Dear Selection Committee of the NBER Digital Economics and AI Tutorial,

I am writing to apply to the NBER Digital Economics and AI Tutorial. My name is Yiquan Ji, and I am a third-year PhD student in Decision, Risk, and Operations at Columbia Business School. I work with Professors Carri W. Chan, Hannah Li, and David Holtz. My research focuses on how algorithms and information technologies influence decision-making and resource allocation. In particular, I am interested in when algorithmic information improves outcomes, how operational constraints affect its value, and how AI systems can better account for differences across users.

My graduate training combines optimization and stochastic modeling with econometrics and statistical machine learning. My coursework has included linear programming, convex optimization, stochastic models, econometrics, empirical models, and statistical machine learning, and I am currently studying natural language processing. This training has shaped my interest in combining mathematical models that explain mechanisms with empirical analysis that evaluates their practical relevance. I hope to use this combination to study not only what AI systems can predict or generate, but also how their deployment changes decisions and outcomes.

In an ongoing project with Professors Carri W. Chan and Hannah Li, I study proactive care in hospitals with snooze-enabled early warning systems. These systems identify patients at risk of deterioration, but limited physician capacity constrains the interventions that can be provided. Alerts for patients not prioritized for immediate intervention may be snoozed, temporarily suppressing further warnings. We develop stochastic and fluid models to characterize optimal proactive-care policies and examine how snoozing affects system performance. Our analysis shows that proactive care can remain optimal despite the harm introduced by snoozing. We also identify conditions under which proactive care should cease beyond a snooze-duration threshold, as well as conditions under which it remains optimal even for arbitrarily long snooze durations. This project highlights a distinction that motivates my broader research: the value of a prediction technology depends not only on the information it produces, but also on the capacity and decision rules governing how that information is used.

In a second, early-stage project with Professor David Holtz, I am exploring whether explicitly modeling users' latent states, such as engagement and intent, can improve the decisions made by large language model agents in multi-turn conversations. The central question is whether explicitly matching actions to a user's evolving state improves assistance beyond what an agent achieves from the conversation history alone. For example, the appropriate next action might be to ask a clarifying question rather than immediately provide a recommendation. I am particularly interested in developing evaluations that distinguish improvements from state–action matching from those attributable to the underlying language model. This research direction reflects my broader interest in how the design of human–AI interactions affects the usefulness of AI-enabled services.

The tutorial would be especially valuable as I develop these projects into a broader research agenda. I hope to deepen my understanding of how economists study the effects of AI and digital technologies on productivity and resource allocation, and how theoretical mechanisms can be evaluated through experiments and other empirical approaches. Learning from researchers in digital economics and AI, exchanging ideas with other doctoral students, and attending the associated research meetings would help me refine both my research questions and evaluation strategies. I would also welcome the opportunity to contribute an operations perspective on the capacity constraints and dynamic trade-offs that arise when AI is deployed in practice.

Thank you for considering my application. I would be grateful for the opportunity to participate.

Sincerely,
Yiquan Ji
Third-Year PhD Student, Decision, Risk, and Operations
Columbia Business School
yj2836@columbia.edu