My research focuses on questions in **causal inference** arising in applications such as healthcare, e-commerce, policy-making, and finance. In particular, I am interested in **individual-level what-if** questions. For example, *what* happens to a patient's health *if* we prescribe them a drug, or *what* happens to a consumer's behavior *if* we expose them to a product? These individualized inferential tasks, i.e., determining *what happens to an individual's outcome if we do an intervention*, enable personalized data-driven decision-making, and complement conventional causal inference methods that focus on population-level inference, e.g., *what happens to the outcome averaged across the population if we do an intervention*?

Data-Rich Environments. Experimental data, collected through controlled trials where interventions are systematically assigned to subjects, are the gold standard for causal inference. However, trials are often costly, impractical, or ethically concerning. In such scenarios, we need to rely on **observational data**, i.e., data acquired without researcher manipulation. Deriving causal insights from such data is difficult, as the reasons for intervention assignment are unknown. However, the richness of modern observational data, **abundant in number of samples and auxiliary information** such as demographics and online behavior, presents a timely opportunity to develop methods for data-rich causal inference.

Biases. Observational data presents its own challenges. First, unobserved factors create **spurious associations** between interventions and outcomes known as **confounding bias**. This necessitates an emphasis on the principle that *correlation is not causation*, as we can never be certain of recording all the relevant factors. Second, **errors in data recording** lead to inaccurate estimates which introduce **measurement bias**. Consequently, we need principled methods to answer causal questions from observational data.

Beyond Conventional Causal Inference. I build methods for **fine-grained** causal inference, overcoming challenges of observational data by harnessing its richness, e.g., high dimensionality of an individual's sample. These methods, summarized in Figure 1, enable answering **individual-level rather than population-level** as well as **distribution-related instead of mean-related** what-if questions, e.g., *what happens to the distribution of an individual's outcome if we do an intervention*? To tackle confounding bias and measurement bias, I leverage tools and frameworks from machine learning and statistics. Below, I provide an overview of my contributions across three threads, each exploiting a different kind of structure to adjust for the unobserved factors.

1. Structure in distribution. I develop **computationally efficient alternatives to maximum likelihood estimation** for learning exponential family distributions [1, 2, 3] and use this framework to estimate

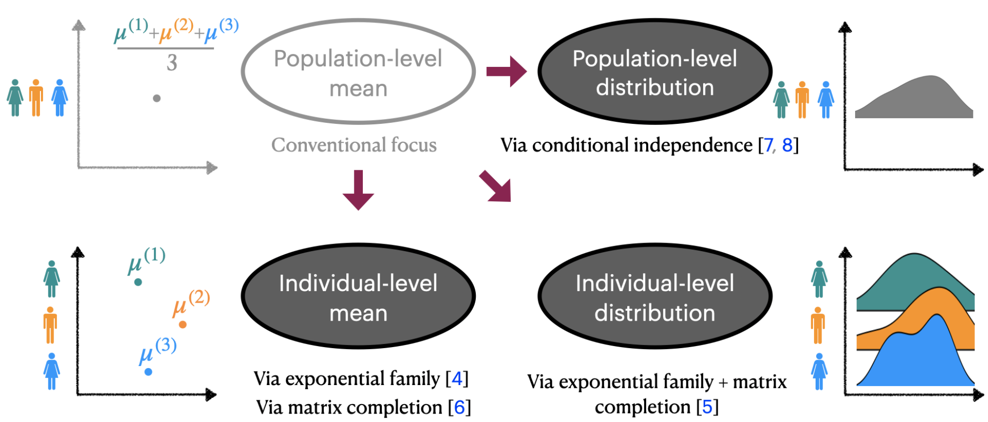

Figure 1: My works (in black) push beyond the conventional focus on population-level mean causal questions

individual-level mean of outcomes with just one sample per individual [4]. These estimators can also **denoise data with sparse measurement errors**. Furthermore, I integrate this method with matrix completion to infer the individual-level distribution of outcomes [5].

2. Structure in causal factors. I provide **a new matrix completion algorithm** to estimate **individual-level means of outcomes** by exploiting the relationships between outcomes and unobserved factors, as well as between interventions and unobserved factors [6]. The resulting estimates are **doubly robust**, i.e., they remain accurate even if either of the two aforementioned relationships is mis-specified.
3. Structure in causal graph. I design **data-driven conditional independence tests** to infer **population-level distribution of the outcome** [7, 8]. These tests identify subsets of observed covariates that account for unobserved factors, even with **limited knowledge** of the underlying causal graph.

Below, I describe my work in these threads in more detail along with ongoing and future directions.

1. Individual-level Counterfactual Inference with Exponential Family

Given an action-outcome pair, counterfactuals reveal the potential outcome, i.e., the outcome if an intervention (with a different action) had been implemented. Consider a movie streaming platform interacting with a customer, over many days, who watches a movie on the platform daily based on observed and unobserved factors. Given historical data of many customers, the platforms seeks to maximize every customer's viewing time and asks: *what would have happened to each customer's viewing time if they were exposed to a different sequence of movies?* The econometrics literature on panel data investigates such questions, representing potential outcomes as a tensor with units (individuals), measurements (days), and interventions (movies) as different axes. However, these works **do not allow outcomes and actions to depend on past outcomes and actions**. They also **require special structure** in the observed component of the tensor, e.g., there exists an intervention with observations for some measurements for all units.

Key Takeaway. We use exponential family to estimate the potential outcome tensor without imposing any structure on its observed component [4]. Our method accommodates (a) sequential dependence of outcomes and actions on past outcomes and actions, and (b) unseen interventions from a compact set.

Technical Contributions. We model the conditional distribution of outcomes as an exponential family and reduce learning the potential outcome tensor (with n units and p measurements) to learning parameters of n different distributions from the same exponential family, each with only one p-dimensional sample. Our convex estimator jointly learns all n parameter vectors and results in finite sample recovery rate of $O(p^{-1/2})$ for individual-level mean of outcomes. Our framework enables imputing sparsely missing unobserved factors and denoising data with sparse measurement errors. We integrate our estimator with matrix completion to infer individual-level distributions of outcomes when unobserved factors have a low-rank factorization [5]. Our method generalizes the estimators in our prior work [1, 2], which are computationally efficient alternatives to the maximum likelihood estimator (MLE) for learning exponential family distributions. We show that these alternative estimators are instances of a Bregman score and have many desirable properties, e.g., asymptotic normality and finite sample parameter recovery rate of $O(n^{-1/2})$ [3]. In ongoing work [9], we provide a unified framework encompassing various estimators for learning exponential family, filling a gap in the literature by offering finite sample rate of $O(n^{-1/2})$ for MLE.

2. Doubly Robust Causal Inference using Matrix Completion

In causal inference, model-based and design-based are two complementary identification strategies. The former employs restrictions on the process that determines how observed/unobserved factors affect potential outcomes, while the latter employs restrictions on the process that determines how observed/unobserved factors affect interventions. Doubly robust estimators combine model-based and design-based strategies to provide estimators that remain consistent as long as either of the two sets of restrictions is correct. In settings with no unobserved factors, these relationships are often estimated using machine learning, a technique known as double machine learning. However, despite their popularity, doubly robust estimators are **unavailable for settings with unobserved factors**, such as the panel data setting described earlier.

Key Takeaway. We develop a new matrix completion algorithm to provide doubly robust estimates of individual-level mean of potential outcomes in the presence of unobserved factors/confounding [6].

Technical Contributions. We model the relationships between potential outcomes and unobserved factors, as well as between interventions and unobserved factors by low-rank matrices. Under generic conditions on learning these matrices and for p-dimensional outcomes, we establish finite sample guarantees, asymptotic normality, and confidence intervals of our doubly robust estimator with rate $O(p^{-1/2})$.

3. Causal Effect Estimation with Limited Graph via Conditional Independence

Causal effect estimation involves learning the population-level distribution of the potential outcome. This differs from the conditional distribution due to the spurious associations caused by unobserved factors. Economists typically approach this task by assuming cause-effect relations implicitly through ignorability, involving conditional independence between counterfactual variables. By contrast, computer scientists typically assume cause-effect relations explicitly via causal graphs. However, ignorability is **an untestable assumption**, and the **full causal graph is often unknown**, especially in domains with many variables.

Key Takeaway. We estimate the population-level distribution of the potential outcome using data-driven conditional independence tests, even with limited knowledge of the causal graph.

Technical Contributions. Our tests identify a subset of observed variables and use them to construct a function that coincides with the population-level distribution of the potential outcome. These tests require only limited knowledge of the graph: either a causal parent of the intervention [7] or all causal children of the intervention [8]. To scale these tests in high-dimensions, we employ an out-of-distribution generalization framework as a continuous optimization-based approximation for conditional independence. We apply these tests to construct **fair credit risk models**, reducing the need for complete causal graph knowledge. In ongoing work [10], we generalize these tests to estimate the population-level distribution of the potential outcome for sequences of interventions, with applications to **clinical healthcare**.

4. Future Agenda: Leveraging Robustness and Fairness for Causal Inference

I am enthusiastic about **developing causal inference methods that are robust and fair**. Robustness is pivotal for ensuring reliable predictions, as it involves a comprehensive assessment of how various factors influence causal estimates. Likewise, algorithmic fairness plays a crucial role in promoting equity by mitigating disparities in causal predictions across protected groups, such as race or gender. Next, I outline **new problems and applications in causal inference** that I am excited to explore.

Spillover effects. Spillover effects are indirect impacts of interventions on individuals that were not the primary targets, e.g., the effect of a marketing strategy on unintended customer segments or a healthcare intervention on non-targeted patients. I plan to develop methods to quantify these spillover effects, enhancing the robustness of causal models and offering insights into broader impacts of interventions. I anticipate the exponential family framework to provide a systematic way to model these effects.

Sensitivity analysis. I am also keen to understand the impact of different structures/assumptions used to model unobserved factors on the robustness of causal estimates. Building on my prior work, I plan to investigate how causal estimates vary when unobserved confounding is only approximately captured by an exponential family, a low-rank matrix, or a conditional independence. I foresee the necessity for domain-specific sensitivity analyses in fields such as healthcare and policy-making, and I am enthusiastic about collaborating with experts in these domains to conduct rigorous and contextually relevant analyses.

Causality for Fairness. Many applications have inherent biases due to underlying causal relationships, such as health disparities based on genetic factors where individuals have no control over their genetics. I am excited to **integrate my research on causal inference and algorithmic fairness** [11, 12] to identify these biases and develop methods that acknowledge legitimate disparities while addressing unfair ones. To start, I intend to apply tools from my research on fair credit risk analysis to such problems.

Fair and causal selective prediction. Models that predict selectively, i.e., abstain from prediction, can magnify disparities between protected groups. To tackle this, we introduce a novel fairness notion for models that make fewer predictions, requiring improvements not only in overall performance but also in performance of every protected group [11]. Our method, featured in MIT News, reduces disparities in real-world **healthcare and criminal-justice** settings. Building on this foundation, I want to design *causal models* that are fair and can abstain from prediction, offering a new perspective on selective prediction.

Causal Fairness with uncertain protected attributes. In fair prediction, learners often lack access to accurate protected attributes due to collection bias, limited annotation, or legal regulations (e.g., California Consumer Privacy Act). To ensure fair predictions when protected attributes are uncertain, we propose a bootstrap-based algorithm [12] applicable to a variety of fairness notions and types of protected attributes. This algorithm achieves fairness comparable to the scenario with access to accurate protected attributes in real-world **healthcare and criminal-justice** settings. Leveraging this work, I plan to develop the first systematic framework to understand the effect of uncertain protected attributes in causal fairness.

References

[1] **A. Shah**, D. Shah, and G. Wornell, "A computationally efficient method for learning exponential family distributions," *Advances in Neural Information Processing Systems*, 2021.

[2] **A. Shah**, D. Shah, and G. Wornell, "On learning continuous pairwise Markov random fields," in *International Conference on Artificial Intelligence and Statistics*, PMLR, 2021.

[3] **A. Shah**, D. Shah, and G. W. Wornell, "On computationally efficient learning of exponential family distributions," *arXiv preprint arXiv:2309.06413*, 2023.

[4] **A. Shah**, R. Dwivedi, D. Shah, and G. W. Wornell, "On counterfactual inference with unobserved confounding," *arXiv preprint arXiv:2211.08209*, 2022.

[5] **A. Shah**, R. Dwivedi, D. Shah, and G. W. Wornell, "Learning counterfactual distribution under unobserved confounding with exponential family," *Working paper*, 2023.

[6] A. Abadie, A. Agarwal, R. Dwivedi, and **A. Shah**, "Doubly robust inference for causal latent factor models," *Working paper*, 2023.

[7] **A. Shah**, K. Shanmugam, and K. Ahuja, "Finding valid adjustments under non-ignorability with minimal dag knowledge," in *International Conference on Artificial Intelligence and Statistics*, PMLR, 2022.

[8] **A. Shah**, K. Shanmugam, and M. Kocaoglu, "Front-door adjustment beyond Markov equivalence with limited graph knowledge," *Advances in Neural Information Processing Systems*, 2023.

[9] **A. Shah***, J. J. Ryu*, D. Shah, and G. W. Wornell, "A unified view on learning exponential family distributions via noise-contrastive estimation," *Working paper*, 2023.

[10] **A. Shah**, S. Joshi, M. Kocaoglu, and K. Shanmugam, "Treatment effect estimation beyond sequential ignorability: An application to clinical healthcare," *Working paper*, 2023.

[11] **A. Shah***, Y. Bu*, J. K. Lee, S. Das, R. Panda, P. Sattigeri, and G. W. Wornell, "Selective regression under fairness criteria," in *International Conference on Machine Learning*, PMLR, 2022.

[12] **A. Shah**, M. Shen, J. J. Ryu, S. Das, P. Sattigeri, Y. Bu, and G. W. Wornell, "Group fairness with uncertainty in sensitive attributes," *arXiv preprint arXiv:2302.08077*, 2023.