua Letchford, Vincent Conitzer, and Kamesh Munagala. Learning and Approximating the Optimal Strategy to Commit To. In the *Second International Symposium on Algorithmic Game Theory (SAGT-09)*, pp. 250-262, Paphos, Cyprus, 2009.
62. Vincent Conitzer. Approximation Guarantees for Fictitious Play. In the *Proceedings of the 47th Annual Allerton Conference on Communication, Control, and Computing (Allerton-09)*, pp. 636-643, Allerton Retreat Center, Monticello, IL, USA, 2009. (Invited.)
61. Naoki Ohta, Vincent Conitzer, Ryo Ichimura, Yuko Sakurai, Atsushi Iwasaki, and Makoto Yokoo. Coalition Structure Generation Utilizing Compact Characteristic Function Representations. In the *Fifteenth International Conference on Principles and Practice of Constraint Programming (CP-09)*, pp. 623-638, Lisbon, Portugal, 2009. (Acceptance rate 42%.)
60. Vincent Conitzer. Prediction Markets, Mechanism Design, and Cooperative Game Theory. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence (UAI-09)*, pp. 101-108, Montreal, Canada, 2009. (Acceptance rate 12% for oral presentations.)
59. B. Paul Harrenstein, Mathijs M. de Weerd, and Vincent Conitzer. A Qualitative Vickrey Auction. In *Proceedings of the Tenth ACM Conference on Electronic Commerce (EC-09)*, pp. 197-206, Stanford, CA, USA, 2009. (Acceptance rate 25%.)
58. Vincent Conitzer, Matthew Rognlie, and Lirong Xia. Preference Functions That Score Rankings and Maximum Likelihood Estimation. In the *Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09)*, pp. 109-115, Pasadena, CA, USA, 2009. (Acceptance rate 26%.)
57. Vincent Conitzer, Jérôme Lang, and Lirong Xia. How hard is it to control sequential elections via the agenda? In the *Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09)*, pp. 103-108, Pasadena, CA, USA, 2009. (Acceptance rate 26%.)
56. Erik Halvorson, Vincent Conitzer, and Ronald Parr. Multi-step Multi-sensor Hider-Seeker Games. In the *Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09)*, pp. 159-166, Pasadena, CA, USA, 2009. (Acceptance rate 26%.)
55. Lirong Xia and Vincent Conitzer. Finite Local Consistency Characterizes Generalized Scoring Rules. In the *Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09)*, pp. 336-341, Pasadena, CA, USA, 2009. (Acceptance rate 26%.)

54. Lirong Xia, Michael Zuckerman, Ariel D. Procaccia, Vincent Conitzer, and Jeffrey Rosenschein. Complexity of Unweighted Coalitional Manipulation Under Some Common Voting Rules. In *Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09)*, pp. 348-353, Pasadena, CA, USA, 2009. (Acceptance rate 26%.)
53. Joseph Farfel and Vincent Conitzer. Turing Trade: A Hybrid of a Turing Test and a Prediction Market. In *Proceedings of the First Conference on Auctions, Market Mechanisms and Their Applications (AMMA-09)*, pp. 61-73, Boston, MA, USA, 2009. (Acceptance rate 47%.)
52. Vincent Conitzer. Anonymity-Proof Voting Rules. In *Proceedings of the Fourth Workshop on Internet and Network Economics (WINE-08)*, pp. 295-306, Shanghai, China, 2008. (Acceptance rate 33% for full-length papers.)
51. Krzysztof Apt, Vincent Conitzer, Mingyu Guo, and Evangelos Markakis. Welfare Undominated Groves Mechanisms. In *Proceedings of the Fourth Workshop on Internet and Network Economics (WINE-08)*, pp. 426-437, Shanghai, China, 2008. (Acceptance rate 33% for full-length papers.)
50. Joshua Letchford, Vincent Conitzer, and Kamal Jain. An “Ethical” Game-Theoretic Solution Concept for Two-Player Perfect-Information Games. In *Proceedings of the Fourth Workshop on Internet and Network Economics (WINE-08)*, pp. 696-707, Shanghai, China, 2008. (Acceptance rate 33% for full-length papers.)
49. Mingyu Guo and Vincent Conitzer. Better Redistribution with Inefficient Allocation in Multi-Unit Auctions with Unit Demand. In *Proceedings of the Ninth ACM Conference on Electronic Commerce (EC-08)*, pp. 210-219, Chicago, IL, USA, 2008. (Acceptance rate 19%.)
48. Lirong Xia and Vincent Conitzer. A Sufficient Condition for Voting Rules to Be Frequently Manipulable. In *Proceedings of the Ninth ACM Conference on Electronic Commerce (EC-08)*, pp. 99-108, Chicago, IL, USA, 2008. (Acceptance rate 19%.)
47. Lirong Xia and Vincent Conitzer. Generalized Scoring Rules and the Frequency of Coalitional Manipulability. In *Proceedings of the Ninth ACM Conference on Electronic Commerce (EC-08)*, pp. 109-118, Chicago, IL, USA, 2008. (Acceptance rate 19%.)
46. Liad Wagman and Vincent Conitzer. Optimal False-Name-Proof Voting Rules with Costly Voting. In *Proceedings of the 23rd National Conference on Artificial Intelligence (AAAI-08)*, pp. 190-195, Chicago, IL, USA, 2008. (Acceptance rate 24%. One of the two Outstanding Paper Awards.)
45. Lirong Xia and Vincent Conitzer. Determining Possible and Necessary Winners under Common Voting Rules Given Partial Orders. In *Proceedings of the 23rd National Conference on Artificial Intelligence (AAAI-08)*, pp. 196-201, Chicago, IL, USA, 2008. (Acceptance rate 24%.)
44. Lirong Xia, Vincent Conitzer, and Jérôme Lang. Voting on Multiattribute Domains with Cyclic Preferential Dependencies. In *Proceedings of the 23rd National Conference on Artificial Intelligence (AAAI-08)*, pp. 202-207, Chicago, IL, USA, 2008. (Acceptance rate 24%.)
43. Mingyu Guo and Vincent Conitzer. Optimal-in-Expectation Redistribution Mechanisms. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-08)*, pp. 1047-1054, Estoril, Portugal, 2008. (Acceptance rate 22%.)
42. Mingyu Guo and Vincent Conitzer. Undominated VCG Redistribution Mechanisms. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-08)*, pp. 1039-1046, Estoril, Portugal, 2008. (Acceptance rate 22%.)
41. Liad Wagman and Vincent Conitzer. Strategic Betting for Competitive Agents. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-08)*, pp. 847-854, Estoril, Portugal, 2008. (Acceptance rate 22%.)

40. Naoki Ohta, Vincent Conitzer, Yasufumi Satoh, Atsushi Iwasaki, and Makoto Yokoo. Anonymity-Proof Shapley Value: Extending Shapley Value for Coalitional Games in Open Environments. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-08)*, pp. 927-934, Estoril, Portugal, 2008. (Acceptance rate 22%. Winner of the Pragnesh Jay Modi Best Student Paper Award.)
39. Vincent Conitzer. Limited Verification of Identities to Induce False-Name-Proofness. In *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-07)*, pp. 102-111, Brussels, Belgium. (Acceptance rate 32%.)
38. Mingyu Guo and Vincent Conitzer. Worst-Case Optimal Redistribution of VCG Payments. In *Proceedings of the Eighth ACM Conference on Electronic Commerce (EC-07)*, pp. 30-39, San Diego, CA, USA. (Acceptance rate 27%.)
37. Vincent Conitzer. Eliciting Single-Peaked Preferences Using Comparison Queries. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-07)*, pp. 408-415, Honolulu, HI, USA, 2007. (Acceptance rate 22%.)
36. Vincent Conitzer and Tuomas Sandholm. Incremental Mechanism Design. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-07)*, pp. 1251-1256, Hyderabad, India, 2007. (Acceptance rate 35%.)
35. Tuomas Sandholm, Vincent Conitzer, and Craig Boutilier. Automated Design of Multistage Mechanisms. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-07)*, pp. 1500-1506, Hyderabad, India, 2007. (Acceptance rate 16% for oral presentations.)
34. Vincent Conitzer. Computing Slater Rankings Using Similarities Among Candidates. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI-06)*, pp. 613-619, Boston, MA, USA, 2006. (Acceptance rate 30%. Early version: IBM Research Report RC23748.)
33. Vincent Conitzer, Andrew Davenport, and Jayant Kalagnanam. Improved Bounds for Computing Kemeny Rankings. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI-06)*, pp. 620-626, Boston, MA, USA, 2006. (Acceptance rate 22% for oral presentations.)
32. Vincent Conitzer and Tuomas Sandholm. Nonexistence of Voting Rules That Are Usually Hard to Manipulate. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI-06)*, pp. 627-634, Boston, MA, USA, 2006. (Acceptance rate 22% for oral presentations.)
31. Naoki Ohta, Atsushi Iwasaki, Makoto Yokoo, Kohki Maruono, Vincent Conitzer, and Tuomas Sandholm. A Compact Representation Scheme for Coalitional Games in Open Anonymous Environments. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI-06)*, pp. 697-702, Boston, MA, USA, 2006. (Acceptance rate 30%.)
30. Vincent Conitzer and Nikesh Garera. Learning Algorithms for Online Principal-Agent Problems (and Selling Goods Online). In *Proceedings of the 23rd International Conference on Machine Learning (ICML-06)*, pp. 209-216, Pittsburgh, PA, USA, 2006. (Acceptance rate < 20%.)
29. Vincent Conitzer and Tuomas Sandholm. Computing the Optimal Strategy to Commit To. In *Proceedings of the Seventh ACM Conference on Electronic Commerce (EC-06)*, pp. 82-90, Ann Arbor, Michigan, USA, 2006. (Acceptance rate 29%.)
28. Vincent Conitzer and Tuomas Sandholm. A Technique for Reducing Normal-Form Games to Compute a Nash Equilibrium. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-06)*, pp. 537-544, Hakodate, Japan, 2006. (Acceptance rate 11% for oral presentations. One of four runners-up for the Best Student Paper Award.)

27. Vincent Conitzer and Tuomas Sandholm. Failures of the VCG Mechanism in Combinatorial Auctions and Exchanges. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-06)*, pp. 521-528, Hakodate, Japan, 2006. (Acceptance rate 23%.)
26. Vincent Conitzer and Tuomas Sandholm. Common Voting Rules as Maximum Likelihood Estimators. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI-05)*, pp. 145-152, Edinburgh, Scotland, UK, 2005. (Acceptance rate 9% for plenary presentations.)
25. Vincent Conitzer and Tuomas Sandholm. Expressive Negotiation in Settings with Externalities. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-05)*, pp. 255-260, Pittsburgh, PA, USA, 2005. (Acceptance rate 28%.)
24. Vincent Conitzer and Tuomas Sandholm. A Generalized Strategy Eliminability Criterion and Computational Methods for Applying It. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-05)*, pp. 483-488, Pittsburgh, PA, USA, 2005. (Acceptance rate 28%.)
23. Vincent Conitzer, Tuomas Sandholm, and Paolo Santi. Combinatorial Auctions with k -wise Dependent Valuations. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-05)*, pp. 248-254, Pittsburgh, PA, USA, 2005. (Acceptance rate 18% for oral presentations.)
22. Makoto Yokoo, Vincent Conitzer, Tuomas Sandholm, Naoki Ohta, and Atsushi Iwasaki. Coalitional Games in Open Anonymous Environments. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-05)*, pp. 509-514, Pittsburgh, PA, USA, 2005. (Acceptance rate 18% for oral presentations. This paper was also presented at the 19th Annual Conference of the Japan Society for Artificial Intelligence (JSAI-05) where it was one of five Awarded Papers.)
21. Tuomas Sandholm, Andrew Gilpin, and Vincent Conitzer. Mixed-Integer Programming Methods for Finding Nash Equilibria. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-05)*, pp. 495-501, Pittsburgh, PA, USA, 2005. (Acceptance rate 18% for oral presentations.)
20. Vincent Conitzer and Tuomas Sandholm. Complexity of (Iterated) Dominance. In *Proceedings of the Sixth ACM Conference on Electronic Commerce (EC-05)*, pp. 88-97, Vancouver, Canada, 2005. (Acceptance rate 29%.)
19. Vincent Conitzer and Tuomas Sandholm. Communication Complexity of Common Voting Rules. In *Proceedings of the Sixth ACM Conference on Electronic Commerce (EC-05)*, pp. 78-87, Vancouver, Canada, 2005. (Acceptance rate 29%.)
18. Paolo Santi, Vincent Conitzer, and Tuomas Sandholm. Towards a Characterization of Polynomial Preference Elicitation with Value Queries in Combinatorial Auctions. In *Proceedings of the 17th Annual Conference on Learning Theory (COLT-04)*, pp. 1-16, Banff, Alberta, Canada, 2004. (Acceptance rate 40%.)
17. Vincent Conitzer and Tuomas Sandholm. Communication Complexity as a Lower Bound for Learning in Games. In *Proceedings of the 21st International Conference on Machine Learning (ICML-04)*, pp. 185-192, Banff, Alberta, Canada, 2004. (Acceptance rate 32%.)
16. Vincent Conitzer and Tuomas Sandholm. An Algorithm for Automatically Designing Deterministic Mechanisms without Payments. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-04)*, pp. 128-135, New York, NY, USA, 2004. (Acceptance rate 24%.)
15. Vincent Conitzer and Tuomas Sandholm. Computing Shapley Values, Manipulating Value Division Schemes, and Checking Core Membership in Multi-Issue Domains. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI-04)*, pp. 219-225, San Jose, California, USA, 2004. (Acceptance rate 27%.)

14. Vincent Conitzer, Jonathan Derryberry, and Tuomas Sandholm. Combinatorial Auctions with Structured Item Graphs. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI-04)*, pp. 212-218, San Jose, California, USA, 2004. (Acceptance rate 27%.)
13. Vincent Conitzer and Tuomas Sandholm. Expressive Negotiation over Donations to Charities. In *Proceedings of the Fifth ACM Conference on Electronic Commerce (EC-04)*, pp. 51-60, New York, NY, USA, 2004. (Acceptance rate 16%.)
12. Vincent Conitzer and Tuomas Sandholm. Self-Interested Automated Mechanism Design and Implications for Optimal Combinatorial Auctions. In *Proceedings of the Fifth ACM Conference on Electronic Commerce (EC-04)*, pp. 132-141, New York, NY, USA, 2004. (Acceptance rate 16%.)
11. Vincent Conitzer and Tuomas Sandholm. Automated Mechanism Design: Complexity Results Stemming from the Single-Agent Setting. In *Proceedings of the 5th International Conference on Electronic Commerce (ICEC-03)*, pp. 17-24, Pittsburgh, PA, USA, 2003.
10. Vincent Conitzer and Tuomas Sandholm. AWESOME: A General Multiagent Learning Algorithm that Converges in Self-Play and Learns a Best Response Against Stationary Opponents. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pp. 83-90, Washington, DC, USA, 2003. (Acceptance rate 32%.)
9. Vincent Conitzer and Tuomas Sandholm. BL-WoLF: A Framework For Loss-Bounded Learnability In Zero-Sum Games. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pp. 91-98, Washington, DC, USA, 2003. (Acceptance rate 32%.)
8. Vincent Conitzer, Jérôme Lang, and Tuomas Sandholm. How Many Candidates Are Needed to Make Elections Hard to Manipulate? In *Proceedings of the Ninth Conference on Theoretical Aspects of Rationality and Knowledge (TARK-03)*, pp. 201-214, Bloomington, Indiana, USA, 2003.
7. Vincent Conitzer and Tuomas Sandholm. Universal Voting Protocol Tweaks to Make Manipulation Hard. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, pp. 781-788, Acapulco, Mexico, 2003. (Acceptance rate 21%.)
6. Vincent Conitzer and Tuomas Sandholm. Complexity of Determining Nonemptiness of the Core. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, pp. 613-618, Acapulco, Mexico, 2003. (Acceptance rate 21%.)
5. Vincent Conitzer and Tuomas Sandholm. Complexity Results about Nash Equilibria. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, pp. 765-771, Acapulco, Mexico, 2003. (Acceptance rate 21%.)
4. Vincent Conitzer and Tuomas Sandholm. Definition and Complexity of Some Basic Metareasoning Problems. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, pp. 1099-1106, Acapulco, Mexico, 2003. (Acceptance rate 21%.)
3. Vincent Conitzer and Tuomas Sandholm. Complexity of Mechanism Design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI-02)*, pp. 103-110, Edmonton, Canada, 2002. (Acceptance rate 34%.)
2. Vincent Conitzer and Tuomas Sandholm. Complexity of Manipulating Elections with Few Candidates. In *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI-02)*, pp. 314-319, Edmonton, Canada, 2002. (Acceptance rate 6% for plenary presentations.)
1. Vincent Conitzer and Tuomas Sandholm. Vote Elicitation: Complexity and Strategy-Proofness. In *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI-02)*, pp. 392-397, Edmonton, Canada, 2002. (Acceptance rate 26%.)

Short archival conference papers

10. Vincent Conitzer. Using Human Cognitive Limitations to Enable New Systems. In *the Eighth AAAI Conference on Human Computation and Crowdsourcing (HCOMP-20)*, Blue Sky Ideas track, Hilversum, the Netherlands (virtually), 2020. (Best Blue Sky Idea Award. Note: due to a timing issue, the Blue Sky Ideas track papers do not appear in the proceedings.)
9. Vincent Conitzer. Designing Preferences, Beliefs, and Identities for Artificial Intelligence. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19) Senior Member / Blue Sky Track*, pp. 9755-9759, Honolulu, HI, USA, 2019.
8. Vincent Conitzer, Walter Sinnott-Armstrong, Jana Schaich Borg, Yuan Deng, and Max Kramer. Moral Decision Making Frameworks for Artificial Intelligence. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17) Senior Member / Blue Sky Track*, pp. 4831-4835, San Francisco, CA, USA, 2017. (Received a CCC Blue Sky Award.)
7. Vincent Conitzer, Markus Brill, and Rupert Freeman. Crowdsourcing Societal Tradeoffs. In *Proceedings of the Fourteenth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-15) Blue Sky Ideas track*, pp. 1213-1217, Istanbul, Turkey, 2015. (Counted as a full paper in the acceptance statistics; acceptance rate for full papers 25%.)
6. Sayan Bhattacharya, Dmytro Korzhyk, and Vincent Conitzer. Computing a Profit-Maximizing Sequence of Offers to Agents in a Social Network. Short (7-page) paper in *Proceedings of the Eighth Workshop on Internet and Network Economics (WINE-12)*, pp. 482-488, Liverpool, UK, 2012.
5. Vincent Conitzer. Discussion of “A conditional game for comparing approximations.” Discussion paper in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS-11)*, pp. 72-73, Fort Lauderdale, FL, USA, 2011.
4. Mingyu Guo and Vincent Conitzer. Computationally Feasible Automated Mechanism Design: General Approach and Case Studies. In the NECTAR track of the *Proceedings of the 24th National Conference on Artificial Intelligence (AAAI-10)*, pp. 1676-1679, Atlanta, GA, USA, 2010. (Acceptance rate 25%.)
3. Mingyu Guo and Vincent Conitzer. False-Name-Proofness with Bid Withdrawal. Short paper in *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-10)*, pp. 1475-1476, Toronto, ON, Canada, 2010. (Acceptance rate 44%.)
2. Vincent Conitzer. Computational Aspects of Mechanism Design. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-05)* (Doctoral Consortium overview abstract), pp. 1642-1643, Pittsburgh, PA, USA, 2005.
1. Vincent Conitzer and Tuomas Sandholm. Computational Criticisms of the Revelation Principle. Short paper at the *Fifth ACM Conference on Electronic Commerce (EC-04)*, pp. 262-263, New York, NY, USA, 2004. (Acceptance rate 32%.) Also orally presented at the *Sixth Conference on Logic and the Foundations of Game and Decision Theory (LOFT-04)*, Leipzig, Germany, 2004.

Nonarchival conference papers

2. Vincent Conitzer. Should Stackelberg Mixed Strategies Be Considered a Separate Solution Concept? The *Eleventh Conference on Logic and the Foundations of Game and Decision Theory (LOFT-14)*, Bergen, Norway, 2014. (Plenary presentation, acceptance rate 40%.)
1. Vincent Conitzer. Comparing Multiagent Systems Research in Combinatorial Auctions and Voting. The *Tenth International Symposium on Artificial Intelligence and Mathematics (ISAIM-08)*, invited to the Special Session on Computation and Social Choice, Fort Lauderdale, FL, USA.

Workshop papers

5. Vincent Conitzer. Why should we ever automate moral decision making? *Ethics and Trust in Human-AI Collaboration: Socio-Technical Approaches* (writeup for invited talk in the corresponding workshop), Macao, 2023.
4. Steven Jecmen, Nihar B. Shah, Fei Fang, and Vincent Conitzer. Tradeoffs in Preventing Manipulation in Paper Bidding for Reviewer Assignment. Workshop on ML Evaluation Standards at ICLR 2022. Outstanding Paper Award.)
3. Vincent Conitzer, Jana Schaich Borg, and Walter Sinnott-Armstrong. Using Human Subjects' Judgments for Automated Moral Decision Making. Whitepaper for the *Workshop on Trustworthy Algorithmic Decision-Making*, Arlington, VA, USA, 2017.
2. Vincent Conitzer. Using a Memory Test to Limit a User to One Account. The *10th International Workshop on Agent Mediated Electronic Commerce (AMEC-08)*, Estoril, Portugal, 2008. Appears in LNBP 44, Agent-Mediated Electronic Commerce and Trading Agent Design and Analysis, pp. 60-72.
1. Vincent Conitzer and Tuomas Sandholm. Applications of Automated Mechanism Design. The *UAI-03 Bayesian Modeling Applications Workshop*, Acapulco, Mexico, 2003.

Newsletter articles and puzzles

5. Vincent Conitzer. Puzzle: Does Occasional Simulation Enable Cooperation? (Puzzle in honor of Joe Halpern's 70th birthday.) *SIGecom Exchanges*, Vol. 21.1, June 2023, pages 62-63.
4. Vincent Conitzer. Puzzle: Communicating to Plan Noam Nisan's 60th Birthday Workshop. (Puzzle in honor of Noam Nisan's 60th birthday.) *SIGecom Exchanges*, Vol. 20.1, July 2022.
3. Vincent Conitzer. Puzzle: The AI Circus. (Puzzle in honor of Tuomas Sandholm's 50th birthday.) *SIGecom Exchanges*, Vol. 17.2, October 2019.
2. Vincent Conitzer and David Easley. Notes from the EC'14 Program Chairs. *SIGecom Exchanges*, Vol. 13, No. 1, June 2014, pp. 2-4.
1. Joshua Letchford, Liam MacDermid, Vincent Conitzer, Ronald Parr, and Charles Isbell. Computing Stackelberg Strategies in Stochastic Games. *SIGecom Exchanges*, Vol. 11, No. 2, December 2012, pp. 36-40.

Writing for broader audiences

6. Yoshua Bengio and Vincent Conitzer. What do large language models tell us about ourselves? *Institute for Ethics in AI blog*, 8 July 2024. <https://www.oxford-aiethics.ox.ac.uk/blog/what-do-large-language-models-tell-us-about-ourselves/>
5. Vincent Conitzer and Derek Leben. How ChatGPT has been prompted to respect safety, fairness, and copyright. *Institute for Ethics in AI blog*, 26 Feb 2024. <https://www.oxford-aiethics.ox.ac.uk/blog/how-chatgpt-has-been-prompted-to-respect-safety-fairness-and-copyright/>
4. Vincent Conitzer. Natural Intelligence Still Has Its Advantages. *The Wall Street Journal*, August 28, 2018. <https://www.wsj.com/articles/natural-intelligence-still-has-its-advantages-1535495256> A version of the article is also available on arXiv as Can Artificial Intelligence Do Everything That We Can?
3. Vincent Conitzer. The AI debate must stay grounded in reality. *Prospect* (in association with the British Academy), March 6, 2017. <http://www.prospectmagazine.co.uk/britishacademy/the-ai-debate-must-stay-grounded-in-reality>
2. Vincent Conitzer. Today's Artificial Intelligence Does Not Justify Basic Income. *MIT Technology Review*, October 31, 2016. <https://www.technologyreview.com/s/602747/todays-artificial-intelligence-does-not-justify-basic-income/>

1. Vincent Conitzer. Artificial intelligence: where's the philosophical scrutiny? *Prospect*, 2016-5-4.
<http://www.prospectmagazine.co.uk/science-and-technology/artificial-intelligence-where-s-the-philosophical-scrutiny>
An unedited version of the article is also available on arXiv as Philosophy in the Face of Artificial Intelligence.

Edited Volumes

- Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society (AIES'22). Co-edited with John Tasioulas, Matthias Scheutz, Ryan Calo, Martina Mara, and Annette Zimmermann.
- Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI'20). Co-edited with Fei Sha and Francesca Rossi.
- Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (AIES'19). Co-edited with Gillian Hadfield and Shannon Vallor.
- Proceedings of the Twelfth Workshop on the Economics of Networks, Systems and Computation (NetEcon'17). Co-edited with Roch Guérin.
- Proceedings of the Seventeenth ACM Conference on Economics and Computation (EC'16). Co-edited with Dirk Bergemann and Yiling Chen.
- ACM Transactions on Economics and Computation (TEAC), Volume 5, Issue 1, November 2016. Co-edited with Preston McAfee.
- ACM Transactions on Economics and Computation (TEAC) - Special Issue on EC'14, Volume 4, Issue 4, August 2016. Co-edited with Preston McAfee. Guest editors for this issue: Vincent Conitzer and David Easley.
- ACM Transactions on Economics and Computation (TEAC) - Special Issue on EC'13, Volume 4, Issue 3, June 2016. Co-edited with Preston McAfee. Guest editors for this issue: Preston McAfee and Éva Tardos.
- ACM Transactions on Economics and Computation (TEAC), Volume 4, Issue 2, February 2016. Co-edited with Preston McAfee.
- ACM Transactions on Economics and Computation (TEAC), Volume 4, Issue 1, December 2015. Co-edited with Preston McAfee.
- ACM Transactions on Economics and Computation (TEAC), Volume 3, Issue 4, July 2015. Co-edited with Preston McAfee. Contains special section on WINE'13 guest-edited by Yiling Chen and Nicole Immorlica.
- ACM Transactions on Economics and Computation (TEAC), Volume 3, Issue 3, June 2015. Co-edited with Preston McAfee.
- ACM Transactions on Economics and Computation (TEAC) - Special Issue on EC'12, Part 2; Volume 3, Issue 2, April 2015. Co-edited with Preston McAfee. Guest editors for this issue: Kevin Leyton-Brown and Panos Ipeirotis.
- ACM Transactions on Economics and Computation (TEAC) - Special Issue on EC'12, Part 1; Volume 3, Issue 1, March 2015. Co-edited with Preston McAfee. Guest editors for this issue: Kevin Leyton-Brown and Panos Ipeirotis.
- Proceedings of the Fifteenth ACM Conference on Economics and Computation (EC'14). Co-edited with Moshe Babaioff and David Easley.

- ACM Transactions on Economics and Computation (TEAC), Volume 2, Issue 4, October 2014. Co-edited with Preston McAfee.
- ACM Transactions on Economics and Computation (TEAC), Volume 2, Issue 3, July 2014. Co-edited with Preston McAfee.
- ACM Transactions on Economics and Computation (TEAC), Volume 2, Issue 2, June 2014. Co-edited with Preston McAfee.
- ACM Transactions on Economics and Computation (TEAC), Volume 2, Issue 1, March 2014. Co-edited with Preston McAfee.
- ACM Transactions on Economics and Computation (TEAC), Volume 1, Issue 4, December 2013. Co-edited with Preston McAfee.
- ACM Transactions on Economics and Computation (TEAC), Volume 1, Issue 3, September 2013. Co-edited with Preston McAfee.
- ACM Transactions on Economics and Computation (TEAC) - Special Issue on Algorithmic Game Theory, Volume 1, Issue 2, May 2013. Co-edited with Preston McAfee. Guest editors for this issue: Michal Feldman and Noam Nisan.
- ACM Transactions on Economics and Computation (TEAC), Volume 1, Issue 1, January 2013. Co-edited with Preston McAfee.
- Proceedings of the Eleventh International Conference on Autonomous Agents and Multi Agent Systems (AAMAS 2012). Co-edited with Michael Winikoff, Wiebe van der Hoek, and Lin Padgham.
- SIGecom Exchanges, Volume 10.1 (March, 2011). Co-edited with Yiling Chen. (Co-contributed introduction.)
- Vincent Conitzer and Jörg Rothe (editors). *Proceedings of the Third International Workshop on Computational Social Choice (COMSOC-2010)*, printed by Düsseldorf University Press, September 2010.
- SIGecom Exchanges, Volume 9.1 (June, 2010). (Contributed introduction and puzzle.)
- SIGecom Exchanges, Volume 8.2 (December 2009). (Contributed introduction and puzzle.)
- SIGecom Exchanges, Volume 8.1 (July 2009). (Contributed introduction and puzzle.)
- SIGecom Exchanges, Volume 7.3 (November 2008). (Contributed introduction and puzzle.)
- SIGecom Exchanges, Volume 7.2 (June 2008). (Contributed introduction and puzzle.)
- SIGecom Exchanges, Volume 7.1 (December 2007), special issue on combinatorial auctions. (Contributed introduction and puzzle.)

Competition Entries

- Vincent Conitzer. Theory of Conscious Experience. Meritorious Prize Winner (among top 7 entries out of 800+ submissions) in the NSF 2026 Idea Machine competition.

Tutorials

- Vincent Conitzer and Caspar Oesterheld. *Foundations of Cooperative AI*. Tutorial at the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'25), Detroit, MI, USA, May 2025.
- Vincent Conitzer. *Designing Agents' Preferences, Beliefs, and Identities*. Tutorial at the Safe and Trusted AI (STAI) summer school, London, UK, July 2023.
- Vincent Conitzer. *Designing Agents' Preferences, Beliefs, and Identities*. Tutorial at the 31st International Joint Conference on Artificial Intelligence (IJCAI-ECAI'22), Vienna, Austria, July 2022.
- Vincent Conitzer. *Designing Agents' Preferences, Beliefs, and Identities*. Tutorial at the 22nd ACM Conference on Economics and Computation (EC'21), virtual, July 2021.
- Vincent Conitzer. *Designing Agents' Preferences, Beliefs, and Identities*. Tutorial at the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-21), virtual, May 2021.
- Vincent Conitzer. *Designing Agents' Preferences, Beliefs, and Identities*. Tutorial at the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21), virtual, February 2021.
- Vincent Conitzer. *Designing Agents' Preferences, Beliefs, and Identities*. Tutorial at the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-19), Montréal, Québec, Canada, May 2019.
- Vincent Conitzer. *Computational Social Choice and Moral Artificial Intelligence*. Invited tutorial at the 27th International Joint Conference on Artificial Intelligence and the 23rd European Conference on Artificial Intelligence (IJCAI-ECAI-18), Stockholm, Sweden, July 2018.
- Vincent Conitzer. *Crash Course on Computational Social Choice and Fair Division*. Invited tutorial at the mentoring workshop at the 19th ACM Conference on Economics and Computation (EC-18), Ithaca, NY, USA, June 2018.
- Vincent Conitzer. *Computing Game-Theoretic Solutions*. Invited tutorial at the International Summer School on Autonomous Agents and Multiagent Systems, Singapore, May 2016.
- Vincent Conitzer. *Some Game-Theoretic Aspects of Voting*. Invited tutorial at WINE, Amsterdam, the Netherlands, December 2015.
- Vincent Conitzer. *Computing Game-Theoretic Solutions*. Tutorial at AAAI, Austin, TX, USA, January 2015.
- Vincent Conitzer. *Computing Game-Theoretic Solutions*. Invited tutorial at NeurIPS, Montréal, Québec, Canada, December 2014.
- Vincent Conitzer. *Computational Social Choice*. Tutorial at the Summer School on Algorithmic Economics, Carnegie Mellon University, Pittsburgh, PA, USA, August 2012.
- Vincent Conitzer. *A Brief Introductory Tutorial on Computational Social Choice*. Tutorial at the Third International Workshop on Computational Social Choice (COMSOC-10), Düsseldorf, Germany, 2010.
- Vincent Conitzer and Ariel Procaccia. *Computational Voting Theory*. Tutorial at the 11th ACM Conference on Electronic Commerce (EC-10), Cambridge, MA, USA, 2010.
- Vincent Conitzer and Ariel Procaccia. *Computational Voting Theory*. Tutorial at the 9th International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-10), Toronto, ON, Canada, 2010.

- Vincent Conitzer and Yevgeniy Vorobeychik. *Automated Mechanism Design: Methods and Applications*. Tutorial at the 21st International Joint Conference on Artificial Intelligence (IJCAI-09), Pasadena, CA, USA, 2009.
- Vincent Conitzer and Yevgeniy Vorobeychik. *Automated Mechanism Design: Methods and Applications*. Tutorial at the 8th International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-09), Budapest, Hungary, 2009.
- Vincent Conitzer. *Game Theoretic Solutions and How to Compute Them*. Invited tutorial at the DIMACS/LAMSADE workshop on algorithmic decision theory, Université Paris Dauphine, Paris, France, October 2008.
- Vincent Conitzer and Yevgeniy Vorobeychik. *Automated Mechanism Design: Approaches and Applications*. Tutorial at the 9th ACM Conference on Electronic Commerce (EC-08), Chicago, IL, USA, 2008.
- Vincent Conitzer. *Mechanism Design for Multi Agent Systems*. Invited tutorial at the Dubai Agents and Multi Agent Systems School (DAMAS), Dubai, UAE, January 2008.
- Vincent Conitzer. *Mechanism Design for Computer Scientists*. Invited tutorial at the North East Student Colloquium on Artificial Intelligence (NESCAI), Ithaca, NY, USA, April 2006.

Invited talks

(Talks in workshops and conferences that were in large part organized by invitation and had many speakers are in the next category. Invited tutorials are in the previous category.)

- The First International Workshop on Agentic Intelligence: Risks, Ethics, and Trust (AIRET), Pittsburgh, PA, USA, November 2025, *Game Theory for AI Agents*.
- The First International Workshop on Agentic Intelligence: Risks, Ethics, and Trust (AIRET), Pittsburgh, PA, USA, November 2025, panel on *The Human Role in Agentic AI Futures* (with Ross McGowan and Thanh Tran; moderated by Yu-Ru Lin).
- UK AISI Alignment Conference, London, UK, October 2025, *Game Theory for AI Agents*.
- Philosophy of AI: A Workshop at Columbia University, New York, NY, USA / online, October 2025, *What, if anything, can we say about consciousness in LLMs?*
- Cooperative AI summer school, Marlow, UK, July 2025, *Foundations and Frontiers of Cooperative AI*.
- Cooperative AI summer retreat, Marlow, UK, July 2025, *Cooperative AI via Simulations*.
- Nash 75 (75 years of Nash equilibrium) conference, Oxford, UK, July 2025, *Game Theory for AI Agents*.
- Trajan House Seminar Series, Oxford, UK, June 2025, *What (if anything) can we say about consciousness in large language models?*
- AAMAS Detroit Cooperative AI Panel, Detroit, MI, USA, May 2025, *What research will matter most for an age of autonomous agents in mixed-motive settings?* (with Kate Larson, Peter Railton, and Virginia Dignum; moderated by Cecilia Tilli).
- Social Choice for AI (SC4AI) workshop at AAMAS, Detroit, MI, USA, May 2025, panel discussion (with Smitha Milli and Zoi Terzopoulou; moderated by Jobst Heitzig).
- Neuroethics Working Group (NEWG) Meeting, virtual, May 2025, *Ethical considerations related to moral decision making in digital brain twins* (with Jana Schaich Borg and Walter Sinnott-Armstrong).

- Conference on Tests of Consciousness, Duke University, Durham, NC, USA, April 2025, *Do existing LLMs already have consciousness? Reasons to be skeptical.*
- AI, Systems, and Society Conference at Emory University, Atlanta, GA, USA, March 2025, *Getting AI systems (collectively) to behave the way we (collectively) want them to.*
- International Association for Safe & Ethical AI Conference, Paris, France, February 2025, *Getting AI systems (collectively) to behave the way we (collectively) want them to.* (Both a talk and a poster presentation.)
- Online Social Choice and Welfare Seminar Series, October 2024, *Getting AI systems (collectively) to behave the way we (collectively) want them to.*
- Nasscom Future Forge 2024 keynote talk, Bangalore, India (talk delivered remotely), October 2024, *What should we want from AI, and how can India lead?*
- University at Buffalo Institute for Artificial Intelligence and Data Science Distinguished Speaker Series, Buffalo, NY, USA, October 2024, *Getting AI Systems (Collectively) to Behave the Way We (Collectively) Want Them To.*
- Vienna Alignment Workshop, Vienna, Austria, July 2024, *Game Theory and Social Choice for Cooperative AI* (invited lightning talk).
- Pittsburgh Humanities Festival, Pittsburgh, PA, USA, April 2024, *Moral AI And How We Get There.*
- The University of Graz, Graz, Austria, inaugural event in series on engaging with artificial intelligence and related technologies in higher education, March 2024, *The Impact of Recent Progress in AI on Education and Society* (keynote speaker; lecture delivered remotely).
- 22nd International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT), Stony Brook, NY, USA, March 2024, *AI and Ethics* (invited plenary talk). (Followed by Q&A panel that also included Aleksandra M. Ciprijanovic, Elizabeth S. Sexton-Kennedy, and Gordon T. Watts; moderated by Jennifer Ngadiuba.)
- National AI Advisory Committee (NAIAC) public discussion on AI Safety, online, March 2024, panel 1 on *Concepts in AI Safety: How should we be scoping “safety” in AI?* (with Inioluwa Deborah Raji, Chris Meserole, Arvind Narayanan, Julia Angwin, John C. Inglis, Suresh Venkatasubramanian; facilitated by Janet Haven).
- University of Toronto Schwartz Reisman Institute for Technology and Society (SRI) Seminar Series, online, February 2024, *Social choice and game theory for AI alignment.*
- Summit on Existential Security, January 2024, *Social Choice for AI Ethics and Safety* (invited lightning talk).
- Keystone Strategy, San Francisco, CA, USA, January 2024, *Moral AI and How We Get There.*
- Carnegie Mellon University Faculty Dialogues, *Democratizing the Future*, online, December 2023 (with Subha Das and Rema Padman; moderated by Martial Hebert).
- New Orleans Alignment Workshop, New Orleans, LA, USA, December 2023, *Foundations of Cooperative AI* (invited lightning talk).
- *Artificially Intelligent, Naturally Fair? Panel Discussion on Equitable Algorithms*, virtual (organized by UMass), October 2023 (with Alexandra Olteanu, Monideepa Tarafdar, and Yair Zick; moderated by Ethan Zuckerman).
- Oxford in New York, New York City, NY, USA, October 2023, *Ethics in AI - Who Should Set the Rules, and How?*

- Oxford welcome lecture for philosophy joint schools involving computer science, mathematics, and physics, October 2023, Oxford, UK / online, *Ethics in AI and Cooperative AI*.
- Ethics and Trust in Human-AI Collaboration: Socio-Technical Approaches (EthAIcs@IJCAI), Macao, August 2023, *Automated Moral Decision Making by Learning from Humans: Why and How*.
- Keystone Advanced Technology Services (K.ATS), London, UK, July 2023, *Managing the Societal Impact of AI Systems*.
- Summer School on Computational Social Choice, Amsterdam, the Netherlands, July 2023, *Social Choice for Moral AI*.
- Cooperative AI summer school, London, UK, July 2023, *Foundations of Cooperative AI*.
- Cooperative AI summer retreat, London, UK, July 2023, panel on *Criteria for Cooperative AI* (with Joel Leibo, Nika Haghtalab, Thore Graepel; moderated by Jesse Clifton).
- Brookings CRM Roundtable on Frontier AI Regulation/Government Capacity Building, June 2023, panel on *Frontier AI Regulation* (with Jade Leung, Greg Allen, and Adam Thierer; moderated by Anton Korinek).
- Research!America Alliance discussion, May 2023.
- Keystone Strategy event at the London Business School Institute of Entrepreneurship and Private Capital – Competition, Digital Regulation and Technology: Algorithmic Auditing in the Age of Generative AI & Privacy Enhancing Technologies, April 2023, panel on *The future of algorithmic audit: what role will it play in competition and regulation?* (with Amy Jordan, Jessica Rusu, Simon McDougall, and Aparna Surendra; chaired by Stefan Hunt).
- AI in Business and Economics: The Economic Perspective on Artificial Intelligence (EPEAI), March 2023, Ruhr West University of Applied Sciences, Germany (and online), invited keynote, *AI, Economics and Philosophy*.
- San Francisco Alignment Workshop, San Francisco, CA, USA, February 2023, *Foundations of Cooperative AI Lab* (invited lightning talk).
- Workshop on Artificial Intelligence Safety (SafeAI) at the Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI-23), Washington, DC, USA, February 2023, invited speaker, *Foundations of Cooperative AI*.
- AI for Behavior Change Workshop (AI4BC) at the Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI-23), Washington, DC, USA, February 2023, invited panel on *Reinforcing Preferences vs Behavior Change* (with Lisa Singh, Bennett Butler, and Dylan Hadfield-Menell; moderated by Sabina Tomkins).
- The Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI-23), Washington, DC, USA, February 2023, plenary speaker, *New Design Decisions for Modern AI Agents*.
- The 18th Conference on Web and Internet Economics (WINE-22), Troy, NY, USA, December 2022, invited speaker, *Automated Mechanism Design for Strategic Classification*.
- University of Nebraska-Lincoln School of Computing Colloquium Series, Lincoln, NE, USA, November 2022, *Designing Agents' Preferences, Beliefs, and Identities*.
- ICDS Fall 2022 Symposium: Data Science, AI, and a Sustainable, Resilient, and Equitable Future, Pennsylvania State University, State College, PA, USA, October 2022, *Fireside Chat: Responsible AI & AI Bill of Rights* (moderated by Hadi Hosseini and Jennifer Wagner).

- Special Distinguished AI Seminar, Pennsylvania State University, State College, PA, USA, October 2022, *Designing Agents' Preferences, Beliefs, and Identities*.
- Keystone Strategy FATE (Algorithmic Fairness, Auditing, Transparency, and Externalities) Speaker Series, San Francisco, CA, USA, October 2022, *Managing the Societal Impact of AI Systems*.
- Principles of Intelligent Behavior in Biological and Social Systems (PIBBSS) speaker series, July 2022, *AI Agents May Cooperate Better If They Don't Resemble Us*.
- Artificial Intelligence and Business Analytics Workshop, Frankfurt School of Finance & Management, Frankfurt, Germany, July 2022, *Automated Mechanism Design for Strategic Classification*.
- ACM's 75th Anniversary Celebration, San Francisco, CA, June 2022, panel on *Incentives and Markets* (together with Susan Athey, Nicole Immorlica, Sendhil Mullainathan, and Hal Varian; moderated by Kevin Leyton-Brown).
- Center for Human-Compatible AI (CHAI) 2022 workshop, Asilomar Conference Grounds, CA, June 2022, invited session talk, *AI Agents May Cooperate Better If They Don't Resemble Us*.
- AAMAS Workshop on Autonomous Agents for Social Good (AASG), May 2022, *Designing Agents for Social Good*.
- FHI DeepMind AI Safety Seminar, May 2022, *AI Agents May Cooperate Better If They Don't Resemble Us*.
- University of Toronto Department of Computer Science Distinguished Lecture Series, April 2022, *AI Agents May Cooperate Better If They Don't Resemble Us*.
- Duke Kunshan University Computational Microeconomics Seminar Series of Undergraduate Studies, April 2022, *AI Agents May Cooperate Better If They Don't Resemble Us*.
- Existential and Societal-Scale Safety for AI (ESSAI) Colloquium, April 2022, *AI Agents May Cooperate Better If They Don't Resemble Us*.
- University of Virginia Human and Machine Intelligence Seminar, March 2022, *AI Agents May Cooperate Better If They Don't Resemble Us*.
- New Directions in Cooperative AI seminar series (inaugural seminar), January 2022, *AI Agents May Cooperate Better If They Don't Resemble Us*. (With Edith Elkind and Joe Halpern as discussants.)
- NeurIPS Workshop on Learning in Presence of Strategic Behavior, December 2021, *Automated Mechanism Design for Strategic Classification*. (Including discussion afterwards with Dorsa Sadigh, moderated by Peter Stone.)
- Absolutely Interdisciplinary, November 2021, *Computational Ethics: Does AI create radically new problems in ethics?* (Together with Johanna Thoma; moderated by Jennifer Nagel.) (Also participated on panel "Introducing the Cooperative AI Foundation and the Collective Intelligence Journal" with Jessica Flack, Eric Horvitz, Ruairi Donnelly, and Jesse Clifton; moderated by Gillian Hadfield.)
- Conference on Decision and Game Theory for Security (GameSec) 2021, keynote talk, October 2021, *AI Agents May Cooperate Better If They Don't Resemble Us*.
- Duke Forever Learning event, September 2021, *Morality in AI*. (Together with Jana Schach Borg; moderated by Walter Sinnott-Armstrong.)
- The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD) 2021, keynote talk, August 2021, *Automated Mechanism Design for Strategic Classification*.

- Summer School on Game Theory and Social Choice (organized by City University of Hong Kong), June 2021, *New Design Decisions for Modern AI Agents*.
- AI Ethics Panel at ML-LABS Ireland, June 2021, *Living with AI: Should machines make decisions for us?* (Together with Susan Leavy, Margaret Mitchell, Eugenia Siapera; moderated by Marie Boran.)
- Israel Algorithmic Game Theory Seminar, May 2021, *Automated Mechanism Design for Strategic Classification*.
- AAMAS 2021 keynote talk, May 2021, *New Design Decisions for Modern AI Agents*. (Talk associated with ACM/SIGAI Autonomous Agents Research Award.)
- The City University of New York (CUNY) Graduate Center, City of Science series, April 2021, *The Ethics of AI*. (Together with Gillian Hadfield, Francesca Rossi, Julie Shah; moderated by Susan Epstein.)
- Carnegie Mellon University School of Computer Science Special Seminar, April 2021, *New Design Decisions for Modern AI Agents*.
- STEM World II, March 2021, *Artificial Intelligence from Societal, Ethical, and Philosophical Perspectives*.
- Lindgren Lecture, Shared Learning Association of Chapel Hill, February 2021, *Artificial Intelligence from Societal, Ethical, and Philosophical Perspectives*.
- SIGecom Winter Meeting, February 2021, *Automated Mechanism Design for Correlated Valuations*.
- Drexel University College of Computing and Informatics, Department of Computer Science Dr. Jay Modi Lecture, February 2021, *Designing Agents' Preferences, Beliefs, and Identities*.
- Carnegie Mellon University Artificial Intelligence Seminar, December 2020, *Automated Mechanism Design for Strategic Classification*.
- MTL Connect, online, October 2020, *Moral Artificial Intelligence*.
- Center for Human-Compatible AI (CHAI) 2020 workshop, online over Zoom, June 2020, invited plenary talk, *New Directions in Belief Formation and Decision Theory for AI*.
- COMSOC Video Seminar (International Seminar Series on Social Choice), online over Zoom, April 2020, *Computational Social Choice for Moral Artificial Intelligence*.
- North Carolina State University ECE Interdisciplinary Distinguished Seminar Series (IDSS), Raleigh, NC, USA, January 2020, *Designing Agents' Preferences, Beliefs, and Identities*.
- Duke Radiology Grand Rounds Lecture Series, Durham, NC, USA, September 2019, *The Future of AI in Radiology*. (Together with Walter Sinnott-Armstrong, moderated by Lawrence Ngo.)
- University of Oxford Department of Computer Science Departmental Seminar, Oxford, England, UK, September 2019, *Designing Agents' Preferences, Beliefs, and Identities*.
- Theoretical Aspects of Rationality and Knowledge (TARK) conference, invited speaker, Toulouse, France, July 2019, *Designing Agents' Preferences, Beliefs, and Identities*.
- DARPA SI3-CMD program meeting guest speaker, Durham, NC, USA, July 2019, *Designing Agents and Their Objectives: A Game-Theoretic Perspective*.
- STOC workshop on New Frontiers of Automated Mechanism Design for Pricing and Auctions, Phoenix, AZ, USA, June 2019, *New Directions in Automated Mechanism Design*.

- University of Chicago CS Distinguished Lecturer Series, Chicago, IL, USA, May 2019, *Moral Artificial Intelligence: How to Learn Objectives from People*.
- North Carolina State University, Distinguished Lecture on the Ethics of AI, Raleigh, NC, USA, April 2019, *Should Moral AI Help Us Decide Who Gets a Kidney?* (Together with Jana Schaich Borg and Walter Sinnott-Armstrong.)
- Bahamas, March 2019, *Artificial Intelligence and Its Impact on Society*.
- Congressional briefing on the ethics of AI, Washington, DC, USA, February 2019. (Together with Walter Sinnott-Armstrong and Nita Farahany; organized by Duke in DC.)
- Envision Conference, Princeton, NJ, USA, December 2018, *Moral Artificial Intelligence*.
- Ninth Annual McGowan Symposium on Business Leadership & Ethics, Durham, NC, USA, November 2018, *Building Ethics into Robust Artificial Intelligence*.
- Cary Academy National Honor Society Speaker Series, October 2018, Cary, NC, USA, event discussing the benefits of and concerns about AI. (Together with Behnam Kia.)
- Université de Montréal, département d'informatique et de recherche opérationnelle, September 2018, Montréal, Québec, Canada, *Moral Artificial Intelligence: How to Learn Objectives from People*.
- Block Seoul, September 2018, Seoul, South Korea, *Moral Artificial Intelligence*. (Also served on the panel *Designing Our Future: Technology and Humanity*.)
- Centre for European Policy Studies (CEPS) Task Force: Artificial Intelligence: Ethics, Governance, and Policy Challenges, July 2018, Brussels, Belgium, *Moral AI: at the crossroads between philosophy, computer science, and economics*.
- AAAI 2018 Spring Symposium on AI and Society: Ethics, Safety and Trustworthiness in Intelligent Agents, March 2018, Stanford University, CA, USA, *Moral Artificial Intelligence, Kidney Exchanges, and Societal Tradeoffs*.
- Workshop on AI and Marketing (at AAAI 2018), February 2018, New Orleans, LA, USA, *Multiplicative Pacing Equilibria in Auction Markets*.
- Chicago Council on Global Affairs, January 2018, Chicago, IL, USA, *Artificial Intelligence: Hype, Hope, or Hazard?* (Together with Dave Gunning; moderated by Cécile Shea.)
- Northwestern University Computer Science Seminar, January 2018, Evanston, IL, USA, *Moral Artificial Intelligence, Kidney Exchanges, and Societal Tradeoffs*.
- Chinese American Friendship Association, October 2017, Raleigh, NC, USA, *Artificial Intelligence: Past, Present, and Future*.
- North Carolina Museum of Natural Sciences, Science Cafe, October 2017, Raleigh, NC, USA, *Who Needs a Driver?* (Together with Walter Sinnott-Armstrong and Jana Schaich Borg.)
- Applied Mathematics, Modeling and Computational Science (AMMCS) Plenary Talk, August 2017, Waterloo, Ontario, Canada, *Moral Artificial Intelligence and the Societal Tradeoffs Problem*.
- Duke in DC, May 2017, Washington, DC, USA, *Ethical Policy Conversations: Jobs and Artificial Intelligence at an Intersection*. (Together with Andrea D. Martin and M-H. Carolyn Nguyen; moderated by Sean Illing.)

- Centre for European Policy Studies (CEPS) Ideas Lab, February 2017, Brussels, Belgium, Prime Talk: *Ex machina: Towards an ethics of algorithms*. (Together with Tim Hwang, Jérôme Perrin, Walter Sinnott-Armstrong, and Toby Walsh; moderated by Andrea Renda.)
- Meeting of the Professional Engineers of North Carolina (PENC) and the American Society of Highway Engineers (ASHE), Raleigh, NC, USA, January 2017, *Driverless Cars: Are We Ready to Hand Over the Wheel?* (Together with Michael Clamann and Jana Schaich Borg.)
- McGill University CS Colloquium, Montréal, Québec, Canada, December 2016, *Moral Artificial Intelligence and the Societal Tradeoffs Problem*.
- Vanderbilt University EECS Lecture, Nashville, TN, USA, September 2016, *Crowdsourcing Societal Tradeoffs*.
- RTI International, Durham, NC, USA, September 2016, *Computing Game-Theoretic Solutions and Applications to Security*.
- Booking.com, Amsterdam, the Netherlands, July 2016, *Vickrey Auctions with Flexible Types of Payment*.
- The Tenth Workshop on Advances in Preference Handling (MPREF) at IJCAI-16, New York, NY, USA, *Mechanism Design in Data-Rich Environments*.
- The AAMAS-16 Workshop on Security and Multi-agent Systems (SecMAS), Singapore, May 2016, *Nash Equilibria of Security Games*.
- Stanford University RAIN (Research on Algorithms and Incentives in Networks) seminar, October 2015, *Crowdsourcing Societal Tradeoffs*.
- Simons Institute EconCS Survey Seminar, Berkeley, CA, September 2015, *Computing Stackelberg Mixed Strategies and Applications to Security*.
- Facebook, Menlo Park, CA, USA, August 2015, *Vickrey Auctions with Flexible Types of Payment*.
- Amazon Tech Talk, Seattle, WA, USA, August 2015, *Vickrey Auctions with Flexible Types of Payment*.
- The first IJCAI workshop on Algorithmic Game Theory (AGT@IJCAI-2015), Buenos Aires, Argentina, July 2015, *Crowdsourcing Societal Tradeoffs: Strategic Aspects*.
- IJCAI workshop on Behavioral, Economic and Computational Intelligence for Security (BECIS@IJCAI-2015), Buenos Aires, Argentina, July 2015, *Time and Signaling in Stackelberg and Security Games*.
- Lorentz Center workshop on Clusters, Games and Axioms, Leiden, the Netherlands, June 2015, *Crowdsourcing Societal Tradeoffs*.
- Carnegie Mellon University AI seminar, Pittsburgh, PA, USA, April 2015, *Computational Social Choice: A Journey from Basic Complexity Results to a Brave New World for Social Choice*.
- University of Toronto AI seminar, Toronto, ON, Canada, March 2015, *Computational Social Choice: A Journey from Basic Complexity Results to a Brave New World for Social Choice*.
- Museum of Life and Science, Durham, NC, USA, November 2014, *Introduction to Game Theory... by Playing Games!* (Two distinct sessions, each repeated twice, where first participants played a game and then I used this game to introduce basic concepts in game theory.)
- Columbia University, New York, NY, USA, September 2014, *Computing Game-Theoretic Solutions for Security*.

- The 12th Meeting of the Society for Social Choice and Welfare, Boston, MA, USA, June 2014, *Computational Social Choice: A Journey from Basic Complexity Results to a Brave New World for Social Choice*. (2014 Social Choice and Welfare Prize talk.)
- Illinois Institute of Technology, Chicago, IL, USA, April 2014, *Tearing Down the Wall between Mechanism Design with and without Money*.
- Western Illinois University, Macomb, IL, USA, April 2014, *Voting in Highly Anonymous Environments*. New joint research seminar series of the School of Computer Sciences, Department of Mathematics, and Department of Economics and Decision Sciences.
- Association of Private College and University Alumni Directors (PCUAD) Winter Conference January 2014 (held at Duke), *Crowdfunding, Public Goods, and Optimization*. Session on “Big Data: The Impact on Customer Management and Incentives.”
- University of Southern California CS Colloquium, Los Angeles, CA, USA, November 2013, *Tearing Down the Wall between Mechanism Design with and without Money*.
- Carnegie Mellon University Intelligence Seminar, September 2013, *Tearing Down the Wall between Mechanism Design with and without Money*.
- Université de Montréal, département de sciences économiques, séminaire de microéconomie, April 2013, *Mechanism Design in Highly Anonymous Environments*.
- Duke Provost’s Lecture Series, March 2013, *Algorithmic Economics: How Computer Science Lets Us Put Economic Theory to Work*.
- University of Illinois at Urbana-Champaign (UIUC) Cognitive & Algorithmic Decision Making Seminar (Coordinated Science Laboratory/Electrical & Computer Engineering), January 2013, *Computing Game-Theoretic Solutions for Security*.
- Nanyang Technological University Division of Mathematical Sciences Colloquium, Singapore, October 2012, *Computing Game-Theoretic Solutions for Security*.
- DSO National Laboratories, Singapore, October 2012, *Computing Stackelberg Strategies*.
- Workshop on Protecting Critical Infrastructure using Strategic Planning and Analytics, Singapore, October 2012, *Computing Stackelberg Strategies*.
- University of Waterloo Computational Mathematics Colloquium, September 2012, *Computing Game-Theoretic Solutions for Security*.
- Thomas Langford lecture at Duke, September 2012, *Algorithmic Game Theory: The Automation of Strategic Play and Recent Applications to the Security of Airports and Harbors*.
- The 26th National Conference on Artificial Intelligence (AAAI), July 2012, *Computing Game-Theoretic Solutions and Applications to Security*. (A “What’s hot” talk in the Subarea Spotlight Track.)
- The 10th Conference on Logic and the Foundations of Game and Decision Theory (LOFT-12), plenary speaker, University of Sevilla, Spain, June 2012, *Some unconventional game-theoretic solution concepts*.
- UCLA Center for Engineering Economics, Learning, and Networks Seminar Series, May 2012, *Computing Game-Theoretic Solutions for Security*.
- Union College CS Seminar Series (co-sponsored with the Math department), May 2012, *Voting on the Internet and in Other Highly Anonymous Environments*.

- Vanderbilt University EECS Distinguished Lecture Series, March 2012, *Computing Game-Theoretic Solutions for Security*.
- The University of Otago, Dunedin, New Zealand, December 2011, *Computing Game-Theoretic Solutions for Security*.
- Carnegie Mellon University Intelligence Seminar, November 2011, *Computing Game-Theoretic Solutions for Security*.
- Stanford University RAIN (Research on Algorithms and Incentives in Networks) seminar, October 2011, *Computing Game-Theoretic Solutions for Security*.
- New York Computer Science and Economics Day (NYCE-11), New York, NY, USA, September 2011, *Computing Game-Theoretic Solutions for Security*.
- The 22nd International Joint Conference on Artificial Intelligence (IJCAI-11), Barcelona, Catalonia, Spain, July 2011, *One Equilibrium Is Not Enough: Computing Game-Theoretic Solutions to Act Strategically*. Plenary talk for Computers and Thought Award.
- The 11th Bar-Ilan Symposium on Foundations of Artificial Intelligence (BISFAI-11), Ramat Gan, Israel, June 2011, *Computational Methods for Acting Strategically*. Keynote speaker.
- Ben-Gurion University of the Negev, Beersheba, Israel, May 2011, *Mechanism Design in Highly Anonymous Environments*.
- Tilburg University Social Choice Colloquium, Tilburg, The Netherlands, March 2011, *Solving Complete-Information Voting Games by Backward Induction*.
- CWI Algorithmic Game Theory Seminar, Amsterdam, The Netherlands, March 2011, *Computational Methods for Acting Strategically*.
- University of Amsterdam ILLC Computational Social Choice Seminar, Amsterdam, The Netherlands, February 2011, *Solving Complete-Information Voting Games by Backward Induction*.
- University of Southern California, Los Angeles, CA, USA, October 2010, *Computational Methods for Acting Strategically*.
- Advancing Robotics Technology for Societal Impact (ARTSI) faculty workshop, Durham, NC, USA, June 2010, *Attracting Students to Computer Science Using Artificial Intelligence, Economics, and Linear Programming*.
- University of Southern California, Los Angeles, CA, USA, September 2009, *Game theory and security: Computing, approximating, and learning optimal mixed strategies to commit to*.
- Yahoo! Research, New York, NY, USA, August 2009, *Game Theory and Security: Finding the Optimal Mixed Strategy to Commit To*.
- Hebrew University, Jerusalem, Israel, May 2009, *Finding the Optimal Mixed Strategy to Commit To*.
- Kyushu University, Japan, December 2008, *Mechanism Design in Highly Anonymous Environments*.
- University of Southern California, Los Angeles, CA, USA, November 2008, *Progress in Optimal Redistribution Auctions*.
- University of Southern California CS Colloquium, Los Angeles, CA, USA, November 2008, *Computing Game-Theoretic Solutions*.

- California Institute of Technology, Pasadena, CA, USA, November 2008, *Computing Game-Theoretic Solutions*.
- Technische Universiteit Delft Agent Colloquium, Delft, the Netherlands, October 2008, *Game-Theoretic Solutions and How to Compute Them*.
- Rice University Microeconomic Theory Workshops, Houston, TX, USA, October 2008, *Mechanism Design in Highly Anonymous Environments*.
- Duke Computer Science Graduate Student Retreat invited talk, Beaufort, NC, USA, September 2008, *Game-Theoretic Solutions and How to Compute Them*.
- AAAI Spring Symposium on Using AI to Motivate Greater Participation in Computer Science, Stanford University, CA, USA, March 2008, *Attracting Students to Computer Science Using Artificial Intelligence, Economics, and Linear Programming*.
- California Institute of Technology CS/Economics Seminar, Pasadena, CA, USA, March 2008, *Computer Science and Mechanism Design: Two Case Studies*.
- Bar-Ilan University, Ramat Gan, Israel, March 2008, *Computer Science and Mechanism Design: Two Case Studies*.
- Hebrew University, Jerusalem, Israel, March 2008, *Computer Science and Mechanism Design: Two Case Studies*.
- Center for Mathematics and Computer Science (CWI), Amsterdam, the Netherlands, October 2007, *Worst-Case Optimal Redistribution of VCG Payments in Multi-Unit Auctions*.
- Microsoft Research, Redmond, WA, USA, July 2007, *Disincentivizing the Use of False Names in Preference Aggregation*.
- Victor Lesser Distinguished Dissertation Award invited talk, AAMAS 2007, *Automated Mechanism Design*.
- Kyushu University, Japan, March 2007, *Improved VCG Redistribution Mechanisms*.
- Wayne State University Graduate Seminar, Detroit, MI, USA, February 2007, *Techniques for Computing Game-Theoretic Solutions*.
- Technical University of Munich, Germany, December 2006, *Voting Rules That Are Hard to Manipulate*.

Nonarchival conference/workshop presentations

- AI-SDM Workshop: Human-AI Complementarity for Decision Making, Pittsburgh, PA, USA, September 2025, facilitator for *Societal and ethical alignment of AI in decision making* (with Atoosa Kasirzadeh).
- IDEAL Annual Meeting, Chicago, IL, USA / online, June 2025, Industry Panel on “Opportunities for Engagement between IDEAL and Industry” (with Alekh Agarwal, Eunye Koh, Ravi Kumar, Samir Khuller, Arunava Majumdar; moderated by Elena Zheleva and Aravindan Vijayaraghavan).
- Agency, Intentions, and Artificial Intelligence workshop, Chapman University, Orange, CA, USA, June 2025. (Responder for sessions on defining intentions and on alignment, and chair of one on mechanisms and representation of intention.)
- CMU Values in the Loop: AI Alignment & Human Feedback workshop, Pittsburgh, PA, USA, May 2025, *Social Choice for AI Alignment*.

- AAMAS Joint Event For Faculty Retirement (JEFFR) workshop, Detroit, MI, USA, May 2025, *From Jeff to Computational Social Choice to Cooperative AI*.
- 17th Meeting of the Society for Social Choice and Welfare (SSCW-24), Paris, France, July 2024, *Social choice for AI ethics and safety*.
- 14th Oxford Workshop on Global Priorities Research, Oxford, UK, June 2024, *Social Choice for AI Alignment*.
- International Symposium on AI and Mathematics (ISAIM), Special Session on Alternative Models for Fairness in AI, Fort Lauderdale, FL, USA, January 2024, *Social Choice for AI Alignment*.
- Social Choice for AI Ethics and Safety workshop, Berkeley, CA, USA, December 2023, *Manipulation in social choice*.
- The Many Facets of Preference-based Learning (workshop at ICML), July 2023, *Eliciting Human Judgments for Moral Artificial Intelligence*. (Invited talk.)
- 12th Oxford Workshop on Global Priorities Research, June 2023, *Foundations of Cooperative AI*.
- Halpernfest (in honor of Joe Halpern's 70th birthday), June 2023, *Foundations of Cooperative AI*.
- CredibleCommitments.wtf, May 2023, *Fireside Chat: Anthropic Alignment*. (Together with Georgios Piliouras.)
- CRA Research Integrity Working Group, March 2023, *Recognizing and Addressing Attempts to Manipulate Peer Review in Conferences*.
- 11th Oxford Workshop on Global Priorities Research, Oxford, UK, December 2022, *AI Agents May Cooperate Better If They Don't Resemble Us*.
- CIREQ Montreal Microeconomic Theory Conference (Economics and Computation), Montréal, Québec, Canada, October 2022, *Automated Mechanism Design for Strategic Classification*.
- NSF Convergence Accelerator Workshop: Provably Safe and Beneficial Artificial Intelligence (PSBAI), Berkeley, CA, USA, October 2022, *Compositional AI Safety and Game Theory*. (Lightning talk.)
- Workshop in honor of Prof. Noam Nisan's 60th Birthday, the Hebrew University of Jerusalem, Israel, June 2022, *AI Agents May Cooperate Better If They Don't Resemble Us*.
- Agents behaviour in combinatorial game theory workshop, November 2021, *Automated Mechanism Design for Strategic Classification*.
- Center for Human-Compatible AI (CHAI) workshop, June 2021, invited brief talk, *Computational Social Choice for Determining Objectives for AI*.
- Market Algorithms workshop (part of the Google Algorithms Workshop Series on Markets, Mobility, and the Mind), May 2021, *First-price and second-price pacing equilibria in auction markets*.
- Simons Institute workshop on Games and Equilibria in System Design and Analysis, March 2021, *Designing Agents' Preferences, Beliefs, and Identities*.
- The Self-Organizing Conference on Machine Learning (SOCML-20), online, November 2020, *AI, Social Impact and Public Goods* (panel with Gillian Hadfield, Eliane Ubalijoro, Joanna Bryson, and Virginia Dignum; moderated by Yoshua Bengio).
- PETAL conference on AI and Ethics, online, November 2020, *AI and Justice* (moderator; panelists were Timnit Gebru, Rui Guo, Hoda Heidari, and Crystal Yang).

- Duke Law and Technology for Judges program, online, September 2020, *Ethical and Policy Issues with Artificial Intelligence* (with Jana Schaich Borg and Walter Sinnott-Armstrong).
- Simons Institute: Algorithm Design, Law, and Policy, online, July 2020. Moderated and presented on behalf of a breakout session on Transparency and Accountability of COVID-19 Tracing Apps.
- CEPS-Brookings Transatlantic AI Dialogue, multiple meetings; spoke about research collaboration in AI, including issues such as regulation, society, ethics.
- Marketplace Algorithms and Design seminar, online, 2020, served as a guest (with Nicole Immorlica and Ramesh Johari) for Preston McAfee's "Ask Me Anything."
- Decision Theory & the Future of Artificial Intelligence, Australian National University, Canberra, Australia, August 2019, invited, *Designing Belief Formation and Decision Theories for AI*.
- Morality and Machine Intelligence, Australian National University, Canberra, Australia, August 2019, invited, *How Artificial Intelligence Can Improve Human Moral Judgments*.
- TTIC Workshop on Automated Algorithm Design, Chicago, IL, USA, August 2019, invited, *New Directions in Automated Mechanism Design*.
- Duke Law and Technology for Judges program, Durham, NC, USA, May 2019, *Ethical and Policy Issues with Artificial Intelligence* (with Jana Schaich Borg and Walter Sinnott-Armstrong).
- Center for Human-Compatible AI (CHAI) 2019 workshop, Asilomar Conference Grounds, CA, May 2019, invited brief talk, *Designing Agents' Preferences, Beliefs, and Identities*.
- Kenan Institute for Ethics Conference on Ethics of Emerging Tech, Durham, NC, USA, April 2019, panelist on "Ethical Considerations in Tech Regulation and Governance" (with Chaz Arnett, Nita Farahany, Kristin Johnson, and Corinna Lain; moderated by Mark Borsuk and Jonathan Wiener).
- Duke Health Leadership Summit (Feagin Leadership Program), Durham, NC, USA, March 2019, *Are We Losing Our Humanity? Health Care Leadership in the Age of Artificial Intelligence*. (Together with Jana Schaich Borg and Walter Sinnott-Armstrong.)
- Computing Community Consortium (CCC) Artificial Intelligence Roadmap Workshop 1 - Integrated Intelligence, Chicago, IL, USA, November 2018. Presented on behalf of a breakout session on Open Repositories of World Knowledge.
- The AAMAS-IJCAI Workshop on Agents and Incentives in Artificial Intelligence (AI³), Stockholm, Sweden, July 2018, *Multiplicative Pacing Equilibria in Auction Markets*.
- Seventh International Workshop on Computational Social Choice (COMSOC), Troy, NY, USA, June 2018, *Computational Social Choice for Automated Moral Decision Making*. (Rump session talk.)
- AI in the Administrative State, Duke Law School, Durham, NC, USA, May 2018, panelist (discussant) in the overview panel (with Ed Felten and Cary Coglianese; moderated by Lori Benneer).
- Second Annual Philosophy, Politics, and Economics (PPE) Society Meeting, New Orleans, LA, USA, March 2018, *Moral Artificial Intelligence and the Societal Tradeoffs Problem*.
- International Symposium on AI and Mathematics (ISAIM), Fort Lauderdale, FL, USA, January 2018, *Moral Decision Making Frameworks for Artificial Intelligence*.
- Social Responsibility of Algorithms (SRA), Paris, France, December 2017, *Moral Decision Making Frameworks for Artificial Intelligence*.

- Workshop on Trustworthy Algorithmic Decision-Making, Arlington, VA, USA, December 2017, *Using Human Subjects' Judgments for Automated Moral Decision Making* (presented as part of a panel).
- Workshop on Morality and Social Trust in Autonomous Robots (at RSS-17), Cambridge, MA, USA, July 2017, *Moral Decision Making Frameworks for Artificial Intelligence*.
- First Workshop on Mechanism Design for Social Good, Cambridge, MA, USA, June 2017, *Rules for Choosing Societal Tradeoffs*.
- Simons Institute Economics and Computation Reunion Workshop, Berkeley, CA, USA, April 2017, *Moral Artificial Intelligence and the Societal Tradeoffs Problem*.
- ISAT/DARPA Communicating Intent for Autonomy (CIA) Workshop, Arlington, VA, USA, March 2017, *How to measure specification effort / expressiveness?*
- Incorporating Ethics into Artificial Intelligence, Washington, DC, USA, March 2017, introductory speaker / dialogue starter for the session on AI and the economy (and participant in other sessions).
- Third International Workshop on AI, Ethics and Society (at AAAI'17), February 2017, *Moral Decision Making Frameworks for Artificial Intelligence*.
- AI Safety workshop at the Asilomar Beneficial AI conference, January 2017, *Moral Decision Making Frameworks for Artificial Intelligence*.
- Hastings Center, Control & Responsible Innovation in the Development of Autonomous Machines Meeting II, Garrison, NY, USA, October 2016, *Moral Decision Making Frameworks for Artificial Intelligence*.
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, September 2016, *Bayesian Security Games*.
- Fifth World Congress of the Game Theory Society (GAMES), Maastricht, the Netherlands, July 2016, *Computing Possible and Necessary Equilibrium Actions (and Bipartisan Set Winners)*.
- EXPLORE-2016: The Third Workshop on Exploring Beyond the Worst Case in Computational Social Choice (at AAMAS 2016), Singapore, May 2016, *Fair Social Choice in Dynamic Settings*.
- Workshop on AI, Ethics and Society (at AAAI 2016), Phoenix, AZ, USA, February 2016, *How to Build Ethics into Robust Artificial Intelligence*. (Session on FLI grants; also participated on panel "What are the most promising research directions for keeping AI beneficial?")
- International Symposium on AI and Mathematics (ISAIM), Fort Lauderdale, FL, USA, January 2016, *Rules for Choosing Societal Tradeoffs*.
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, September 2015, *Structuring Payoffs, Information Flow over Networks, and Timing of Games*.
- Computing Community Consortium (CCC) Theoretical Foundations of Social Computing Workshop, Washington, DC, July 2015. Led a breakout session on key problems for the community and presented the results.
- ACM Conference on Economics and Computation, Portland, OR, USA, June 2015, *Timeability of Extensive-Form Games*. (Poster presentation.)
- Amazon Faculty Invitational Research Symposium, Seattle, WA, USA, May 2015, *Classification in Game-Theoretic Environments*. (Poster presentation.)

- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, December 2014, *Game-Theoretic Approaches to Stable Coalitions and Structuring the Game*.
- Eleventh Conference on Logic and the Foundations of Game and Decision Theory (LOFT-14), Bergen, Norway, 2014, *Should Stackelberg Mixed Strategies Be Considered a Separate Solution Concept?*
- Fifth International Workshop on Computational Social Choice, Pittsburgh, PA, USA, June 2014, *The Maximum Likelihood Approach to Voting on Social Networks*.
- The 12th Meeting of the Society for Social Choice and Welfare, Boston, MA, USA, June 2014, *The Maximum Likelihood Approach to Voting on Social Networks*.
- First Workshop on Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE), Paris, France, May 2014, *The Maximum Likelihood Approach to Voting on Social Networks*.
- International Symposium on AI and Mathematics (ISAIM) 2014, *The Maximum Likelihood Approach to Voting on Social Networks*. (Special session on computational social choice, invited.)
- WINE 2013 poster session, December 2013, *The Maximum Likelihood Approach to Voting on Social Networks*. (Poster presentation.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, November 2013, *New Directions in Algorithmic Game Theory: Scalability, Time, Evolution, Stochasticity, and Coalitions*.
- Fifteenth Chinese-American Kavli Frontiers of Science Symposium, Irvine, CA, USA, October 2012, *Computing Game-Theoretic Solutions and Applications to Security*. (Poster presentation. Symposium co-sponsored by the Chinese Academy of Sciences and the U.S. National Academy of Sciences.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Marina del Rey, CA, USA, September 2012, *Solving for Stackelberg Strategies in Stochastic Games*.
- EC 2012 poster session, June 2012, *Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases*. (Poster presentation.)
- Dagstuhl Seminar on Computation and Incentives in Social Choice, March 2012, *Evaluating Resistance to False-Name Manipulations in Elections*.
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Kickoff Meeting, Marina del Rey, CA, USA, September 2011, *Stochastic Coalitional Game Theory*.
- Workshop on Innovations in Algorithmic Game Theory (iAGT), May 2011, *Algorithms for Security Games*. (During short-term visit to the Hebrew University Institute of Advanced Studies special semester on Algorithmic Game Theory.)
- V CORE/Maastricht Workshop on Advances in Collective Choice, Maastricht, The Netherlands, April 2011, *Solving Complete-Information Voting Games by Backward Induction*.
- DIMACS/CCICADA Workshop on Adversarial Decision Making, October 2010, *Finding Optimal Mixed Strategies to Commit to in Security Games*.
- Guanajuato Workshop on Prior-free Mechanism Design, May 2010, *New Results and Approaches in the Design of False-Name-Proof Mechanisms*.

- Dagstuhl Seminar on Computational Foundations of Social Choice, March 2010, *False-name-proof voting in social networks*.
- Bellairs Workshop on Algorithmic Game Theory, March 2009, *Finding the Optimal Mixed Strategy to Commit To*.
- INFORMS Annual Meeting 2008, *New Directions in False-name-proofness*. (Session on auctions and mechanism design, invited.)
- INFORMS Annual Meeting 2008, *Progress in Optimal Redistribution Auctions*. (Session on extending auction theory: computational perspectives, invited.)
- INFORMS Annual Meeting 2008, *Progress on Using Computational Hardness to Prevent Manipulation*.
- World Congress of the Game Theory Society (GAMES) 2008, *Mechanism Design in Highly Anonymous Environments*.
- International Symposium on AI and Mathematics (ISAIM) 2008, *Comparing Multiagent Systems Research in Combinatorial Auctions and Voting*. (Special session on computation and social choice, invited.)
- Dagstuhl Seminar on Computational Issues in Social Choice 2007, *Disincentivizing the Use of False Names in Social Choice*.
- INFORMS Annual Meeting 2007, *Optimal VCG Redistribution Mechanisms*. (Session on computer science and auctions, invited.)
- INFORMS Annual Meeting 2007, *False-name-proofness in social choice, and limited verification of identities*.
- Dagstuhl Seminar on Computational Social Systems and the Internet 2007, *False-name-proofness in social choice, and limited verification of identities*.
- NCSU Mini-Workshop on Selected Topics in E-Commerce 2007, *Limited Verification of Identities to Induce False-Name-Proofness*.
- Warwick Workshop on Algorithmic Game Theory 2007, *Preprocessing Techniques for Computing Nash Equilibria*.
- INFORMS Annual Meeting 2006, *Common Voting Rules as Maximum Likelihood Estimators*. (Session on information aggregation, invited.)
- Stony Brook Game Theory Festival 2006, *Nonexistence of Voting Rules That Are Usually Hard to Manipulate*.
- Stony Brook Game Theory Festival 2005, *A Generalized Strategy Eliminability Criterion and Computational Methods for Applying It*.
- Stony Brook Workshop on Game Theory and Computer Science 2005, *Complexity of (Iterated) Dominance and Other Solution Concepts*.
- Dagstuhl Seminar on Computing and Markets 2005, *Expressive Negotiation over Donations to Charities*.
- DIMACS Workshop on Computational Issues in Auction Design 2004, *Elicitation in Combinatorial Auctions with Restricted Preferences and Bounded Interdependency Between Items*.
- New and Alternative Directions for Learning Workshop 2004, *Communication Complexity as a Lower Bound for Learning in Games*.

- Conference on Logic and the Foundations of Game and Decision Theory 2004, *Computational Criticisms of the Revelation Principle*.

... by my students and postdocs: *(not kept up to date)*

- 15th Workshop on the Economics of Networks, Systems and Computation (NetEcon) live highlighted-paper session in 2020, online, July 2020, *Learning Opinions in Social Networks*. (Given by Hanrui Zhang.)
- Incentives in Machine Learning workshop (at ICML 2020), online, July 2020, *Classification with Strategically Withheld Data*. (Given by Anilesh Krishnaswamy.)
- Incentives in Machine Learning workshop (at ICML 2020), online, July 2020, *Classification with Few Tests through Self-Selection*. (Given by Yu Cheng.)
- Incentives in Machine Learning workshop (at ICML 2020), online, July 2020, *Incentive-Aware PAC Learning*. (Given by Hanrui Zhang.)
- Conference on Mechanism and Institution Design 2020 (CMID20), online, June 2020, *Eliciting information for decision making*. (Given by Caspar Oesterheld.)
- Seventh International Workshop on Computational Social Choice (COMSOC), Troy, NY, USA, June 2018, *Dynamic proportional sharing: A game-theoretic approach*. (Given by Rupert Freeman.)
- Seventh International Workshop on Computational Social Choice (COMSOC), Troy, NY, USA, June 2018, *A Better Algorithm for Societal Tradeoffs*. (Poster given by Hanrui Zhang.)
- Third Workshop on Algorithmic Game Theory and Data Science, Cambridge, MA, USA, June 2017, *Automated Design of Robust Mechanisms*. (Given by Michael Albert.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, September 2016, *Public Good Game with Parallel Bilateral Relationships*. (Given by Catherine Moon.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, September 2016, *Disarmament Games*. (Given by Yuan (Eric) Deng.)
- Fifth World Congress of the Game Theory Society (GAMES), Maastricht, the Netherlands, July 2016, *Maximal Cooperation in Repeated Games on Social Networks*. (Given by Catherine Moon.)
- Second Workshop on Algorithmic Game Theory and Data Science (at EC'16), Maastricht, the Netherlands, July 2016, *Maximizing Revenue with Limited Correlation: The Cost of Ex-Post Incentive Compatibility*. (Given by Michael Albert.)
- Second Workshop on Algorithmic Game Theory and Data Science (at EC'16), Maastricht, the Netherlands, July 2016, *Robust Automated Mechanism Design*. (Given by Michael Albert.)
- International Symposium on AI and Mathematics (ISAIM), Fort Lauderdale, FL, USA, January 2016, *False-Name-Proof Recommendations in Social Networks*. (Given by Rupert Freeman.)
- International Symposium on AI and Mathematics (ISAIM), Fort Lauderdale, FL, USA, January 2016, *Computing Possible and Necessary Equilibrium Actions (and Bipartisan Set Winners)*. (Given by Rupert Freeman.)
- INFORMS Annual Meeting 2015, Philadelphia, PA, USA, November 2015, *Role Assignment for Game-Theoretic Cooperation*. (Given by Catherine Moon.)

- INFORMS Annual Meeting 2015, Philadelphia, PA, USA, November 2015, *Maximal Cooperation in Repeated Games on Social Networks*. (Given by Catherine Moon.)
- Joint ADT/LPNMR Doctoral Consortium, Lexington, KY, USA, September 2015, *Axiomatic and Computational Social Choice*. (Given by Rupert Freeman.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, September 2015, *Role Assignment for Game-Theoretic Cooperation in Repeated Games*. (Given by Catherine Moon.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, September 2015, *Classification of Strategic Agents*. (Given by Andrew Kephart.)
- The first IJCAI workshop on Algorithmic Game Theory (AGT@IJCAI-2015), Buenos Aires, Argentina, July 2015, *Role Assignment for Game-Theoretic Cooperation*. (Given by Catherine Moon.)
- Dagstuhl Seminar on Computational Social Choice: Theory and Applications, June 2015, *Computing the Optimal Game*. (Given by Markus Brill.)
- IRT workshop at the University of Toulouse, France, May 2015, *Crowdsourcing Societal Trade-offs*. (Given by Markus Brill.)
- Second Workshop on Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE), Istanbul, Turkey, May 2015, *Computing the Optimal Game*. (Given by Markus Brill.)
- AAAI Spring Symposium on Applied Computational Game Theory, Palo Alto, CA, USA, March 2015, *Repeated Advice with a Strategic Adversary*. (Given by Aaron Kolb.)
- AAAI Spring Symposium on Applied Computational Game Theory, Palo Alto, CA, USA, March 2015, *What Goes Around Comes Around: Maximal Cooperation in Repeated Games on Social Networks*. (Given by Catherine Moon.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, December 2014, *Coalitional Game Theory with Uncertainty*. (Given by Yuqian Li.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, December 2014, *Computing Maximal Cooperative Coalitions in Games on Social Networks*. (Given by Catherine Moon.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, December 2014, *Computing the Optimal Game to Play*. (Poster given by Rupert Freeman.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, December 2014, *Classification in Game-Theoretic Environments*. (Poster given by Andrew Kephart.)
- Dagstuhl Seminar on Formal Methods for Coordinating Multi-Agent Systems, August 2014, *Strategic Voting and Strategic Candidacy*. (Given by Markus Brill.)
- Fifth International Workshop on Computational Social Choice, Pittsburgh, PA, USA, June 2014, *Strategic Voting and Strategic Candidacy*. (Given by Markus Brill.)
- Fifth International Workshop on Computational Social Choice, Pittsburgh, PA, USA, June 2014, *On the Axiomatic Characterization of Runoff Voting Rules*. (Given by Rupert Freeman.)

- The 12th Meeting of the Society for Social Choice and Welfare, Boston, MA, USA, June 2014, *Strategic Voting and Strategic Candidacy*. (Given by Markus Brill.)
- The 12th Meeting of the Society for Social Choice and Welfare, Boston, MA, USA, June 2014, *On the Axiomatic Characterization of Runoff Voting Rules*. (Given by Rupert Freeman.)
- EC 2014 poster session, June 2014, *Mechanism Design for Scheduling with Uncertain Execution Time*. (Poster given by Angelina Vidali.)
- EC 2014 poster session, June 2014, *On the Axiomatic Characterization of Runoff Voting Rules*. (Poster given by Rupert Freeman.)
- EC 2014 poster session, June 2014, *Complexity of Stability-based Solution Concepts in Multi-issue and MC-net Cooperative Games*. (Poster given by Yuqian Li.)
- EC 2014 poster session, June 2014, *Beat the Cheater: Computing Game-Theoretic Strategies for When to Kick a Gambler out of a Casino*. (Poster given by Joshua Letchford.)
- The Fifth Workshop on Cooperative Games in Multiagent Systems, Paris, France, May 2014, *Cooperative Game Solution Concepts that Maximize Stability under Noise*. (Given by Yuqian Li.)
- AAAI Spring Symposium on Applied Computational Game Theory, March 2014, *Beat the Cheater: Computing Game-Theoretic Strategies for When to Kick a Gambler out of a Casino*. (Given by Melissa Dalis.)
- AAAI Spring Symposium on Applied Computational Game Theory, March 2014, *Complexity of Stability-based Solution Concepts in Multi-issue and MC-net Cooperative Games*. (Given by Yuqian Li.)
- WINE 2013 poster session, December 2013, *Optimal Internet Auctions with Costly Communication*. (Poster given by Yuqian Li.)
- WINE 2013 poster session, December 2013, *Fast Equilibrium Computation for Infinitely Repeated Games*. (Poster given by Garrett Andersen.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, November 2013, *Game-Theoretic Question Selection for Tests*. (Poster given by Yuqian Li.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Annual Meeting, Los Angeles, CA, USA, November 2013, *Beat the Cheater: Computing Game-Theoretic Strategies for When to Kick a Gambler out of a Casino*. (Poster given by Troels Bjerre Sørensen.)
- New York Computer Science and Economics Day (NYCE), New York, NY, USA, November 2013, *Fast Equilibrium Computation for Infinitely Repeated Games*. (Poster given by Garrett Andersen.)
- New York Computer Science and Economics Day (NYCE), New York, NY, USA, November 2013, *Mechanism Design for Scheduling with Uncertain Execution Time*. (Poster given by Angelina Vidali.)
- New York Computer Science and Economics Day (NYCE), New York, NY, USA, November 2013, *Beat the Cheater: Computing Game-Theoretic Strategies for When to Kick a Gambler out of a Casino*. (Poster given by Troels Bjerre Sørensen.)
- INFORMS Annual Meeting 2013, *Mechanism Design for Scheduling with Uncertain Execution Time*. (Session on Allocation and Market Design, invited; given by Angelina Vidali.)

- INFORMS Annual Meeting 2013, *Optimal Internet Auctions with Costly Communication*. (Session on Mechanism Design, invited; given by Yuqian Li.)
- Applied Mathematics, Modeling and Computational Science (AMMCS) conference, Waterloo, Ontario, Canada, August 2013, *Fast Equilibrium Computation for Infinitely Repeated Games*. (Given by Garrett Andersen.)
- Stony Brook Workshop on Computational Game Theory, July 2013, *Mechanism Design for Scheduling with Uncertain Execution Time*. (Given by Angelina Vidali.)
- EC 2013 poster session, June 2013, *Game-Theoretic Question Selection for Tests*. (Poster presentation, given by Yuqian Li.)
- INFORMS Annual Meeting 2012, *Computing Optimal Strategies to Commit to in Stochastic Games*. (Session on Security Games and Optimization on Networks, invited; given by Joshua Letchford.)
- Bellairs Workshop on Algorithmic Game Theory, Barbados, April 2012, *Computational Aspects of Security Games*. (Given by Dmytro Korzhuk.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Kickoff Meeting, Marina del Rey, CA, USA, September 2011, *Security Games with Multiple Attacker Resources*. (Poster given by Dmytro Korzhuk.)
- Scalable, Stochastic and Spatio-temporal Game Theory for Real-World Human Adversarial Behavior MURI Kickoff Meeting, Marina del Rey, CA, USA, September 2011, *Solving Stackelberg Games with Uncertain Observability*. (Poster given by Dmytro Korzhuk.)
- Workshop on Information in Networks (WIN), September 2010, *False-Name-Proofness in Social Networks*. (Poster given by Joshua Letchford.)
- INFORMS Annual Meeting 2010, *Strategic Sequential Voting in Multi-Issue Domains and Multiple-Election Paradoxes*. (Session on Hard Problems in Learning and Decision-Making, invited; given by Lirong Xia.)
- INFORMS Annual Meeting 2010, *Stackelberg Voting Games: Computational Aspects and Paradoxes*. (Given by Lirong Xia.)
- Workshop on Solution Concepts for Extensive Games (Aarhus, Denmark), June 2010, *Computing Optimal Strategies to Commit to in Extensive-Form Games*. (Given by Joshua Letchford.)
- Guanajuato Workshop on Prior-free Mechanism Design, May 2010, *Undominated Groves Mechanisms*. (Given by Mingyu Guo.)
- Bertinoro Workshop on Frontiers in Mechanism Design, March 2010, *Unique game-theoretic outcomes for strategic voting with complete information*. (Given by Lirong Xia.)
- Dagstuhl Seminar on Computational Foundations of Social Choice, March 2010, *Unique game-theoretic outcomes for strategic voting with complete information*. (Given by Lirong Xia.)
- The second FTC-Northwestern University Microeconomics Conference 2009, *Who Benefits from Online Privacy?* (Given by Liad Wagman; acceptance rate about 10% for the session.)
- The 38th Research Conference on Communication, Information and Internet Policy (TPRC) 2009, *Who Benefits from Online Privacy?* (Given by Liad Wagman.)
- North American Summer Meeting of the Econometric Society 2009, *Who Benefits from Online Privacy?* (Given by Liad Wagman.)

- The Seventh International Industrial Organization Conference 2009, *Optimal False-Name-Proof Voting Rules with Costly Voting*. (Given by Liad Wagman.)
- World Congress of the Game Theory Society (GAMES) 2008, *A sufficient condition for rules to be frequently manipulable for any number of alternatives*. (Given by Lirong Xia.)
- World Congress of the Game Theory Society (GAMES) 2008, *Optimal VCG Redistribution Mechanisms*. (Given by Mingyu Guo.)
- World Congress of the Game Theory Society (GAMES) 2008, *Strategic Betting for Competitive Agents*. (Poster, presented by Liad Wagman.)
- Dagstuhl Seminar on Computational Issues in Social Choice 2007, *Voting with Partial Orders*. (Given by Lirong Xia.)
- DIMACS Workshop on the Boundary between Economic Theory and Computer Science 2007, *Worst-Case Optimal Redistribution of VCG Payments in Multi-Unit Auctions*. (Given by Mingyu Guo.)
- Stony Brook Game Theory Festival 2007, *Improved VCG Redistribution Mechanisms*. (Given by Mingyu Guo.)
- Victor Rothschild Memorial Symposium on Economic Aspects of Communication and Information 2007, *Betting to Come Out on Top: Competitive Gambling with Applications to R&D and Patent Races*. (Given by Liad Wagman.)
- NCSU Mini-Workshop on Selected Topics in E-Commerce 2007, *Improved VCG Redistribution Mechanisms*. (Given by Mingyu Guo.)

Patents

- Items Ratio Based Price/Discount Adjustment in a Combinatorial Auction. Tuomas Sandholm, David Levine, David Parkes, Subhash Suri, Vincent Conitzer, Robert Shields, Yuri Smirnov. (CombineNet, Inc.) United States Patent 8,195,524, June 5, 2012.
- Overconstraint Detection, Rule Relaxation and Demand Reduction in Combinatorial Exchange. Tuomas Sandholm, David Levine, David Parkes, Subhash Suri, Vincent Conitzer, Robert Shields, Yuri Smirnov. (CombineNet, Inc.) United States Patent 8,190,490, May 29, 2012.
- Bid Modification Based on Logical Connections between Trigger Groups in a Combinatorial Exchange. Tuomas Sandholm, David Levine, David Parkes, Subhash Suri, Vincent Conitzer, Robert Shields, Yuri Smirnov. (CombineNet, Inc.) United States Patent 8,190,489, May 29, 2012.
- Method of Determining an Exchange Allocation That Promotes Truthful Bidding and Improves the Obtainment of Exchange Objectives. Vincent Conitzer and Tuomas Sandholm. (CombineNet, Inc.) United States Patent 8,060,433, November 15, 2011.
- Method and Apparatus for Conducting a Dynamic Exchange. Tuomas Sandholm, David Levine, David Parkes, Subhash Suri, Vincent Conitzer, Robert Shields, and Yuri Smirnov. (CombineNet, Inc.) United States Patent 7,577,589, August 18, 2009.
- Dynamic Exchange Method and Apparatus. Tuomas Sandholm, Richard James McKenzie Jr., David Levine, David Parkes, Subhash Suri, Vincent Conitzer, Robert Shields, Benjamin Schmaus, and Christopher Cole. (CombineNet, Inc.) United States Patent 7,499,880, March 3, 2009.

Research Grants and Gifts

Fellowships awarded directly to other researchers in my group are not listed here, but under “Advising.”

- Center for Emerging Risk Research (CERR), *Foundations of Cooperative AI Lab* at Carnegie Mellon. Estimated amount: \$3,000,000. Started January 2022, for 3-5 years.
- Cooperative AI Foundation (CAIF), for *Foundations of Cooperative AI Lab* at Carnegie Mellon. Estimated amount: \$500,000. Started January 2022, for 3-5 years.
- ARO Grant W911NF2110230: *Information Networks: RESUME: Artificial Intelligence, Algorithms, and Optimization for Responsible Reopening*. PI: Conitzer. Co-PI: Debmalaya Panigrahi. Estimated Total Award Amount: \$99,821. Start Date: June, 2021. Projected Duration: 1 year.
- Duke Bass Connections Award (under Brain & Society Theme (B&S)): *How to Build Ethics into Robust Artificial Intelligence*. PIs: Walter Sinnott-Armstrong, Conitzer, Jana Schaich Borg. Budget: \$25,000. 2019-2020.
- NSF Award IIS-1814056: *RI: Small: Designing Preferences, Beliefs, and Identities for Artificial Intelligence*. PI: Conitzer. Estimated Total Award Amount: \$400,000. Start Date: September 1, 2018. Projected Duration: 3 years.
- Templeton World Charity Foundation grant TWCF0321: *How to Use Artificial Intelligence to Enhance Human Moral Intelligence*. PI: Jana Schaich Borg, co-PIs: Walter Sinnott-Armstrong, Conitzer. Budget: \$205,147. Start Date: July 25, 2018. Projected Duration: 2 years.
- Subaward from USC on foundations for strategic mechanism design under their MURI (“Realizing Cyber Inception: Towards a Science of Personalized Deception for Cyber Defense”). Start: April 1, 2018. Projected duration: <1 year. Amount: \$50,417.00.
- Duke Bass Connections Award (under Brain & Society Theme (B&S), as well as the Information, Society & Culture Theme (ISC)): *Moral Artificial Intelligence*. PIs: Conitzer, Jana Schaich Borg. Contributor: Walter Sinnott-Armstrong. Budget: \$25,000. 2018-2019.
- Duke Bass Connections Award (under Information, Society & Culture Theme (ISC), as well as the Brain & Society Theme (B&S)): *How to Build Ethics into Robust Artificial Intelligence*. PI: Walter Sinnott-Armstrong. Co-PIs: Conitzer, Jana Schaich Borg. Budget: \$24,200. 2017-2018.
- Future of Life Institute (FLI) grant: *How to Build Ethics into Robust Artificial Intelligence*. PI: Conitzer. Co-PI: Walter Sinnott-Armstrong (Duke Philosophy). Estimated Total Award Amount: \$200,000. Start Date: September 1, 2015. Projected Duration: 3 years.
- NSF Award IIS-1527434: *RI: Small: New Directions in Computational Social Choice and Mechanism Design*. PI: Conitzer. Estimated Total Award Amount: \$499,972. Start Date: September 1, 2015. Projected Duration: 3 years.
- Guggenheim Fellowship. \$50,000.
- NSF Award CCF-1337215: *XPS: CLCCA: Allocating Heterogeneous Datacenter Hardware to Strategic Agents*. PI: Benjamin Lee. Co-PI: Conitzer. Estimated Total Award Amount: \$700,000. My portion: \$45,974. Start date: September 15, 2013. Projected Duration: 4 years.
- ARO PECASE Grant W911NF-12-1-0550: *Computing Game-Theoretic Solutions for Security in the Medium Term*. PI: Conitzer. Estimated Total Award Amount: \$1,000,000. Start Date: September 13, 2012. Projected Duration: 5 years.
- ARO MURI Grant W911NF-11-1-0332: *Scalable, Stochastic and Spatiotemporal Game Theory for Real-World Human Adversarial Behavior*. PIs/co-PIs: Milind Tambe (USC); Andrea Bertozzi (UCLA), P. Jeffrey Brantingham (UCLA), Vincent Conitzer, Maria Rita D’Orsogna (CSUN),

Richard John (USC), Rajiv Maheswaran (USC), Michael McBride (UCI), Yoav Shoham (Stanford), Martin Short (UCLA), Richard Dekmejian (USC). August 2011-July 2014 (with the possibility to continue through 2016). Award Amount: \$1,250,000 per year. My projected portion over 5 years: \$571,625 (subaward from USC).

- NSF Award CCF-1101659: *ICES: Small: Mechanism Design for Highly Anonymous Environments*. PI: Conitzer. Award Amount: \$400,000. Start Date: September 1, 2011. Projected Duration: 36 Months.
- NSF Award IIS-0953756: *CAREER: New Directions in Computing Game-Theoretic Solutions: Commitment and Related Topics*. PI: Conitzer. Estimated Total Award Amount: \$500,002. Start Date: March 1, 2010. Projected Duration: 60 months.
- ARO Grant W911NF-09-1-0459 (proposal 56698-CI): *Algorithms for Computing Stackelberg Strategies in Large, Realistic Games* (awarded September 2009). PI: Conitzer, co-PI: Parr. Estimated Total Award Amount: \$240,000.
- Yahoo! Faculty Research Grant for *Prediction Markets, Mechanism Design, and Computation* (awarded December 2008). \$25,000.
- NSF Award IIS-0812113: *(RI+hcc)-Small: Computational Social Choice: Aggregating Preferences in Combinatorial Domains*. PI: Conitzer. Estimated Total Award Amount: \$429,212. Start Date: September 1, 2008. Projected Duration: 48 Months. (\$8,000 REU supplement in 2009.)
- Alfred P. Sloan Research Fellowship. \$50,000.
- US-Israel Binational Science Foundation (BSF) 2006216, *Computational Aspects of Social Choice in Multiagent Systems*, with Jeff Rosenschein (Hebrew University), \$70,000 over 4 years (apart from travel, all the money went to the Israeli side).

... not for my own group:

- ARO *Workshop on Reasoning in Adversarial and Noncooperative Environments (2010)*. PI: Conitzer, co-PI: Parr. Amount: \$30,071.
- NSF Award IIS-1026617: *Doctoral Mentoring Consortium at the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2010)*. PI: Conitzer. Amount: \$45,000.

Advising

Postdoctoral researchers

- Ivan Geffner. Working on alternating commitments to get to better equilibria. 2024. Now an Assistant Professor at Utrecht University (Mathematics).
- Vojtech (Vojta) Kovarik. Working on game theory with players that can be simulated. 2022-present.
- Anilesh Krishnaswamy. Worked on classification with strategically withheld and transformed data. 2019–2021. Now at Google.
- Yu Cheng. Worked on algorithms for quantitative judgment aggregation, and machine learning in contexts where samples are selected by a strategic agent. 2017–2019. Now an Assistant Professor of Computer Science at Brown University.
- Michael Albert. Worked on automated design of robust mechanisms in settings with correlation. 2016–2018. Now an Assistant Professor at the University of Virginia (Darden School of Business and SEAS).

- Jana Schaich Borg. Worked on our moral AI project. 2016. Now a Research Assistant Professor at Duke (still working with us).
- Aaron Kolb. Briefly a postdoc in my group after graduating with his Ph.D. (more detail below). 2016.
- Mentored Simina Brânzei during the Simons Institute semester on Economics and Computation, 2015. (Now an Assistant Professor at Purdue CS.)
- Markus Brill. Worked on computational social choice and algorithmic game theory. Supported by a Feodor Lynen Research Fellowship of the Alexander von Humboldt Foundation. 2013—2015. Now an Assistant Professor and Emmy Noether Research Group Leader at TU Berlin.
- Troels Bjerre Lund (formerly Sørensen). Worked on algorithms for computing game-theoretic solutions. 2012—2013. Now an Assistant Professor in the Theoretical Computer Science section at the IT-University of Copenhagen (Denmark).
- Angelina Vidali. Worked on mechanism design for computational task allocation. 2012—2014. Now Academic Fellow at De Montfort University (UK).
- Taiki Todo. Worked on mechanism design in highly anonymous environments. Supported by the Japan Society for the Promotion of Science (JSPS). 2012—2013. Now an Assistant Professor at the Department of Informatics, and a faculty member of the Innovative Market Design Research Center, Kyushu University (Fukuoka, Japan).

Ph.D. students

- Ratip Emin Berker, CMU Computer Science Department Ph.D. student. Working on repeated games with restarts. Cooperative AI PhD Fellow, 2025. 2023-present.
- Jiayuan Liu, CMU Computer Science Department Ph.D. student. Working on ultimatum games and mechanism design, and on sequential program equilibrium. 2023-present.
- Emanuel Tewolde, CMU Computer Science Department Ph.D. student. Working on computing solutions for finitely repeated games, and games with imperfect recall. Cooperative AI PhD Fellow, 2025. 2022-present.
- Christiaan (Chris) van Merwijk, CMU Computer Science Department Ph.D. student. Working on tool-like AI. 2022-present.
- Caspar Oesterheld, CMU Computer Science Department Ph.D. student. Working on decision markets, program equilibrium, and other aspects of decision theory. Vitalik Buterin PhD Fellowship in AI Existential Safety. 2018—present.
- Hanrui Zhang, CMU Computer Science Department Ph.D. student. Worked on algorithms for societal tradeoffs, judgment aggregation from queries, and machine learning with strategically selected and transformed samples. Finalist for the 2021 Economics and Computation Facebook Fellowship. Dissertation title: *Designing and Analyzing Machine Learning Algorithms in the Presence of Strategic Behavior*. Defended May 1, 2023. Honorable mention for CMU School of Computer Science Distinguished Dissertation Award.
- Yuan (Eric) Deng, Duke Computer Science Ph.D. student. Worked on computational aspects of game theory, ethical solution concepts, and mechanism design. Google Ph.D. Fellowship, 2018. Dissertation title: *Dynamic Mechanism Design in Complex Environments*. Defended March 28, 2020. Received the department's Outstanding Ph.D. Dissertation Award in 2020, an honorable mention for the 2020 IFAAMAS Victor Lesser Distinguished Dissertation Award, and an honorable mention for the 2020 SIGecom Doctoral Dissertation Award. Now a Research Scientist at Google Research New York City.

- Rupert Freeman, Duke Computer Science Ph.D. student. Worked on (computational) social choice and fair decision making. Facebook Ph.D. Fellowship, 2017. Dissertation title: *Eliciting and Aggregating Information for Better Decision Making*. Defended May 14, 2018. Received the department's Outstanding Ph.D. Dissertation Award in 2018. After a postdoc at Microsoft Research NYC, now an Assistant Professor at the University of Virginia (Darden School of Business).
- Yuqian Li, Duke Computer Science Ph.D. student. Worked on Internet auction mechanism design when communication is costly, and algorithms for computing game-theoretic solutions (both cooperative and noncooperative). Dissertation title: *Game-Theoretically Allocating Resources to Catch Evaders and Payments to Stabilize Teams*. Defended March 21, 2016. Now at Pony.ai (previously at Google).
- Aaron Kolb, Duke Economics Ph.D. student. (Primary advisor: Attila Ambrus in Economics.) Worked on game-theoretic advice in the presence of a strategic attacker. Dissertation title: *Games of Private Information and Learning*. Defended April 5, 2016. (Also briefly a postdoctoral researcher in my group.) Now an Assistant Professor at Indiana University Kelley School of Business.
- Catherine Moon, Duke Economics Ph.D. student. Worked on various methods to induce cooperation in game theory. Dissertation title: *Essays on Identification and Promotion of Game-Theoretic Cooperation*. Defended March 20, 2018. Now an Economist at Keystone Strategy.
- Dmytro Korzhyk, Duke Computer Science Ph.D. Worked on computing optimal mixed strategies to commit to and Nash equilibria in settings with exponentially many pure strategies, especially security games. Dissertation title: *Security Games: Solution Concepts and Algorithms*. Defended April 18, 2013. Now at Google (previously at Two Sigma).
- Joshua Letchford, Duke Computer Science Ph.D. Worked on learning and computing optimal mixed strategies to commit to, security games, and ethical solution concepts in game theory. Dissertation title: *Computational Aspects of Stackelberg Games*. Defended April 2, 2013. Now at Sandia National Laboratories.
- Lirong Xia, Duke Computer Science Ph.D. Worked on computational aspects of social choice/voting, including manipulation/game-theoretic aspects and voting in combinatorial domains. Dissertation title: *Computational Voting Theory: Game-Theoretic and Combinatorial Aspects*. Defended July 29th, 2011. Received the department's Outstanding Ph.D. Dissertation Award in 2011. Continued as a Computing Innovation Fellow (CIFellow) at Harvard's Center for Research on Computation and Society (CRCS); now an Associate Professor in RPI's CS department.
- Mingyu Guo, Duke Computer Science Ph.D. Worked on computationally feasible automated mechanism design, redistribution of auction revenue to bidders, mechanism design without money, false-name-proof mechanisms, and prediction markets. Dissertation title: *Computationally Feasible Approaches to Automated Mechanism Design*. Defended July 30th, 2010. Received the department's Outstanding Ph.D. Dissertation Award in 2010 and one of two runners-up for the IFAAMAS-10 Victor Lesser Distinguished Dissertation Award. Continued as a Lecturer (in the UK sense) in the University of Liverpool's Economics and Computation Group; now a Lecturer (in the Australian sense) in the School of Computer Science, University of Adelaide.
- Liad Wagman, Duke Economics Ph.D. Worked on gambling to end up with more money than one's opponent, false-name-proofness when there is a cost to using additional identifiers, and paying for anonymity in the context of price discrimination. Co-advised with Curtis Taylor in Economics. Dissertation title: *Essays on Privacy, Information, and Anonymous Transactions*. Defended April 6th, 2009. Now a Professor at the Illinois Institute of Technology Stuart School of Business.

Long(ish)-term visitors to FOCAL

- Mingyu Guo, full calendar year 2024 (on sabbatical from the University of Adelaide).
- Emery Cooper, Fall 2023-present.
- Nathaniel Sauerberg, Spring 2023.
- Lewis Hammond, Fall 2022.

M.S. students

- Haoming Li, Duke Economics and Computation M.S., 2020. Worked on classification with strategically modified data. Now a Ph.D. student at USC CS.
- Andrew Kephart, Duke Computer Science M.S., 2017. Worked on mechanism design with costly signaling. Now Senior Data Scientist on the Marketplace DS team at Instacart.
- Garrett Andersen, Duke Computer Science M.S., 2015. Worked on algorithms for computing game-theoretic solutions, using cryptography in game theory, and preventing a single person from signing up for multiple accounts. Now ML Software Engineer at Google.
- Michael Albert, Duke Computer Science M.S., 2013. (Also a Ph.D. student at Fuqua, not advised by me there.) Worked on linear programming techniques in mechanism design. Title: *Approximately Optimal Mechanisms With Correlated Buyer Valuations*.
- Peter Franklin, Duke Computer Science M.S., 2010. Worked on linear programming techniques in mechanism design (with Pino Lopomo at Fuqua) and on applying game-theoretic techniques to wireless LANs. Title: *Internet-based Fault Recovery for Home Wireless LANs*. Received the department's Outstanding Master's Award in 2010. Now Chief Executive Officer at OnwardMobility.
- Joseph Farfel, Duke Computer Science M.S., 2007. Worked on preference elicitation in single-peaked domains, and hybridizing Turing Tests and prediction markets. Title: *Getting agents to agree on numbers: Aggregating value ranges and Turing Trade*. Now at Google.

Undergraduate students

- Alexander Heckett, CMU undergraduate student.
- Jerick Shi, CMU undergraduate student.
- David Rein, Duke undergraduate student. Working on classification with strategically withheld data. 2019—present.
- Nikhil Ravi, Duke undergraduate student. Worked on building a practical implementation of our societal tradeoffs work. 2016—2017.
- Rachel Freedman, Duke undergraduate student. Interdepartmental B.A. in Artificial Intelligence Systems, 2017. Senior thesis advisor, *Adapting a Kidney Exchange Algorithm to Incorporate Human Values*, high distinction. Now a Ph.D. student at UC Berkeley.
- Jeremy Fox, Duke undergraduate student. Worked on moral AI, with a focus on aggregating different people's moral classifiers in a social-choice-theoretic sense. 2016—2017. Now at Google.
- Max Kramer, Duke undergraduate student. Working on moral AI, with a focus on identifying the circumstances under which people are comfortable with AI making decisions. 2016—2017. Now a Ph.D. student in Philosophy at the University of Arizona.

- Melissa Dalis, Duke undergraduate student. B.S., Computer Science, Mathematics, Minor in Economics, 2014. Senior thesis advisor, *Beat the Cheater: Computing Game-Theoretic Strategies for When to Kick a Gambler out of a Casino*, high distinction and Alex Vasilos award. Now a Data Science Manager at Instagram.
- Siyang Chen, Duke undergraduate student. Math B.S., CS B.S. 2012. Senior thesis co-advisor (with Kamesh Munagala), *New Directions in Mechanism Design for Facility Location*, highest distinction. (Also summer research in 2012.) Now at Facebook.
- Bo Waggoner, Duke undergraduate student. Math+CS B.S. 2011. Did research with me in Summer 2011 (resulting in the AAAI 2012 paper *Evaluating Resistance to False-Name Manipulations in Elections*, listed above), now Assistant Professor of Computer Science at UC Boulder.
- Chengyu Li, Duke undergraduate student. Senior thesis co-advisor (with Huseyin Yildirim), *Matching Offers in the Provision of Public Goods*, 2011.
- Peng Shi, Duke undergraduate student. Senior thesis advisor & C-SURF. Worked on prediction markets and related topics. Thesis title: *Prediction Mechanisms That Do Not Incentivize Undesirable Actions* (highest distinction). CRA Outstanding Undergraduate Research Award Finalist. Math B.S. + CS B.A. + Econ minor 2010, now an Assistant Professor at USC Marshall School of Business' Data Science and Operations group.
- Matthew Rognlie, Duke undergraduate student. Johnson RA and C-SURF. Worked on interpreting voting rules as maximum likelihood estimators, and repeated games in which players can restart the game. Duke Faculty Scholar award. Econ+Math B.S. 2010, now an Assistant Professor of Economics at Northwestern.
- Maggie Bashford, Duke undergraduate student. C-SURF. Worked on how hard it is typically to control elections. CS+Econ B.S. 2010. Now at Deloitte.
- Everett Wetchler, Duke undergraduate student. Senior-year independent study on computer poker. Spring 2007.
- Zhijian Lim, CMU undergraduate student. Senior project on building an expressive negotiation server for donations to charities. 2003—2004.

Committees of other students

... outside of home institution:

- Duncan McElfresh, Ph.D. student in Applied Mathematics, University of Maryland, College Park, MD, USA. Committee member; defended 8/27/2021. Thesis title: *Practical Algorithms for Resource Allocation and Decision Making*.
- Zhe Feng, Ph.D. student in Computer Science, Harvard University, Cambridge, MA, USA. Committee member; defended 4/28/2021. Thesis title: *Machine Learning-Aided Economic Design*.
- Anson Kahng, Ph.D. student in Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA. Committee member; defended 7/9/2021. Thesis title: *Computational Perspectives on Democracy*.
- Dominik Peters, Ph.D. student in Computer Science, University of Oxford, Oxford, England, UK. Examiner; defended (“viva”) 9/2/2019. Thesis title: *Fair Division of the Commons*.
- Christian Kroer. Ph.D. student in Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA. Committee member; defended 9/12/2018. Thesis title: *Large-Scale Sequential Imperfect-Information Game Solving: Theoretical Foundations and Practical Algorithms with Guarantees*.

- Haifeng Xu. Ph.D. student in Computer Science, University of Southern California, Los Angeles, CA, USA. Committee member; defended 5/9/2018. Thesis title: *Information as a Double-Edged Sword in Strategic Interactions*.
- Nisarg Shah. Ph.D. student in Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA. Committee member; defended 6/28/2016. (Also visited my group in the summer of 2015.) Thesis title: *Optimal Social Decision Making*.
- Bo Waggoner. Ph.D. student in Computer Science, Harvard University, Cambridge, MA, USA. Committee member; defended 5/27/2016. Thesis title: *Acquiring and Aggregating Information from Strategic Sources*.
- Hau Chan, Ph.D. student in Computer Science, Stony Brook University, Stony Brook, NY, USA. Committee member; defended 6/5/2015. Thesis title: *Game-Theoretic Models for Interdependent Security: Modeling, Computing, and Learning*.
- Sam Ganzfried, Ph.D. student in Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA. Committee member; defended 4/14/2015. Thesis title: *Computing Strong Game-Theoretic Strategies and Exploiting Suboptimal Opponents in Large Games*.
- Tyler Lu. Ph.D. student in Computer Science, University of Toronto, Toronto, ON, Canada. External examiner; defended 3/9/2015. Thesis title: *Group Decision Making with Partial Preferences*.
- Rong Yang, Ph.D. student in Computer Science, University of Southern California, Los Angeles, CA, USA. Thesis committee member; defended 4/3/2014. Thesis title: *Addressing Human Decision Making in Security Games: Models and Algorithms*.
- Manish Jain, Ph.D. student in Computer Science, University of Southern California, Los Angeles, CA, USA. Thesis committee member; defended 4/10/2013. Thesis title: *Thwarting Adversaries with Unpredictability: Massive-scale Game-Theoretic Algorithms for Real-world Security Deployments*. Best Dissertation Award, Computer Science Department, Viterbi School of Engineering, University of Southern California.
- John K. Lai, Ph.D. student in Computer Science, Harvard University, Cambridge, MA, USA. Thesis committee member; defended 4/29/2013. Thesis title: *Truthful and Fair Resource Allocation*.
- Travis Service, Ph.D. student in Computer Science, Vanderbilt University, Nashville, TN, USA. Thesis committee member; defended 3/22/2012. Thesis title: *Coalition Structure Generation in Characteristic Function Games*.
- Michael Zuckerman, Department of Computer Science, Hebrew University, Jerusalem, Israel. Committee member, Ph.D. thesis proposal title: *Algorithmic Game Theory and Computational Social Choice*, 4/2010. M.S. thesis evaluator, title: *Algorithms for the Coalitional Manipulation Problem*, 9/2007.
- Reshef Meir, Department of Computer Science, Hebrew University, Jerusalem, Israel. M.S. thesis evaluator, 12/2008. Thesis title: *Strategy Proof Classification*.
- Troels Bjerre Lund (formerly Sørensen), Ph.D. student at the Department of Computer Science, University of Aarhus, Denmark. Thesis committee member; defended 3/10/08. Thesis title: *Computing Equilibria of Two Player Games*.

... at home institution:

- Brian Hu Zhang, Ph.D. student in Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA. Thesis committee member; defended 7/3/2025. Title: *New Solution Concepts and Algorithms for Equilibrium Computation and Learning in Extensive-Form Games and Beyond*.

- James Michelson, Ph.D. student in Philosophy, Carnegie Mellon University, Pittsburgh, PA, USA. Thesis committee member; defended 7/26/2024. Title: *Essays in Design Economics*.
- Chun Kai Ling, Ph.D. student in Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA. Thesis committee member; defended 5/16/2023. Title: *Scalable Learning and Solving of Extensive-Form Games*.
- Gabriele Farina, Ph.D. student in Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA. Thesis committee member; defended 4/6/2023. Title: *Game-Theoretic Decision Making in Imperfect-Information Games: Learning Dynamics, Equilibrium Computation, and Complexity*.
- Waqar Aqeel, Ph.D. student in Duke's Department of Computer Science. Thesis committee member; defended 11/16/2021. Title: *The Latency Budget: How to Save and What to Buy*.
- Shao-Heng Ko, Ph.D. student in Duke's Department of Computer Science.
- Gayan Seneviratna, M.S. student in the Department of Electrical and Computer Engineering, 2019. Title: *Pomegranate: A Graphical User Interface for Synchrony Analysis in Video*.
- Dmitry Vagner, M.S. student in the Department of Computer Science, Duke University, 2019. Title: *Towards Type-Driven Compositional Learning*.
- Atefeh Mehrabi, Ph.D. student in Duke's ECE department.
- Kangning Wang, Ph.D. student in the Department of Computer Science, Duke University. Thesis committee member; defended 7/19/2022. Thesis title: *Approximations for Economic Efficiency and Fairness*.
- Reza Alijani, Ph.D. student in the Department of Computer Science, Duke University. Thesis committee member; defended 7/20/2020. Thesis title: *Algorithms for Two-sided Online Marketplaces*.
- Brandon Fain, Ph.D. student in the Department of Computer Science, Duke University. Thesis committee member; defended 3/8/2019. Title: *Algorithms for Public Decision Making*.
- Lok C. Chan, Ph.D. student in the Department of Philosophy, Duke University. Thesis committee member; defended 09/28/2018. Thesis title: *Deliberative Bayesianism: Abduction, Reflection, and the Weight of Evidence*.
- Seyed Majid Zahedi, Ph.D. student in the Department of Computer Science, Duke University. Thesis committee member; defended 03/19/2018. Thesis title: *Managing Shared Resources in the Data Center Era: Computer Architecture Meets Game Theory*. Joining the faculty at the University of Waterloo (ECE).
- Jason Pazis, Ph.D. student in the Department of Computer Science, Duke University. Thesis committee member; defended 10/6/2015. Thesis title: *PAC-optimal, Non-parametric Algorithms and Bounds for Exploration in Concurrent MDPs with Delayed Updates*.
- Peng Guan, Ph.D. student in the Department of Electrical and Computer Engineering, Duke University. Thesis committee member; defended 8/21/2015. Thesis title: *Topics in Online Markov Decision Processes*.
- Janardhan Kulkarni, Ph.D. student in the Department of Computer Science, Duke University. Thesis committee member; defended 5/1/2015. Thesis title: *Design of Scheduling Algorithms Using Game Theoretic Ideas*.
- Xiaoming (Nate) Xu, Ph.D. student in the Department of Computer Science, Duke University. Thesis committee member; defended 4/6/2015. Thesis title: *Network Externality and Mechanism Design*.

- Rolando Estrada, Ph.D. student in the Department of Computer Science, Duke University. Thesis committee member; defended 7/16/2013. Thesis title: *Tree Topology Estimation*.
- Sergiu Ungureanu, Ph.D. student in the Department of Economics, Duke University. Thesis committee member; defended 4/10/2013. Thesis title: *Essays on Prospect Theory, Dynamic Contracting and Procurement*.
- Christopher Painter-Wakefield, Ph.D. student in the Department of Computer Science, Duke University. Thesis committee member; defended 4/3/2013. Thesis title: *Sparse Value Function Approximation for Reinforcement Learning*.
- Sayan Bhattacharya, Ph.D. student in the Department of Computer Science, Duke University. Thesis committee member; defended 11/15/2012. Thesis title: *Auctions, Equilibria and Budgets*.
- Gavin Taylor, Ph.D. student in the Department of Computer Science, Duke University. Thesis committee member; defended 4/8/2011. Thesis title: *Feature Selection for Value Function Approximation*.
- Kareem Dana, M.S. student in the Department of Computer Science, Duke University, 2009. Title: *The Problem of Honest Shirking in Content Swarms*.
- Amita Devaraj, M.S. student in the Department of Computer Science, Duke University, 2009. Title: *Uncertainty Propagation in Analytic Models using Monte Carlo methods*.

Teaching

- CMU 15-784: Foundations of Cooperative AI. Spring 2025.
- CMU 15-281: Artificial Intelligence: Representation and Problem Solving. Spring 2025. Taught together with Tuomas Sandholm.
- CMU 15-326: Computational Microeconomics. Fall 2024.
- CMU 15-281: Artificial Intelligence: Representation and Problem Solving. Fall 2023. Taught together with Aditi Raghunathan. (Agreed to teach this as an extra course at the last minute due to another instructor becoming unavailable.)
- CMU 15-326: Computational Microeconomics. Fall 2023.
- CMU 15-784: Foundations of Cooperative AI. Fall 2022. Taught together with Caspar Oesterheld and Tuomas Sandholm.
- Duke COMPSCI 323 and 590.09: Computational Microeconomics. Instructor. Spring 2022. (Because this was my last semester at Duke and students at multiple levels still wanted to take a course like this, I taught the undergraduate and graduate versions of the course simultaneously.) Undergraduate version: Quality of course: 4.19/5.00, quality of instructor: 4.34/5.00. Graduate version: Quality of course: 4.17/5.00, quality of instructor: 4.58/5.00.
- Duke COMPSCI 570: Artificial Intelligence. Instructor. Fall 2021. Duke Trinity Arts and Sciences unexpectedly decided, near the end of the semester, to push the management of graduate course evaluations to the departments, so this course was rated only through a department-run Qualtrics survey. On that survey, 76% rated the course as excellent, and 24% as very good; 81% rated the instructor as excellent, and 19% as very good.
- Duke COMPSCI 590.7: Computational Microeconomics: Game Theory, Social Choice, and Mechanism Design. Instructor. Fall 2020. Quality of course: 4.67/5.00, quality of instructor: 4.83/5.00. **(Top 5% of all undergraduate instructors teaching in the Natural Sciences at Duke.)**

- Duke COMPSI 323: Computational Microeconomics. Instructor. Spring 2020. Quality of course: 4.25/5.00, quality of instructor: 4.54/5.00.
- Duke Philosophy 590.2: Ethics and AI. Co-instructor with Walter Sinnott-Armstrong and Jana Schaich Borg. Fall 2019.
- Duke COMPSI 270: Introduction to Artificial Intelligence. Instructor. Spring 2019. Quality of course: 4.11/5.00, quality of instructor: 4.30/5.00.
- Duke COMPSI 590.2: Computational Microeconomics: Game Theory, Social Choice, and Mechanism Design. Instructor. Fall 2018. Quality of course: 4.33/5.00, quality of instructor: 4.58/5.00.
- Duke COMPSI 223: Computational Microeconomics. Instructor. Spring 2018. Quality of course: 3.91/5.00, quality of instructor: 4.55/5.00.
- Duke Philosophy 590.1: Ethics and AI. Co-instructor with Walter Sinnott-Armstrong and Jana Schaich Borg. Spring 2018.
- Duke COMPSI 570: Artificial Intelligence. Instructor. Fall 2017. Quality of course: 4.69/5.00, quality of instructor: 4.84/5.00. **(Top 5% of all undergraduate instructors teaching in the Natural Sciences at Duke.)**
- Duke Philosophy 590.3: Ethics and AI. Co-instructor with Walter Sinnott-Armstrong and Jana Schaich Borg. Fall 2017.
- Duke COMPSI 590.2: Computation, Information, and Learning in Market Design. Instructor (taught together with postdoc Michael Albert). Spring 2017. Quality of course: 4.86/5.00, quality of instruction: 4.93/5.00, quality of instructor: 4.86/5.00.
- Duke COMPSI 590.4: Computational Microeconomics: Game Theory, Social Choice, and Mechanism Design. Instructor. Spring 2016. Quality of course: 4.50/5.00, quality of instruction: 4.78/5.00, quality of instructor: 4.88/5.00.
- Duke COMPSI 290.4/590.4: Crowdsourcing Societal Tradeoffs. Instructor. Spring 2015. Quality of course: 4.44/5.00 (290.4) and 4.75/5.00 (590.4), quality of instruction: 4.67/5.00 (290.4) and 5.00/5.00 (590.4), quality of instructor: 4.90/5.00 (290.4) and 5.00/5.00 (590.4).
- Duke COMPSI 570: Artificial Intelligence. Instructor. Fall 2014. Quality of course: 4.45/5.00, quality of instruction: 4.58/5.00, quality of instructor: 4.63/5.00.
- Duke COMPSI 590.4: Computational Microeconomics: Game Theory, Social Choice, and Mechanism Design. Instructor. Spring 2014. Quality of course: 4.76/5.00, quality of instruction: 4.88/5.00, quality of instructor: 4.81/5.00.
- Duke COMPSI 590.01: Linear and Integer Programming. Instructor. Fall 2012. Quality of course: 4.88/5.00, quality of instructor: 4.88/5.00.
- Duke COMPSI 173: Computational Microeconomics. Instructor. Spring 2012. Quality of course: 4.67/5.00, quality of instructor: 4.83/5.00.
- Duke COMPSI 296.1: Computational Microeconomics: Game Theory, Social Choice, and Mechanism Design. Instructor. Fall 2011. Quality of course: 4.79/5.00, quality of instructor: 4.93/5.00.
- Duke COMPSI 296.1: Linear and Integer Programming. Instructor. Fall 2010. Quality of course: 4.73/5.00, quality of instructor: 4.82/5.00.
- Duke COMPSI 173: Computational Microeconomics. Instructor. Spring 2010. Quality of course: 4.20/5.00, quality of instructor: 4.60/5.00.

- Duke COMPSCI 196.1/296.1: Computational Microeconomics: Game Theory, Social Choice, and Mechanism Design. Instructor. Fall 2009. Quality of course: 4.89/5.00, quality of instructor: 4.90/5.00. (**Top 5% of all undergraduate instructors at Duke.**)
- Duke COMPSCI 170: Introduction to Artificial Intelligence. Instructor. Spring 2009. Quality of course: 4.10/5.00, quality of instructor: 4.20/5.00.
- Duke COMPSCI 270: Artificial Intelligence. Instructor. Fall 2008. Quality of course: 4.62/5.00, quality of instructor: 4.77/5.00.
- Duke COMPSCI 196.1/296.2: Linear and Integer Programming. Instructor. (New graduate course.) Spring 2008. Quality of course: 4.78/5.00, quality of instructor: 4.89/5.00.
- Duke COMPSCI 196.2: Introduction to Computational Economics. Instructor. (New undergraduate course.) Fall 2007. Quality of course: 4.78/5.00, quality of instructor: 4.78/5.00.
- Duke COMPSCI 296.3: Topics in Computational Economics. Instructor. (New graduate course.) Spring 2007. Quality of course: 4.64/5.00, quality of instructor: 4.70/5.00.
- Duke COMPSCI 192.06: Computer Poker. Guided Everett Wetchler in an independent study. Spring 2007.
- Duke COMPSCI 296.2: Computational Game Theory and Mechanism Design. Instructor. (New graduate course.) Fall 2006. Quality of course: 4.63/5.00, quality of instructor: 4.63/5.00.
- CMU 15681: Machine Learning. Teaching Assistant. Fall 2003.
- CMU West eBusiness Technology program, task 16: Negotiation. Vertical Mentor. Summer 2003.
- CMU 16731/15780: Advanced AI Concepts. Teaching Assistant. Spring 2003.
- Harvard Math 20: Introduction to Linear Algebra and Multivariable Calculus. Teaching Assistant. Spring 1999.
- Guest lectures in Duke BA 513 (Choice Theory, at Fuqua), Duke CPS 102 (Discrete Math for Computer Science), Duke CPS 300 (Intro to Graduate Study), Duke CPS Curious Lunch (seminar series for summer research undergraduates), CMU 15780/16731 (Advanced AI), CMU 15892 (Foundations of Electronic Marketplaces), CMU 15681 (Machine Learning).

Professional Service

Editorial & program committee service

- Chair, One Hundred Year Study on Artificial Intelligence (AI100), 2023-2025. (Chair-Elect, 2021-2023.)
- Advisor, Cooperative AI Foundation, 2021-present. (First core advisor to this foundation.)
- Member, AAAI Presidential Panel on the Future of AI, 2024-2025.
- Co-organizer, Social Choice for AI Ethics and Safety workshop, 2023, 2024, 2025.
- Co-organizer, SLMATH (MSRI) program on Algorithms, Fairness, and Equity, 2023. (Also co-organizer for Connections Workshop and Introductory Workshop.)
- AAAI/ACM Conference on AI, Ethics, and Society (AIES) Steering Committee, 2017-present
- AAAI/ACM Conference on AI, Ethics, and Society (AIES) conference co-chair, 2022.

- ACM Transactions on Economics and Computation founding Editor-in-Chief (together with Preston McAfee). Started taking submissions in August 2011; completed second and final term in 2017.
- SIGecom Exchanges Editor-in-Chief, 2007-2010 (stayed on as co-editor for one issue in 2011 to facilitate transition)
- AAAI program co-chair, 2020 (7,737 full paper submissions).
- AAAI/ACM Conference on AI, Ethics, and Society (AIES) program co-chair, 2019.
- EC general chair, 2016 (179 registrations, not counting people who registered for the co-located GAMES instead)
- EC program co-chair, 2014 (271 participants, which includes the 29 Decentralization participants but not the 60 NBER participants)
- Co-editor, special issue of ACM TEAC on EC 2014 (published in 2016)
- AAMAS program co-chair, 2012 (577 participants)
- COMSOC program co-chair, 2010 (92 participants)
- Council of the Society for Social Choice and Welfare (SSCW; elected in 2023)
- AAAI Councilor (elected in 2014 for three-year term)
- Advisor, AI for Global Climate Cooperation, 2022
- Social Choice for AI Ethics and Safety (SC4AI) workshops, co-organizer, 2023-2025
- Workshop on Cooperative AI (held at NeurIPS) co-chair, 2020
- NetEcon program co-chair, 2017
- Workshop on Opinion Aggregation, Dynamics, and Elicitation (WADE) co-organizer, 2018
- WRANE co-organizer, 2010 (40 participants)
- Co-organizer, Dagstuhl Seminar on Computational Foundations of Social Choice, 2010
- ADT Organizing Committee, 2019
- AAAI Conference Committee, 2014-2023
- AAAI PFAB Working Group on the Future of the Conference, 2019-2023
- IJCAI Summer School in Beijing, co-chair, 2012
- AAMAS Blue Sky Ideas track chair, 2017
- FAccT area chair, 2023
- IJCAI-ECAI area chair, 2018
- AAAI area chair, 2010, 2019, 2023, 2026
- AAMAS area chair, 2018
- IJCAI area chair, 2011, 2013
- IJCAI 2013 Advisory Board / Executive Committee
- JAIR AI and Society Track Editor, 2019-present

- IFAAMAS board member, 2010-2016; served on the conference subcommittee
- JAIR Advisory Board, 2012-2015
- COMSOC Steering Committee, 2012-2014
- AMMA Steering Committee, 2009-present
- M-Pref Steering Committee, 2011-present
- Games and Economic Behavior Advisory Editor, 2017-present
- JAIR Associate Editor, 2009-2011.
- JAAMAS Associate Editor, 2007-2023.
- Artificial Intelligence editorial board, 2009-2016
- EC award committee (chair), 2013
- Organizing committee, AAAI Spring Symposium on Game Theory for Security, Sustainability and Health, 2012
- Tutorial chair for AAMAS, 2011
- Scholarship co-chair for AAMAS, 2010
- Tutorial chair for EC, 2009
- Publicity chair for AAMAS, 2008
- NeurIPS Area Chair, 2021
- AAMAS Senior Program Committee, 2008, 2010, 2011, 2013, 2014, 2016, 2025
- IJCAI Senior Program Committee, 2009, 2015, 2016
- AAAI Senior Program Committee, 2011, 2015, 2017, 2018
- AAAI Senior Member Track Program Committee, 2015, 2017
- EC Senior Program Committee, 2010, 2012, 2017, 2018, 2024, 2026
- AISTATS Senior Program Committee, 2011
- UAI Senior Program Committee, 2009
- IFAAMAS Victor Lesser Distinguished Dissertation Award selection committee, 2008, 2009
- Program Co-Chair, Multidisciplinary Workshop on Advances in Preference Handling (at AAAI-08), 2008
- AAAI/ACM Conference on AI, Ethics, and Society (AIES) Program Committee, 2018, 2021
- ADT Program Committee, 2009, 2011
- EC Program Committee, 2007, 2008, 2009, 2011, 2013
- Program committee of the Guanajuato Workshop on prior-free mechanism design, 2010
- LOFT Program Committee, 2010
- ISAIM Program Committee, 2008, 2010

- AISTATS Program Committee, 2009
- AAMAS Program Committee, 2006 (Best PC Member Award), 2007, 2009
- AMEC Program Committee, 2004, 2007, 2008, 2009
- AAAI Program Committee, 2005, 2006, 2007, 2008
- UAI Program Committee, 2005, 2006, 2007, 2008
- COMSOC Program Committee, 2006, 2008
- SOFSEM Program Committee, 2007
- TARK Program Committee, 2007
- NetEcon+IBC Program Committee, 2007
- IBC Program Committee, 2005, 2006

Other reviewing activities

- Journal reviewer for ACM Computing Surveys, ACM Transactions on Economics and Computation (TEAC), ACM Transactions on the Web (TWEB), Algorithmica, American Economic Review (AER): Insights, American Mathematical Monthly, American Philosophical Quarterly, Analysis, Annals of Mathematics and Artificial Intelligence (AMAI), Artificial Intelligence (AIJ), Australasian Journal of Philosophy, Autonomous Agents and Multi-Agent Systems (JAAMAS), Communications of the ACM (CACM), Computational Intelligence, Constraints, Data and Knowledge Engineering, Decision Support Systems, Economic Theory, Engineering, Ergo, Erkenntnis, Ethical Theory and Moral Practice, Games and Economic Behavior (GEB), Global Policy, IEEE Intelligent Systems, IEEE Pervasive Computing, IEEE Transactions on Systems, Man, and Cybernetics, Information Processing Letters, International Game Theory Review (IGTR), International Journal of Systems Science (IJSS), Internet Mathematics, Journal of the ACM (JACM), Journal of Algorithms, Journal of Artificial Intelligence Research (JAIR), Journal of Computer and System Sciences (JCSS), Journal of Economic Theory (JET), Journal of Machine Learning Research (JMLR), Machine Learning, Mathematical Logic Quarterly, Mathematics of Operations Research, Mathematical Social Sciences, Operations Research, Philosophy and Phenomenological Research, Philosophical Studies, Physica A: Statistical Mechanics and its Applications, PLOS ONE, Proceedings of the National Academy of Sciences (PNAS), PNAS Nexus, Public Choice, RAND Journal of Economics, Review of Economic Studies (REStud), Science, Science Advances, Scientific Data, SIAM Journal on Computing, SIGecom Exchanges, Social Choice and Welfare, Synthese, Theoretical Computer Science, Theoretical Economics, Theory and Decision, Theory of Computing Systems.
- Also reviewed papers for AAIM, ECAI, ESA, FCT, FOCS, GameNets, ICALP, ICS, NESCAI, NeurIPS, OSDI, SAGT, SODA, STACS, STOC, WINE.
- NSF panels, CAIF reviews, numerous other reviews of proposals, for positions, etc.

Other

- Organized field trips to Google, Yahoo, and Facebook (together with people in those companies) during the Simons Institute semester on Economics and Computation, 2015
- Site Judge, ACM Collegiate Programming Contest, MidAtlantic region, 2006

Departmental Service

- AI curriculum review committee, 2025-
- CSD Faculty Search Committee, 2024.
- Graduate Affairs Committee, 2020/2021, 2021/2022.
- Faculty Search Committee chair (AI/ML search), Duke Computer Science Department, 2019/2020.
- Served on a promotion committee, 2018.
- Faculty Search Committee chair, Duke Computer Science Department, 2016/2017, 2017/2018.
- Duke CS Quantitative Initiative (QI) committee, 2015/2016.
- Chair of a tenure committee, 2015.
- Graduate Program Committee, 2014/2015, 2017/2018.
- Faculty Search Committee co-chair, Duke Computer Science Department, 2011/2012.
- Executive Committee, Duke Computer Science Department, 2011/2012, 2012/2013, 2013/2014, 2014/2015, 2015/2016, 2016/2017, 2017/2018, 2018/2019, 2019/2020, 2020/2021, 2021/2022.
- Undergraduate Program Committee, Duke Computer Science Department, 2010/2011, 2011/2012, 2013/2014.
- Graduate Admissions Committee, Duke Computer Science Department, 2007, 2008, 2010.
- Graduate Curriculum Committee, Duke Computer Science Department, 2007, 2008, 2009, 2010.
- Faculty Search Committee, Duke Computer Science Department, 2009, 2010/2011 (and as chair or co-chair in later years, see above).
- Ph.D. Admissions Committee, CMU Computer Science Department, 2003, 2004.

Other University Service and Outreach

- Regularly discussing AI and related topics in the media (e.g., USA Today, Rolling Stone, The Washington Post, Fortune, The Independent, Wired, The Daily Beast, Vox, Business Insider, Lifewire, Pittsburgh Post-Gazette, CBS News Pittsburgh/KDKA).
- CMU SCS Presidential Advisory Board (PAB) Visioning, 2025.
- CMU Responsible AI Initiative, interim lead for the School of Computer Science, 2022-present.
- Palo Alto CMU alumni dinner themed around our *Moral AI* book, hosted by Jenny Belardi, 2024.
- Duke Center for Computational Thinking Steering Committee, 2020-2021.
- Duke Science and Society, Faculty Governance Committee, starting in 2019.
- The Lane Family Ethics in Technology program Oversight Committee, starting in 2019.
- Science and Technology Initiative Steering Committee, starting in 2018.
- Graduate Board of Visitors (GBOV) Meeting Panelist, 2018.
- Provost's committee on the sciences at Duke, 2017.

- With the Moral AI group, for Duke Donor Recognition Weekend 2016 we organized an event: *Moral Artificial Intelligence: A Guided Debate with the Kenan Institute for Ethics*. (240 guests registered for the event.) We did a similar event for the 2016 Kenan Advisory Board meeting. Also participated in various other AI-related events on Duke's campus including a NextEdFest panel on AI and Humanity: Ethical Impact and Challenges in October 2018, AI: Risks & Responses – a “Head-to-Head” risk discussion with the Duke Center on Risk in Science & Society in November 2018, a panel on AI and the economy at the 2019 Huang Fellows Symposium – Science & Soc[AI]ty: How AI Will Change the World in March 2019, a talk for Duke Data+ and Code+ undergraduates in June 2019, as a discussant in “Tech Ethics / Corporate Ethics: a dinner roundtable” at the Kenan Institute for Ethics in November 2019, as one of the speakers at “The Inside Shot” (“Do We Want Computers to Read Our Minds?”) in March 2020, and a talk for Duke CS+ undergraduates in June 2020.
- Associate Director, Information Initiative at Duke (iiD), starting in 2014.
- Helped with putting together Duke's Master's of Science in Economics and Computation (went live Dec. 9, 2013)
- Provost Lecture Series committee, 2012/2013
- Workshops on Modeling the Undergraduate Residential System, Spring 2012
- Assessment of Computational/Informational Futures at Duke Committee, Spring 2012
- CS-ECON seminar series, 2011-present

Industry Employment

- **Avata Intelligence, Inc.** (formerly Armorway). Advisor. 2017—2019.
- **Econorithms, LLC**. Founder and President. 2016—present.
- **IBM Research (TJ Watson)**. Summer Co-op as IBM Ph.D. Fellow. Worked on rank aggregation methods with Andrew Davenport and Jayant Kalagnanam. 2005.
- **CombineNet, Inc.** Summer Research Scientist; Consultant. Worked on automated design of market mechanisms and other technical issues. 2003—2005.
- **Oliver, Wyman and Company, LLC**. Summer Intern (Management Consultant). Member of team consulting major Dutch bank on BIS II Capital Accord compliance. 2001.
- **Capital One Services, Inc.** Summer Business Analyst, Internet Team. Assessed marketing value of Capital One's e-mail address database. 2000.
- **Robeco Group**. Summer Intern, Quantitative Research Division. Created and analyzed statistics for performance evaluation of portfolio managers. 1999.

Personal Information

- Born October 1978, the Netherlands. Grew up in Amsterdam. Also a naturalized US citizen. Married; three children.
- *Sports*: current: soccer, running; past: judo, tennis.
- *Languages*: Dutch (native), English, German, French, Modern Hebrew, Latin, Ancient Greek.