Overview

The goal of my research is to develop mathematically rigorous, empirically grounded frameworks to understand and design algorithms for eliciting and aggregating information, preferences, and beliefs. In pursuit of this goal, I seek to develop general methods that allow us to reason formally about the performance of algorithms with human components, in the same way that traditional computer science techniques allow us to formally reason about algorithms that run on machines alone. This agenda is inherently interdisciplinary, and my research draws heavily on ideas from economics, machine learning, optimization, and beyond.

In my early research, I designed foundational machine learning models to capture the most salient features of learning from data generated by heterogeneous crowds, and used these models to design and analyze novel learning algorithms. More recently, my research has focused on elicitation and aggregation using prediction markets, wagering mechanisms, and other crowdsourcing approaches, for which the design of appropriate incentives is key. My research has been supported by a National Science Foundation CAREER award, and has generated significant recognition from the research community, including a Presidential Early Career Award for Scientists and Engineers (PECASE) and Outstanding Paper or Best Student Paper awards at the top-tier computer science conferences ACM EC, COLT (twice), and UAI.

Prediction Markets and Crowdsourcing

A major theme of my research has been the design and analysis of markets for elicitation and aggregation of information from crowds. Here I describe several examples of this research.

Pricing in Combinatorial Prediction Markets: A *prediction market* is a market in which traders buy and sell securities with payments that are contingent on the outcome of a future event. For example, a security may yield a payment of $1 if a Democrat wins the 2016 US Presidential election and $0 otherwise. The *market price* of such a security is thought to reflect the traders' collective belief about the likelihood that a Democrat will win the election.

To facilitate trade, a prediction market can be operated by an *automated market maker*, an algorithmic agent that offers to buy or sell securities at some current market price determined by a *pricing mechanism* that sets prices as a function of the history of trade. The market maker provides liquidity, effectively subsidizing the market and rewarding traders for their private information. This is especially useful in *combinatorial markets*, which offer securities defined on a large or infinite outcome space and propagate information (in the form of prices) appropriately across logically-related securities. For example, in a combinatorial market over the Winter Olympics, the purchase of a security that is worth $1 if and only if the US wins the Pair Skating event should influence the market price of another security with payments contingent on the total number of medals won by the US. Such combinatorial markets maintain a (parameterized) joint probability distribution over the full outcome space (in this case, the space of all possible assignments of athletes to medals in every Olympic event) and can incorporate traders' information about any part of this distribution. Unfortunately, as I proved in joint work with Yiling Chen, Lance Fortnow, Nicolas Lambert, and David Pennock [6], many natural examples of combinatorial markets are intractable (specifically, #P-hard) to price using the most popular and well-studied pricing mechanism, Hanson's Log Market Scoring Rule [19].

In search of tractable pricing rules that would allow combinatorial markets to be run in practice, Jake Abernethy, Yiling Chen, and I proposed a general framework for the design of efficient pricing mechanisms over very large or infinite outcome spaces [1]. We took an axiomatic approach, specifying a set of formal mathematical properties that any reasonable market should satisfy (such as a "no arbitrage" property and an "information incorporation" property) and fully characterized the set of pricing mechanisms that satisfy these properties. Then, using techniques from convex analysis, we provided a method for designing specific pricing mechanisms that satisfy these properties. We described how to balance trade-offs between different features of these pricing mechanisms, such as the amount of money that the market operator loses in the worst case versus the granularity of its predictions. Using this framework, we provided efficient pricing mechanisms for sets of securities that were intractable to price with known techniques. This work grew out of observations that Yiling Chen and I had made about the fundamental mathematical connections between prediction market design and the notion of *regularization* in machine learning [5]. In particular, the mechanisms we derived are mathematically analogous to the "Follow the Regularized Leader" [35] class of learning algorithms, though the semantics are quite distinct.

Our framework [1] provided the foundations for what has grown into an active area of research, inspiring a growing body of follow-up work, both empirical and theoretical. The connections with no-regret learning have received significant attention in the machine learning community. Based on this body of research, I was invited to give tutorials on prediction market design at several top-tier computer science conferences (ICML, KDD, and AAAI) and the Machine Learning Summer School, to give a keynote talk at the 3rd AAAI Human Computation Workshop (HCOMP), and to co-organize workshops on relationships between learning problems and markets at premier machine learning conferences (NIPS and ICML). In addition to its impact on the research community, our framework served as the foundation for several market makers implemented in industry, including those used in the PredictWise election market [13] and previous and upcoming iterations of the Microsoft Prediction Lab.[1]

In order to better understand and characterize the space of possible pricing mechanisms for prediction markets, I studied several extensions of our basic framework. My former UCLA masters student Xiaolong Li and I worked with Jake Abernethy and Rafael Frongillo to show how the market maker can make a guaranteed profit given sufficient volume of trade and disagreement among traders [2, 32]. Yiling Chen, Mike Ruberry, and I extended the framework to allow the market maker to elicit distributions over *continuous* random variables [8]. To make the markets more practical, Miro Dudík, Rafael Frongillo, and I derived a principled method of reducing liquidity along certain dimensions of the market when particular information becomes less valuable to the market maker over time [14]. Our approach, which again relies on tools from convex analysis, allows the market maker to prevent traders from profiting as public information about partial outcomes (such as the Pair Skating medalists in an Olympics market) is revealed, while maintaining incentives to trade logically related securities (such as a security that pays off if the US wins at least ten medals). Finally, my Microsoft intern Hoda Heidari and I worked with Sébastien Lahaie and David Pennock to incorporate the ability to purchase *limit orders* into the framework while maintaining a certain notion of fairness to traders, a challenge in combinatorial prediction markets in which the prices of all interrelated securities change continuously as shares of any particular security are purchased [20].

In theory, the pricing mechanisms described above incentivize utility-maximizing traders with unlimited budget to "push" the market prices to match their beliefs. One important (and still fairly open) question is when and how market prices converge to reflect the traders' collective beliefs. My research has made progress towards this question, investigating the theoretical conditions under which information about a particular event of interest is aggregated at (perfect Bayesian) equilibrium in combinatorial markets [7], and

[1] `http://prediction.microsoft.com/`

studying the dynamics and convergence of market prices when traders are risk averse and have heterogeneous priors [34]. I am currently working on the design of a set of behavioral experiments to explore the ways in which individual traders extract information from the current market price.

Elicitation Via Single-shot Wagering Mechanisms: In practice, simply taking a (weighted) average of agents' forecasts about the likelihood of an event (possibly with additional post-processing) often yields an accurate prediction [37]. Building on the economics literature on *scoring rules*, I worked with Nicolas Lambert and other collaborators to derive a single-shot wagering mechanism for eliciting the individual forecasts of a group of agents with immutable beliefs. We proved that our mechanism is the only wagering mechanism to satisfy a set of desirable properties, including individual rationality (all agents want to participate), incentive compatibility (agents report their true beliefs), budget balance (the total payoff to all agents is $0), and sybilproofness (agents have no incentive to create false identities) [29]. This work received an Outstanding Paper Award at ACM EC, the premier conference on the intersection of computer science and economics, and was invited to appear in a special issue of the *Journal of Economic Theory*, a leading economics journal.

Several years later I worked with collaborators at Microsoft to show how to modify such wagering mechanisms in order to allow the market organizer to profit when the agents disagree, relaxing budget balance but maintaining many of the same desirable incentive properties [9].

Optimizing Crowdsourcing Markets: In crowdsourcing markets such as Amazon Mechanical Turk and Microsoft's Universal Human Relevance System, requesters post short "microtasks" that workers can complete in exchange for a small payment. A typical task might involve verifying the address of a company or labeling the content of an image. Crowdsourcing markets provide a way for requesters to inexpensively obtain distributed labor or information, and have recently become popular among researchers, who use them to conduct user studies, run behavioral experiments, and collect data that is easy for humans to generate but difficult for computers. Unlike in traditional labor markets, requesters interact with many workers rapidly and can adjust their requests in real time as they learn about salient features of the environment, such as workers' skill levels, the difficulty of their tasks, and workers' willingness to accept their tasks at given prices. This leads to a need for smarter repeated decision making in crowdsourcing markets.

With my doctoral student Chien-Ju Ho at UCLA, I addressed the challenge that a requester faces when assigning heterogeneous tasks to workers with unknown, heterogeneous skills [21]. We first formalized the *online task assignment* problem, in which a requester has a fixed set of tasks and a budget that specifies how many times he would like each task to be completed. We then provided a task assignment algorithm that is provably near-optimal if workers return repeatedly and we evaluated this algorithm on data collected from Mechanical Turk. Together with my former student Shahin Jabbari, we extended this line of work to cover classification or labeling tasks in which the quality of work cannot be judged immediately [22].

With Chien-Ju Ho and Alex Slivkins, I investigated algorithms for optimally selecting performance-based payments (for example, bonuses for work well done) for tasks when workers may strategically choose their level of effort [23]. We modeled this problem as a repeated version of the classical *principal-agent* problem from economics, and solved it using a new multi-armed bandit approach.

To understand whether the assumptions of this work are valid in practice, we worked with Sid Suri on behavioral experiments designed to study the effect of performance-based payments on the quality of work produced on Amazon Mechanical Turk [24]. Previous empirical studies of performance-based payments in crowdsourcing markets had produced mixed and somewhat contradictory recommendations. We set out to understand what caused this disparity in results, designing and running a sequence of experiments to understand not just whether performance-based payments improve quality for a specific task or payment scheme, but also the key properties of the payment, payment structure, and the task itself that make performance-based payments effective. Our results offer a simple explanation for why performance-based payments improve quality in some settings but not others, shedding light on the observations from past research. To close the loop, we used the results of these experiments to propose a new model of worker incentives, and showed that our theoretical results continue to hold under this new model.

Learning from Collective Preferences, Behavior, and Beliefs

The increasingly social nature of the Internet has led to novel sources of data on the preferences, behavior, and beliefs of massive populations of users. Naturally, both researchers and engineers are eager to apply techniques from machine learning to aggregate and make sense of this wealth of collective information, but traditional theories of learning fail to capture the complex issues that arise in these settings. In my doctoral dissertation, which received the University of Pennsylvania's Morris & Dorothy Rubinoff Award for "innovative applications of computer technology," I took steps towards addressing these problems by designing machine learning models to capture these complexities, and designing and analyzing new learning algorithms under these models.

As an example, the traditional "probably approximately correct" (PAC) learning model relies on an assumption that data is independent and identically distributed from a single source. When learning a personalized model for a single user of a system, such as a spam filter personalized for an individual email user, it is natural to consider training the model on labeled data collected from similar users. With Koby Crammer and Michael Kearns, I introduced a general PAC-style model to address the problem of learning from multiple sources of data with similar labeling functions [10, 11]. Our work established an algorithm and corresponding error bounds that clearly express a trade-off between three quantities: the sample size used, a weighted average of the distances of the data sources used from the source of interest, and the model complexity. With other collaborators at Penn, I derived new algorithms with uniform convergence bounds for the related but technically distinct *domain adaptation* problem [4].

Another issue ignored in traditional theories of learning is the need for data privacy. Consider a network of social contacts, and imagine that each party would like to compute his probability of having contracted a contagious disease, which depends on the exposures of his network neighbors. If network participants engage in standard belief propagation, each party would learn his probability of exposure conditioned on any evidence, but could also learn sensitive information about the exposure probabilities of his neighbors. Michael Kearns, Jinsong Tan, and I derived provably *privacy-preserving* versions of belief propagation, Gibbs sampling, and other local message-passing algorithms on large distributed networks [27]. Our methods show how to blend powerful tools from cryptography with local message-passing algorithms in a way that preserves the structures of the original computations and ensures that all messages appear to be randomly distributed from the viewpoint of any individual.

As a final example, there is an impressive literature in sociology and economics on mathematical models for collective behavior in settings as diverse as the diffusion of fads in social networks, voting behavior, housing choices, herding behaviors in financial markets, and Hollywood trends. It is natural to ask if there are learning methods tailored to such models and data. With Michael Kearns, I introduced a computational theory of learning from collective behavior, in which the goal is to accurately model and predict the future behavior of a large population after observing their interactions during a training phase of polynomial length [26]. This research not only led to interesting theoretical results, but also inspired us to undertake a set of related behavioral experiments examining the collective behavior of agents playing a biased voting game on networks, which appeared in *Proceedings of the National Academy of Sciences* [28].

Other Contributions in Algorithmic Economics and Machine Learning

In addition to my research on elicitation and aggregation, I have made contributions to other core problems in machine learning, including online learning [12, 15, 17, 18], active learning [3], and evolution [25]. This work was well received in the machine learning community, receiving Best Student Paper awards at COLT (twice) and UAI and an invitation to appear as a "Research Highlight" in *Communications of the ACM*.

I also made several early contributions to the literature on sponsored search auctions. With collaborators at Yahoo Research, I studied the exploration/exploitation trade-off involved in learning click-through rates [30, 31]. With Eyal Even-Dar and Michael Kearns, I initiated the first formal study of sponsored search in settings in which advertisers can bid not only on search terms but on search terms under specific *contexts*, such as queries from users of a specified demographic [16], as is now the industry standard.

Going Forward

To prove elegant and insightful theoretical results, it is often necessary to abstract away the messiness of the real world, modeling only the most salient features of the problem at hand. In order to identify those salient features—or even the right problems to solve—it is necessary to engage with domain experts and empirical or experimental researchers. This is especially true when considering systems with humans in the loop, since to obtain meaningful results, one must accurately model how the humans in the system actually behave. At this stage in my career, I believe that to have the most impact I must challenge myself to form lasting collaborations with empirical and experimental researchers, developing models that are informed by their experience and expertise, and at the same time, creating the opportunity for my algorithmic results to inform their future work.

Future Research: Crowdsourcing is an area in which the gap between theory and practice is especially wide. As Alex Slivkins and I described in a recent note [36], capturing all pertinent aspects of crowdsourcing in a single, coherent model is likely not possible for several reasons. First, there is a tension between the desire to study existing crowdsourcing platforms and the desire to introduce potentially better alternatives. Second, it is not clear how to best model the diversity of available tasks and workers. Third, workers are real human beings who may be strategic or seemingly irrational. As a result, many theoretical papers on crowdsourcing have introduced their own highly specialized models, making it difficult to compare results across papers. Furthermore, the problems addressed in these papers tend to differ from the problems that practitioners feel are the most important. In collaboration with researchers in online behavioral social science and crowdsourcing domain experts, I have begun to perform experiments that will give us insight into how to model crowdworkers and help us to identify the areas in which advances in algorithmic work would likely have the most impact. Our experiments on performance-based payments (described above) were a first small step in this direction. We are currently devising additional experiments to study other aspects of worker behavior that have been observed in anthropological studies of crowdworkers, such as long-term loyalty to requesters or offline collaboration patterns among workers. The ability to properly model these phenomena could impact both the design of algorithms for optimizing the output of existing crowdsourcing systems and the design of entirely new crowdsourcing paradigms.

I have also begun to complement my algorithmic work on prediction market design with empirical studies of trader behavior. To me, the most pressing open problem in the study of prediction markets is to understand when and how prediction markets aggregate the private information and beliefs of traders in practice. Real-world traders are not perfectly rational and Bayesian, as assumed in much of the theory, yet prediction markets have been found to produce accurate forecasts in a variety of domains, even when play money is used. The aggregation process is even less well understood in combinatorial markets. As companies like Microsoft begin to run their own combinatorial markets, it is an important open question whether these markets do indeed yield more accurate predictions than collections of small independent markets, and if so, how. An attempt to answer these questions might involve both empirical or experimental studies aimed at modeling how traders distill information from market prices and make choices about which securities to trade, as well as a theoretical analysis of the circumstances under which these choices lead to aggregation. As a first step in this direction, I am currently designing an experiment to test how traders extract information from market prices, and more specifically, how this depends on the way in which information about the history of trade is presented to them. In addition to informing my algorithmic work, such experiments will immediately impact design decisions made for the Microsoft Prediction Lab and other real-world markets.

Shaping a New Subfield: The method of integrating behavioral studies, theoretical modeling, and algorithmic design that I am advocating here can be applied more broadly to other computational systems with humans in the loop. However, doing this properly requires active participation from researchers across interdisciplinary boundaries. For several years, I have worked to bring together researchers with different backgrounds to formulate and address key problems and begin to build up the foundations of social computing. With Winter Mason and Hanna Wallach, I organized two formative workshops held at NIPS which brought together experts from economics, political science, psychology, sociology, machine learning, and statistics. These led to us guest edit a special issue of *Machine Learning Journal* on Computational Social Science and Social Computing [33], which included submissions from psychologists, cognitive scientists, linguists, and communications scholars, in addition to computer scientists. More recently, I have worked with Yiling Chen, Arpita Ghosh, and Tim Roughgarden to obtain funding from the Computing Community Consortium to hold a "visioning workshop" this June on the Theoretical Foundations of Social Computing,[2] bringing together experts across different fields to discuss and shape the future of this budding research area. My hope is that these efforts will lead us one step closer to developing general methods that allow us to reason formally about the performance of algorithms with human components in the same way that traditional computer science techniques allow us to formally reason about algorithms that run on machines alone.

References

[1] J. Abernethy, Y. Chen, and **J. W. Vaughan**. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation*, 1(2):Article 12, 2013.

[2] J. Abernethy, R. Frongillo, X. Li, and **J. W. Vaughan**. A general volume-parameterized market making framework. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, 2014.

[3] M.-F. Balcan, S. Hanneke, and **J. W. Vaughan**. The true sample complexity of active learning. *Machine Learning Journal*, 80(2–3):111–139, 2010.

[4] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and **J. W. Vaughan**. A theory of learning from different domains. *Machine Learning Journal*, 79(1–2):151–175, 2010.

[5] Y. Chen and **J. W. Vaughan**. A new understanding of prediction markets via no-regret learning. In *Proceedings of the Eleventh ACM Conference on Electronic Commerce (EC)*, 2010.

[6] Y. Chen, L. Fortnow, N. Lambert, D. Pennock, and **J. Wortman**. Complexity of combinatorial market makers. In *Proceedings of the Ninth ACM Conference on Electronic Commerce (EC)*, 2008.

[7] Y. Chen, M. Ruberry, and **J. W. Vaughan**. Designing informative securities. In *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2012.

[2] `http://www.cra.org/ccc/visioning/visioning-activities/social-computing`

[8] Y. Chen, M. Ruberry, and **J. W. Vaughan**. Cost function market makers for measurable spaces. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, 2013.

[9] Y. Chen, N. R. Devanur, D. Pennock, and **J. W. Vaughan**. Removing arbitrage from wagering mechanisms. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, 2014.

[10] K. Crammer, M. Kearns, and **J. Wortman**. Learning from data of variable quality. In *Advances in Neural Information Processing Systems 18 (NIPS)*, 2005.

[11] K. Crammer, M. Kearns, and **J. Wortman**. Learning from multiple sources. *Journal of Machine Learning Research*, 9:1757–1774, 2008.

[12] K. Crammer, E. Even-Dar, Y. Mansour, and **J. W. Vaughan**. Regret minimization with concept drift. In *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, 2010.

[13] M. Dudík, S. Lahaie, D. M. Pennock, and D. Rothschild. A combinatorial prediction market for the U.S. elections. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, 2013.

[14] M. Dudík, R. Frongillo, and **J. W. Vaughan**. Market makers with decreasing utility for information. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2014.

[15] E. Even-Dar, M. Kearns, and **J. Wortman**. Risk-sensitive online learning. In *Proceedings of the 17th International Conference on Algorithmic Learning Theory (COLT)*, 2006.

[16] E. Even-Dar, M. Kearns, and **J. Wortman**. Sponsored search with contexts. In *Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE)*, 2007.

[17] E. Even-Dar, M. Kearns, Y. Mansour, and **J. Wortman**. Regret to the best vs. regret to the average. *Machine Learning Journal*, 72(1–2):21–37, 2008.

[18] K. Ganchev, M. Kearns, Y. Nevmyvaka, and **J. W. Vaughan**. Censored exploration and the dark pool problem. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2009.

[19] R. Hanson. Combinatorial information market design. *Information Systems Frontiers*, (1):105–119, 2003.

[20] H. Heidari, S. Lahaie, D. Pennock, and **J. W. Vaughan**. Integrating market makers, limit orders, and continuous trade in prediction markets. In *Proceedings of the Sixteeth ACM Conference on Economics and Computation (EC)*, 2015 (to appear).

[21] C.-J. Ho and **J. W. Vaughan**. Online task assignment in crowdsourcing markets. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI)*, 2012.

[22] C.-J. Ho, S. Jabbari, and **J. W. Vaughan**. Adaptive task assignment for crowdsourced classification. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, 2013.

[23] C.-J. Ho, A. Slivkins, and **J. W. Vaughan**. Adaptive contract design for crowdsourcing markets: Bandit algorithms for repeated principal-agent problems. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, 2014.

[24] C.-J. Ho, A. Slivkins, S. Suri, and **J. W. Vaughan**. Incentivizing high quality crowdwork. In *Proceedings of the Twenty-Fourth International World Wide Web Conference (WWW)*, 2015 (to appear).

[25] V. Kanade, L. G. Valiant, and **J. W. Vaughan**. Evolution with drifting targets. In *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, 2010.

[26] M. Kearns and **J. Wortman**. Learning from collective behavior. In *Proceedings of the 21st Annual Conference on Learning Theory (COLT)*, 2008.

[27] M. Kearns, J. Tan, and **J. Wortman**. Privacy-preserving belief propagation and sampling. In *Advances in Neural Information Processing Systems 20 (NIPS)*, 2007.

[28] M. Kearns, S. Judd, J. Tan, and **J. Wortman**. Behavioral experiments on biased voting in networks. *Proceedings of the National Academy of Sciences*, 106(5):1347–1352, 2009.

[29] N. S. Lambert, J. Langford, **J. W. Vaughan**, Y. Chen, D. Reeves, Y. Shoham, and D. M. Pennock. An axiomatic characterization of wagering mechanisms. *Journal of Economic Theory*, 156:389–416, 2015.

[30] J. Langford, A. Strehl, and **J. Wortman**. Exploration scavenging. In *Proceedings of the 25th International Conference on Machine Learning (ICML)*, 2008.

[31] J. Langford, L. Li, Y. Vorobeychik, and **J. Wortman**. Maintaining equilibria during exploration in sponsored search auctions. *Algorithmica*, 58(4):990–1021, 2010.

[32] X. Li and **J. W. Vaughan**. An axiomatic characterization of adaptive-liquidity market makers. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, 2013.

[33] W. Mason, **J. W. Vaughan**, and H. Wallach. Computational social science and social computing (Guest Editorial). *Machine Learning Journal*, 95(3):257–260, 2014.

[34] R. Sethi and **J. W. Vaughan**. Belief aggregation with automated market makers. Working paper on SSRN, 2014.

[35] S. Shalev-Shwartz and Y. Singer. A primal-dual perspective of online learning algorithms. *Machine Learning Journal*, 69(2/3):115–142, 2007.

[36] A. Slivkins and **J. W. Vaughan**. Online decision making in crowdsourcing markets: Theoretical challenges (position paper). *ACM SIGecom Exchanges*, 12(2), 2013.

[37] J. Surowiecki. *The Wisdom of Crowds*. Anchor, 2005.