**Applicant Name:** Sahana Subramanyam

**Contact Information:**
Email: sahanasu@stanford.edu
Phone: +1 6502785784
Address: 128 Running Farm Lane, Apartment 102, Stanford CA 94305

**Graduate Study Details:**
I'm a 4th year economics PhD student at the Stanford Economics Department. My fields are behavioral and development economics. My primary advisors are Matthew Jackson, Arun Chandrasekhar and Muriel Niederle.

**Research Interest:**
My ongoing research interests are in trust and communication among human and AI agents. In developing countries, I study the impact of LLM-based information provisioning on social and economic networks, and in the lab, I design experiments to understand how humans make joint decisions with AI.

**Research In Progress (joint with Economics and / or CS PhD students)**

1. **LLM based agricultural advice**
   In collaboration with a startup providing AI-based agricultural support to smallholder farmers across India and Africa, this project examines how users engage with an LLM advisor in a real-world setting. Drawing on a large multilingual dataset of over 6,000 farmers and 120,000+ queries, I study what kinds of questions farmers ask, how their information needs evolve, and how trust in an AI advisor forms through repeated interactions. A key focus is whether these digital advisors act as substitutes or complements to existing sources of knowledge—such as peers, local networks, and human experts—in environments where social ties also shape information exchange, labor, and credit. By comparing AI-assisted advice to traditional pathways, the project explores how LLMs may alter social and economic networks, and what this means for learning, cooperation, and community outcomes.

2. **LLM interviews for embodied knowledge**
   This project develops a voice-based LLM interviewer to elicit open-ended accounts of individuals' experiences and interpretations of gender norms. By drawing on unstructured conversational data across groups, the system identifies locally salient concepts and distinctions that may be missed by researcher-designed instruments. These insights are then used to generate and refine structured survey items, enabling a more precise measurement of how gendered knowledge, expectations, and misperceptions are distributed across people.

3. **Safety evaluation for mental health chatbots**
   This project focuses on developing methods to evaluate the safety of long-form, multi-turn mental health chatbots used by large youth populations across Africa and India. In collaboration with Stanford's Trustworthy AI and Computing and Society groups, it aims to create context-aware safety metrics that can detect subtle forms of conversational failure and safety drift over time. The project also designs synthetic persona–based evaluation pipelines to simulate diverse user interactions and test guardrail responses before real-world deployment. The broader goal is to build a generalizable framework for assessing and improving the safety of generative agents in sensitive, high-stakes settings.

4. **Targeting fathers aspirations**
   This project tests whether engaging fathers, rather than mothers, with their children's academic performance and scholarship opportunities can shift educational aspirations and investment in rural India. Motivated by scoping work showing that fathers make most schooling decisions while mothers interface with schools, we examine whether correcting this information–decision-making mismatch improves outcomes. Through a randomized controlled trial across 16 government secondary schools, we study whether making report cards salient to fathers increases scholarship uptake, aspiration alignment within households, and school attendance. Because the intervention leverages existing school infrastructure, the findings have direct implications for the design and scalability of education programs that aim to reduce dropout and support secondary school completion.

5. **Motherhood penalty revisited: the role of household production**
   This project investigates the persistence of the motherhood penalty by studying how couples' beliefs, preferences, and bargaining processes around labor supply and household production evolve before and after the birth of their first child. Using a novel panel of 1,200 couples in Brazil, combined with an embedded field experiment, we measure how accurately spouses anticipate changes in housework and labor market outcomes, and whether miscalibrated beliefs or inefficient intra-household bargaining contribute to post-childbirth inequalities. By comparing informational and bargaining-based interventions, the project aims to identify which frictions drive the motherhood penalty and to generate evidence that can inform more effective policy design in low- and middle-income countries.

**Committee Members:** Matthew Jackson, Arun Chandrasekhar, Muriel Niederle**,** Doug Bernheim, Matthew Gentzkow