

EC 2017 Workshop on Elicitation and Forecasting

Rafael Frongillo¹, David Rothschild², and Bo Waggoner³

¹CU Boulder. raf@colorado.edu

²Microsoft Research. DavidMR@Microsoft.com

³UPenn. bwag@seas.upenn.edu

Abstract. Forecasting is as old as civilization itself, but the game has changed with the advent of the internet. Gathering, incentivizing, and aggregating information from the crowd, as well as tapping vast sources of other data, has become the cornerstone of modern forecasting techniques. This workshop aims to bring together *theoreticians, empiricists, and practitioners* discuss the elicitation and aggregation of information for prediction making, which has been an emerging theme in the EC community over the last several years. To foster inclusion both across areas and to newcomers, the format mixes evenly between high-level talks from invited speakers and contributed talks on recent work. This will be the first incarnation of the workshop.

Core Topics. The areas of forecasting and elicitation have deep historical roots in game theory, statistics, and computer science. The workshop identifies and focuses on three aspects of this area:

1. Aggregation: Forecasting, polling, prediction, e.g. [6, 9, 15, 4].
2. Elicitation: Evaluating and incentivizing forecasters, e.g. [13, 11, 5].
3. Combining both: Prediction markets [1, 2, 3, 7, 10, 8], wagering mechanisms [12, 6, 14], etc.

Each of these topics will include *both theoretical and empirical perspectives*; e.g. the elicitation block will include talks about correcting bias, designing interfaces, differing risk attitudes, and other practical concerns. The workshop will bring together researchers from the theoretical and practical sides of these areas, as well as interested newcomers.

Schedule. The workshop will span a full day from approximately 9:30am to 4:30pm, divided into three blocks, two before lunch and one after, each corresponding to a bullet point above; following the blocks, the workshop will conclude with a poster session. Each block begins with an invited talk of 45 minutes to one hour. Several short 15-minute talks (selected from submissions) will follow. We will leave ample time for comments and discussion, including a short discussion period with the invited speaker.

1. **Aggregation/Forecasting.** Despite advances in computing and statistics, creating forecasts from historical or polling data has long been a relatively stable process. Recently, however, larger and less controlled datasets from social media or online behavior have challenged researchers to create new tools at the intersection of machine learning and time series econometrics. This block will explore some cutting edge work in how we make modern data-driven forecasts. This audience in particular will appreciate talks at the intersection of machine learning and econometrics, where computer science and economics come together.

2. **Elicitation/Evaluation.** Given a single agent, how can we obtain their best guess or most accurate prediction about an event? This turns out to be a surprisingly deep question in theory and a thorny one in practice. This block explores both angles with a particular goal of bringing out recent challenges where theory and practice meet. Examples include human biases and lay interpretation of probability all the way to theoretical impossibility results in information elicitation. This topic is particularly fun for diverse-interest audience because problems that are easy in theory (e.g. eliciting a probability distribution) are often challenging in practice and vice versa.
3. **Combinations.** Recent years has seen success in theory and practice for *mechanisms* designed to reveal predictions from groups of agents. Prediction markets are an exciting and active field of research and have significant impact both in the public space and within corporations. There are also intriguing alternative approaches such as wagering mechanisms. The goal of this block is to highlight a few cutting-edge directions in design of large-scale mechanisms for prediction, with an eye toward asking how they address the challenges raised in the previous blocks. We expect a vibrant dialog between practitioners and theoreticians, which may spark exciting research directions.

Logistics. Submissions for talks will be of papers or drafts in any format with two-page abstract submissions encouraged, longer submissions acceptable but not necessarily read beyond two pages. Strong preference will be given to more recent or unpublished work. Submissions will be reviewed by the organizers, possibly soliciting outside opinions from domain experts, against criteria of general/broad interest and potential to raise interesting, digestible points in a short talk format.

Organizer bios. Rafael Frongillo is an assistant professor of computer science at the University of Colorado Boulder. His thesis, *Eliciting Private Information from Selfish Agents*, was obtained in 2013 at the University of California at Berkeley advised by Christos Papadimitriou. He also spent time as a postdoctoral researcher at Microsoft Research New York and Harvard University's Center for Research, Computation, and Society and EconCS group.

David Rothschild is an economist at Microsoft Research. He has a Ph.D. in applied economics from the Wharton School of Business at the University of Pennsylvania. His primary body of work is on forecasting, and understanding public interest and sentiment. Related work examines how the public absorbs information. He has written extensively, in both the academic and popular press, on polling, prediction markets, social media and online data, and predictions of upcoming events; most of his popular work has focused on understanding the public's sentiment, an economist take on public policy, and choices in news consumption.

Bo Waggoner is a postdoctoral researcher at the Warren Center (supervisor Michael Kearns) at the University of Pennsylvania. His thesis, *Acquiring and Aggregating Information from Strategic Sources*, was obtained in 2016 at Harvard advised by Yiling Chen.

References

- [1] Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation*, 1(2):12, 2013. URL <http://dl.acm.org/citation.cfm?id=2465777>.

- [2] Jacob Abernethy, Sindhu Kutty, Sbastien Lahaie, and Rahul Sami. Information aggregation in exponential family markets. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 395–412. ACM, 2014. URL <http://dl.acm.org/citation.cfm?id=2602896>.
- [3] Jacob D. Abernethy, Rafael M. Frongillo, Xiaolong Li, and Jennifer Wortman Vaughan. A General Volume-parameterized Market Making Framework. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 413–430, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2565-3. doi: 10.1145/2600057.2602900. URL <http://doi.acm.org/10.1145/2600057.2602900>.
- [4] Yossi Azar, Amir Ban, and Yishay Mansour. When should an expert make a prediction? In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 125–142. ACM, 2016.
- [5] Christopher P. Chambers and Nicolas S. Lambert. Dynamically eliciting unobservable information. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 987–988. ACM, 2014. URL <http://dl.acm.org/citation.cfm?id=2602859>.
- [6] Yiling Chen, Nikhil R Devanur, David M Pennock, and Jennifer Wortman Vaughan. Removing arbitrage from wagering mechanisms. In *Proceedings of the fifteenth ACM conference on Economics and computation*, EC '14, pages 377–394. ACM, 2014.
- [7] Bo Cowgill and Eric Zitzewitz. Corporate prediction markets: Evidence from google, ford, and firm x. In *Proceedings of the fifteenth ACM conference on Economics and computation*, EC '14, pages 507–524. ACM, 2014.
- [8] Rachel Cummings, David M Pennock, and Jennifer Wortman Vaughan. The possibilities and limitations of private prediction markets. In *Proceedings of the 17th ACM Conference on Economics and Computation*, EC '16, pages 143–160. ACM, 2016.
- [9] Daniel G Goldstein, Randolph Preston McAfee, and Siddharth Suri. The wisdom of smaller, smarter crowds. In *Proceedings of the fifteenth ACM conference on Economics and computation*, EC '14, pages 471–488. ACM, 2014.
- [10] Hoda Heidari, Sébastien Lahaie, David M Pennock, and Jennifer Wortman Vaughan. Integrating market makers, limit orders, and continuous trade in prediction markets. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 583–600. ACM, 2015.
- [11] Nicolas S. Lambert and Yoav Shoham. Eliciting truthful answers to multiple-choice questions. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 109–118, 2009.
- [12] Nicolas S Lambert, John Langford, Jennifer Wortman, Yiling Chen, Daniel Reeves, Yoav Shoham, et al. Self-financed wagering mechanisms for forecasting. In *Proceedings of the 9th ACM conference on Electronic commerce*, EC '08, pages 170–179. ACM, 2008.
- [13] Nicolas S. Lambert, David M. Pennock, and Yoav Shoham. Eliciting properties of probability distributions. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 129–138, 2008.
- [14] Nicolas S Lambert, John Langford, Jennifer Wortman Vaughan, Yiling Chen, Daniel M Reeves, Yoav Shoham, and David M Pennock. An axiomatic characterization of wagering mechanisms. *Journal of Economic Theory*, 156:389–416, 2015.
- [15] Johan Ugander, Ryan Drapeau, and Carlos Guestrin. The wisdom of multiple guesses. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 643–660. ACM, 2015.