Privacy concerns are becoming a major obstacle to using data in the way that we want. It's often unclear how current regulations should translate into technology, and the changing legal landscape surrounding privacy can cause valuable data to go unused. For well-intentioned companies holding potentially sensitive data, the easiest and safest mode of compliance is simply to not use the data. For example, it is common practice for major companies to limit the sharing of data across departments, or to avoid storing sensitive data altogether. A recent Scientific American article described this occurrence: "as awareness of these privacy concerns has grown, many organizations have clamped down on their sensitive data, uncertain about what, if anything, they can release" [22]. Given these concerns, how can we continue to make use of potentially sensitive data, while providing rigorous privacy guarantees to the people whose data we are using? Answering this question requires tools to guarantee privacy in data analysis, as well as an understanding of how people reason about their privacy and how privacy concerns affect behavior.

In the last decade, a growing literature on *differential privacy* has emerged to address some of these concerns (see [15, 16] for a survey). First defined by Dwork et al. [17], differential privacy is a parameterized notion of database privacy that gives a mathematically rigorous worst-case bound on the maximum amount of information that can be learned about any one individual's data from the output of a computation. Differential privacy ensures that if a single entry in the database were to be changed, then the algorithm would still have approximately the same distribution over outputs. The privacy community has been prolific in designing algorithms that satisfy this privacy guarantee and maintain usefulness of the computation, resulting in a theoretical toolbox for a wide variety of computational settings, including machine learning, optimization, statistics, and algorithmic economics. Recently, major companies such as Apple and Google, and government organizations such as the United Status Census Bureau, have announced a commitment to implementing differentially private algorithms. However, many theoretical results cannot be directly applied by practitioners whose problems don't match the underlying assumptions.

In order to achieve a comprehensive solution to the privacy challenges we face today, we need to understand how existing tools for differentially private data analysis interact with *strategic and human aspects* of practical privacy guarantees. My work seeks to bridge the gap between theory and practice in the formal study of privacy. This includes problems such as strategic aspects of data generation, incentivizing truthful reporting of data, impacts of privacy policy, human decision making, and algorithm design. More broadly, I take a comprehensive approach to addressing real-world privacy challenges, using a diverse toolkit of both theoretical and practical perspectives.

Privacy concerns affect the interpretability of data

When people know that their current choices may have future consequences, they might modify their behavior to ensure that their data reveal less — or perhaps, more favorable — information about themselves, even if they do not fully understand the consequences of their actions. For example, a person may choose to not download a particular app, change their privacy settings, or even purchase a product from a different website. Indeed, 85% of adult Internet users have taken steps to avoid surveillance by other people or organizations [28]. **If analysts use the data of privacy-aware individuals for learning or inference, will the results still be meaningful?** The classical *revealed preferences problem* from economics (see [23] for a textbook treatment) models an observer making inferences about a consumer based upon her choices; in the absence of privacy concerns, an observer can see these choices and, after enough observations, learn the utility function that guides the consumer's decisions. However, when the consumer can take action to change what is revealed, existing tools are ill-equipped to describe such privacy-aware choices.

Together with Federico Echenique and Adam Wierman, I initiated the study of the testable implications of choice data, in settings where consumers are privacy-aware [10]. The main message of the paper is that **little can be inferred about a consumer's preferences** once we introduce the possibility that she has concerns about privacy. No matter what her behavior, she always has an "alibi" that can explain her choices as a consequence of privacy concerns. We showed that all possible behaviors on the part of the consumer are compatible with all possible preferences she may have over objects, even when her preferences are assumed to satisfy natural economic properties such as separability and monotonicity, which normally place strong restrictions on behavior. My coauthors and I give a constructive proof using tools from *economics* and *graph*

theory. We adapted the standard model of consumer choice theory to a situation where the consumer is aware of, and has preferences over, the information revealed by her choices. We represent privacy-aware preferences as a directed graph, where vertices of the graph are pairs of choice objects and inferences made by an observer. We imposed edges on this graph according to the constraints of our desiderata — rationalizability, monotonicity, and separability — and showed that the graph corresponds to a monotone and separable privacy-aware preference ordering that is consistent with the observed choice behavior.

Eliciting data from individuals

Much of the work in private data analysis starts from the premise that an analyst has access to a fixed database that is representative of the underlying population. But how does the analyst acquire this database? If she were to elicit it, why would anyone truthfully report their data? In many practical settings, people can misreport their data (e.g., change browsing behavior or lie on a survey) or refuse to participate (e.g., delete cookies or opt out of a service). These individuals may wish to influence the outcome of a computation performed on their data, or to mask their input due to privacy concerns. If they could potentially come to harm through the use of their private data, they may require additional compensation for this loss. One line of my work [8, 6, 14] has studied several key challenges an analyst faces when purchasing and aggregating data from strategic individuals with complex incentives and privacy concerns.

Together with my advisor Katrina Ligett, Aaron Roth, and fellow graduate students Steven Wu and Juba Ziani, I studied the procurement problem faced by **an analyst who must purchase and aggregate data from multiple sources** [8]. Individual data providers can sell a noisy, unbiased estimate of the same population statistic, where a more accurate estimate is more costly to provide. The analyst must elicit cost information from the providers, and purchase a collection of estimates that she can aggregate into a single estimator that satisfies a variance constraint. We give a dominant strategy truthful solution to this problem that yields an estimator with optimal expected cost, and violates the variance constraint by at most an additive term that tends to zero as the number of data providers grows large. We modeled the data analyst's problem as a *combinatorial optimization* problem, and showed that relaxing the problem to an LP and rounding the fractional solution only marginally violates the original variance constraint. Pairing our algorithm with VCG payments made truthful reporting of costs a dominant strategy for each data provider.

With Stratis Ioannidis and my advisor Katrina Ligett, I considered a setting where data is unverifiable — such as taste in movies or political beliefs — and **individuals are able to misreport their data to the analyst** [6]. Privacy-aware individuals hold data drawn according to an unknown linear model, which an analyst wishes to learn. The analyst can offer both a privacy guarantee and payments to incentivize players to truthfully report their data, and wishes to minimize her total payments and while still accurately estimating the model. We designed a truthful, individually rational mechanism that produced an asymptotically accurate estimate and allows the analyst's budget to diminish towards zero as the number of participants grows large. The main technical challenge was that differentially private computation of a linear model produces a biased estimate, and existing approaches for eliciting data from privacy-sensitive individuals do not generalize well to biased estimators. We overcame this using tools from *peer prediction* [25] to design our payment scheme, which leveraged the linear correlation of players' data to induce truthfulness.

During an internship at Microsoft Research, my hosts David Pennock and Jennifer Wortman Vaughan and I asked whether **existing techniques for data collection are compatible with differential privacy** [14]. We gave both positive and negative results for the design of private *prediction markets*: financial markets designed to elicit predictions about uncertain events. We first provided a class of private one-shot wagering mechanisms — in which bettors specify a belief about a future event and a monetary wager — that satisfy a number of desirable properties, including truthfulness, budget balance, and differential privacy of the bettors' reported beliefs. We then considered dynamic prediction markets, focusing our attention on the popular cost-function framework in which securities with payments linked to future events are bought and sold by an automated market maker. We showed that it is impossible for such a market maker to simultaneously achieve bounded worst-case loss and differential privacy, without allowing the privacy guarantee to degrade extremely quickly as the number of trades grows.

Impact of privacy policy

Given the promise of differential privacy, one tempting response to privacy concerns is regulation: lawmakers could mandate the use of differentially private algorithms or other privacy technologies, to limit the amount of information that firms can learn about consumers. An implicit assumption in the prior literature is that strengthening privacy protections will both increase utility for the individuals providing data and decrease usefulness of the computation. However, this assumption can fail when strategic concerns affect the impact and guarantees one can get from privacy technologies!

My joint work with Katrina Ligett, Mallesh Pai, and Aaron Roth serves as a cautionary tale against blindly setting privacy policy in strategic settings: the static effects of adding privacy technologies to a system may be the exact opposite of the effects in equilibrium [12]. We study how privacy technologies affect behavior in a simple economic model of data-driven decision making. A lender would like to use a consumer's past purchases to decide the terms of a new loan, but he is given only a differentially private signal about the consumer's behavior — which can range from no signal at all to a perfect signal, as we vary the differential privacy parameter. Using tools from *privacy* and *game theory*, we analyze end-to-end privacy guarantees of this game. We characterize equilibrium behavior as a function of the privacy level promised to the consumer, and show that **the effect of adding privacy in equilibrium can be highly counterintuitive**. Specifically, increasing the level of privacy can actually cause the lender to learn more about the consumer, and can also lead to decreased utility for the consumer and increased utility for the lender. We show that these quantities can generally be non-monotonic and even discontinuous in the privacy level of the signal. Our results demonstrate that even in simple models, privacy exhibits much richer behavior in equilibrium than compared to its static counterpart, and suggest that future policy decisions about privacy technologies ought to consider equilibrium effects.

The human side of privacy

Privacy is more than a mathematical construct; it exists in real-world settings, and its guarantees are impacted by human perceptions and reasoning. Existing work on differentially private algorithm design leaves the choice of privacy parameter as a task for the implementor, and typically assumes that a person's loss from sharing data is monotone and linear in the privacy parameter (e.g., [21, 29, 5, 27]). However, practical valuations of privacy can be highly complex, and **human reasoning about privacy need not match theoretical models**. In a series of ongoing *behavioral experiments* with Katrina Ligett and Ori Heffetz, I am testing how changing the privacy parameter in a simple differentially private algorithm affects human behavior, and how people trade off privacy for money.

A fundamental challenge in testing human perceptions of privacy is choosing which "embarrassing" data to use. People are more likely to value privacy of their real-world sensitive data — such as medical records or school transcripts — but using this data in experiments is typically neither ethical nor legal. Alternatively, we could provide participants with randomly generated data. However, since participants have no intrinsic value for privacy of this data, this approach is unlikely to provide useful insights into human reasoning about privacy. My co-authors and I chose to have participants generate potentially embarrassing data in the lab as a result of their own strategic behavior. In this way, we are simultaneously able to respect the real-world privacy needs of our human subjects, and use data which they may plausibly wish to protect. We have participants play a *dictator game*, where one person must unilaterally divide a fixed amount of money between themselves and a partner, and then we publicly announce a noisy (differentially private) version of they payment they gave to their partner. Participants can signal a taste for fairness or generosity by allocating a high payment to their partner, but doing so reduces their own payment. As we increase privacy of the announcement, their ability to signal is diminished, but the monetary impact of their choices remains the same. By varying the differential privacy parameter across treatments, we can use changes in behavior to measure how people trade off privacy for money. We plan to test for qualitative structural properties of behavior — including monotonicity and linearity in the privacy parameter — which can be used to inform future theoretical models for the value of privacy.

Other highlights

In addition to the publications described above, I've worked on a variety of other topics during graduate school. Here I give a brief description of my other research highlights:

- **Differential privacy as a tool in mechanism design:** With Michael Kearns, Aaron Roth, and Steven Wu [7], we used differential privacy as tool to solve the *equilibrium selection* problem. Our main result implemented private mediators in any large aggregative game, and gave algorithms which can select for the best Nash equilibrium, optimizing any linear objective function. In the process, we also gave the first method for solving a particular class of linear programs under the constraint of joint differential privacy.
- **Information complexity of coordinating distributed agents:** In joint work with my advisor Katrina Ligett, Jaikumar Radhakrishnan, Aaron Roth, and Steven Wu [13], we defined the *coordination complexity* of a problem, to be the amount of information that a fully-informed mediator needs to broadcast in order to coordinate the actions of distributed agents to play a nearly optimal solution. We gave upper and lower bounds on the coordination complexity of computing many-to-one matchings, stable matchings, Nash equilibria in routing games, and solving convex programs.
- **Differential privacy for generalization in machine learning:** A recent line of work [18, 19, 20] showed a connection between differential privacy and *generalization*, which is a guarantee that learning on a sample drawn from an unknown distribution is nearly as effective as learning from the distribution itself. I followed up on these results with Katrina Ligett, Kobbi Nissim, Aaron Roth, and Steven Wu, to study three alternative notions of generalization that have natural interpretations as privacy guarantees [11]. We provided a number of generic techniques for learning under these notions of generalization and proved that many common learning tasks can be carried out subject to these guarantees.
- **Learning algorithms for strategic environments:** Together with Kareem Amin, Lili Dworkin, Michael Kearns, and Aaron Roth [1], I studied learning-theoretic aspects of the classical revealed preferences problem from economics, where data is generated by strategic agents. We considered a setting where a consumer with a linear utility function repeatedly makes purchases from the same seller, who seeks to adapt prices to optimize his profits. Our main result was a polynomial time, no-regret algorithm for the seller's problem.
- **Molecular programming and DNA computing:** In two papers with collaborators Ho-Lin Chen, David Doty, and David Soloveichik, the first of which received a Best Paper Award at the International Symposium on Distributed Computing, we studied that computational power of *chemical reaction networks* [4, 9]. Our results characterized the classes of functions can be computed quickly [4] and correctly [9] under this computational model.

Future directions

Much of my Ph.D. work has emphasized that considering the *strategic and human aspects of privacy* are critical for our ability to collect useful data or interpret data via existing algorithms. In the future, I plan to build upon these insights to develop a formal theory for practical privacy guarantees. The two primary research directions I plan to pursue under this agenda are (1) developing a theory for the economics of data, and (2) understanding how people reason about privacy in practice. My future work will be a comprehensive approach that integrates all aspects of privacy — computational, strategic, human — to provide end-to-end privacy guarantees.

New techniques for selling data

Selling information and data is fundamentally different from selling a physical good, and requires new mathematical techniques to address the novel challenges in this problem space. The large body of work on market design and mechanism design provides many useful tools which can be brought to bear, but is so

far ill-equipped to address the unique challenges specific to selling information and data. For example, data can be freely reproduced or resold, like a *digital good*. However, unlike when purchasing a song or software, there may be sharp *externalities* among purchasers of data. Exclusive access to new market research may be of high value to a firm, but the firm's value for the data may be significantly lower if a competitor also had access to the information. Even for a single buyer, valuations for data are naturally *combinatorial* with a high degree of complementarities, as different data sources can be combined to be more useful in non-linear ways. For example, the two queries "average salary in the dataset" and "average salary in the dataset with Person X removed" are each innocuous on their own, but can be combined to reveal Person X's data. The infamous Netflix Challenge showed that complementarities are possible when combining entire datasets, when researchers were able to re-identify individuals by combining datasets from Netflix viewing history with IMDB movie ratings [26].

In addition, data can have a quantifiable quality level — such as bias or variance of an estimator, or sample size of a dataset — which can be costlessly and smoothly degraded by a data curator. Natural examples include a data analyst adding mean-zero noise to her estimate of a population statistic, or computing the estimate using only a subset of her data. Note that these techniques do not work in the reverse direction: an analyst must typically incur a cost to improve the quality of her data. This leads to a market design problem where buyers have *multi-dimensional preferences* for the data, parameterized by quality, and the seller has multiple copies of the data for each quality level. The seller is constrained by a maximum quality level (i.e., the quality of her unadulterated data), but can violate this constraint at a cost. I aim to develop a formal theory for the economics of information and data, by understanding how markets for data must differ from traditional markets, and combining tools from a diverse collection of relevant bodies of work.

Human decision-making for privacy

Human reasoning about the economic value of privacy plays a critical role in the growing practice of transacting on privacy. For example, Google collects data on its users in exchange for Internet browsing, video streaming, and email services. Grocery store loyalty programs track individual shopping habits in exchange for monetary discounts. Existing behavioral work has observed that people often appear irrational when making decisions about their privacy. For example, Brandimarte et al. [3] observed that the endowment effect exists for privacy: the payment that people require to sell their information is higher that what they will pay to protect it. In a performance art experiment, nearly half of participants gave their Social Security Number, mother's maiden name, and fingerprints in exchange for a cookie [2]. I plan to conduct a formal study of *human decision-making for privacy* through a series of *behavioral experiments*, each designed to test for measurable behavioral biases and offer qualitative and quantitative predictions of how people reason about privacy. The design of these experiments should be informed by both theory and practice: they should test for effects predicted in the theoretical literature on privacy and individual decision-making, as well as behavioral abnormalities observed in the real-world.

An alternative explanation for these apparent irrational privacy decisions is a *lack of understanding* — people may not understand the consequences of a given privacy guarantee, in terms of how their data will be collected, used, and shared. Indeed, the current practice of presenting consumers with impossibly long privacy policies is not designed to ensure understanding. Researchers estimated that if the average American were to read every privacy policy she encountered in a year, it would be her full time job for more than three months [24]. My proposed behavioral experiments can provide insight into which privacy technologies are the most effective in encouraging rational decision-making. For example, the extent to which a particular behavioral bias is observed may depend on the choice of differentially private algorithm. In addition, these experiments can be used to test the efficacy of varying techniques for explaining privacy technologies: the methods that improve understanding will reduce the presence and effect size of behavioral biases, as people will make better informed decisions.

More broadly, my work will continue to bridge the gap between a formal theory of privacy and practical applications of privacy technologies. I anticipate continuing to work on problems that that explore the interwoven threads of data, incentives, privacy, decision making, and human behavior.

References

[1] Kareem Amin, Rachel Cummings, Lili Dworkin, Michael Kearns, and Aaron Roth. Online learning and profit maximization from revealed preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, AAAI '15, 2015.

[2] Lois Beckett. People are willing to give away their personal data for a cinnamon cookie. `http://mashable.com/2014/10/01/data-for-cookies/`, 2014. [Online; accessed Oct. 30, 2016].

[3] Laura Brandimarte, Alessandro Acquisti, and George Loewenstein. Misplaced confidences: Privacy and the control paradox. *Social Psychological and Personality Science*, 4(3):340–347, 2013.

[4] Ho-Lin Chen, Rachel Cummings, David Doty, and David Soloveichik. Speed faults in computation by chemical reaction networks. *Distributed Computing*, pages 1–18, 2015. Preliminary Version appeared in the Proceedings of the 28th International Symposium on Distributed Computing (DISC 2014).

[5] Yiling Chen, Stephen Chong, Ian A. Kash, Tal Moran, and Salil Vadhan. Truthful mechanisms for agents that value privacy. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, EC '13, pages 215–232, 2013.

[6] Rachel Cummings, Stratis Ioannidis, and Katrina Ligett. Truthful linear regression. In *Proceedings of The 28th Conference on Learning Theory*, COLT '15, pages 448–483, 2015.

[7] Rachel Cummings, Michael Kearns, Aaron Roth, and Zhiwei Steven Wu. Privacy and truthful equilibrium selection for aggregative games. In *Proceedings of the 11th International Conference on Web and Internet Economics*, WINE '15, pages 286–299, 2015.

[8] Rachel Cummings, Katrina Ligett, Aaron Roth, Zhiwei Steven Wu, and Juba Ziani. Accuracy for sale: Aggregating data with a variance constraint. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*, ITCS '15, pages 317–324, 2015.

[9] Rachel Cummings, David Doty, and David Soloveichik. Probability 1 computation with chemical reaction networks. *Natural Computing*, 15(2):245–261, 2016. Preliminary Version appeared in the Proceedings of the 20th International Conference on DNA Computing and Molecular Programmings (DNA 2014).

[10] Rachel Cummings, Federico Echenique, and Adam Wierman. The empirical implications of privacy-aware choice. *Operations Research*, 64(1):67–78, 2016. Preliminary Version appeared in the Proceedings of the 15th ACM Conference on Electronic Commerce (EC 2014).

[11] Rachel Cummings, Katrina Ligett, Kobbi Nissim, Aaron Roth, and Zhiwei Steven Wu. Adaptive learning with robust generalization guarantees. In *29th Annual Conference on Learning Theory*, COLT '16, pages 772–814, 2016.

[12] Rachel Cummings, Katrina Ligett, Mallesh M. Pai, and Aaron Roth. The strange case of privacy in equilibrium models. In *Proceedings of the 17th ACM Conference on Economics and Computation*, EC '16, pages 659–659, 2016.

[13] Rachel Cummings, Katrina Ligett, Jaikumar Radhakrishnan, Aaron Roth, and Zhiwei Steven Wu. Coordination complexity: Small information coordinating large populations. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, ITCS '16, pages 281–290, 2016.

[14] 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, 2016.

[15] Cynthia Dwork. Differential privacy: A survey of results. In *Theory and Applications of Models of Computation*, pages 1–19. Springer, 2008.

[16] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(34):211–407, 2014.

[17] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Conference on Theory of Cryptography*, TCC '06, pages 265–284, 2006.

[18] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toni Pitassi, Omer Reingold, and Aaron Roth. Generalization in adaptive data analysis and holdout reuse. In *Advances in Neural Information Processing Systems, NIPS*, pages 2341–2349, 2015.

[19] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. Preserving statistical validity in adaptive data analysis. In *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing, STOC*, pages 117–126, 2015.

[20] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. The reusable holdout: Preserving validity in adaptive data analysis. *Science*, 349(6248):636–638, 2015.

[21] Arpita Ghosh and Aaron Roth. Selling privacy at auction. *Games and Economic Behavior*, 91:334–346, 2015. Preliminary Version appeared in the Proceedings of the 12th ACM Conference on Electronic Commerce (EC 2011).

[22] Erica Klarreich. Privacy by the numbers: A new approach to safeguarding data privacy by the numbers: A new approach to safeguarding data. `https://www.scientificamerican.com/article/privacy-by-the-numbers-a-new-approach-to-safeguarding-data/`, 2012. [Online; accessed Oct. 21, 2016].

[23] Andreu Mas-Colell, Michael D. Whinston, and Jerry R Green. *Microeconomic theory*. Oxford University Press, 1995.

[24] Aleecia M. McDonald and Lorrie Faith Cranor. The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society*, 4(3):540–565, 2008.

[25] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373, 2005.

[26] Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *Proceedings of the 2008 IEEE Symposium on Security and Privacy*, SP '08, pages 111–125, 2008.

[27] Kobbi Nissim, Salil Vadhan, and David Xiao. Redrawing the boundaries on purchasing data from privacy-sensitive individuals. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, ITCS '14, pages 411–422, 2014.

[28] Lee Rainie, Sara Kiesler, Ruogu Kang, and Mary Madden. Anonymity, privacy and security online. Technical report, Pew Research Center, 2013.

[29] David Xiao. Is privacy compatible with truthfulness? In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, ITCS '13, pages 67–86, 2013.