Contact Information:

Email: 14sungre@gmail.com, Phone Number: (513) 508-7654

Short description of Graduate Study:

I am a 5th year economics PhD, enrolled in the Business Economics program at the Olin Business School, Washington University in St Louis. My elective coursework has largely been in industrial organization and public economics. I plan on graduating in 2027.

Short description of research interests:

I am an empirical microeconomist with interests in empirical mechanism design and applied econometrics. I have a number of ongoing research projects related to each of these interests, which I will describe in turn. My interest in empirical mechanism design began when I was working for the chief economist (John List) at the rideshare firm, Lyft. At the time, Lyft was interested in exploring a subscriptions program where consumers could pay a fixed upfront fee each month in exchange for a percent discount off of all rides in the ensuing month. I am currently working on two related papers (with Aaron Bodoh-Creed, Brent Hickman, John List, and Ian Muir) using data from this experience. A first paper explores apparent “mistakes” consumers make in choosing whether or not the subscription program was right for them, and proposes novel methodology to nonparametrically identify both the structural primitives governing this mistake-making as well as the structural primitives of consumer preferences, despite these mistakes. A second paper takes a partial identification approach to tackle empirical challenges that arise when unobserved agent heterogeneity is higher-dimensional than the one-dimensional model typically assumed in many structural principal-agent models.

I also have a number of papers applying similar tools to study optimal income taxation. In a recent working paper with Matthias Rodemeier, we combine a meta-analysis of the literature estimating the elasticity of taxable income as well as a series of survey experiments to understand the extent to which positive versus normative considerations drive disagreement about optimal income tax policy. In a second paper, I study nonparametric identification of agent preferences in optimal income tax models, and in a third paper (which is my intended job market paper), I combine these results with

additional survey evidence to understand political frictions that may prevent the prescriptions of more standard optimal taxation models from being politically feasible.

I am currently also working on a set of projects, with my fellow graduate students at WashU (Kiara Kim and Nathan Mester) about artificial intelligence in the workplace, with an emphasis on how AI use may affect human capital accumulation. Our work here is largely motivated by the idea that in many career ladders, work tasks increase in complexity as workers move on to more advanced job titles, with simpler early career work providing the experiential building necessary to become competent at more complex late career work. Under this setup, there is a risk that the introduction of AI will automate away too much early-career work, thus robbing workers of a chance to build skills necessary for late-career work. We wish to build a microcosm of this sort of a learning process within a field experiment to measure if these concerns are quantitatively important. We have been partnering with the City University of New York to run this experiment with their computer science students. Our proposal received some initial seed funding from Schmidt Sciences and was also selected to be a part of the Griffin Applied Economics AI and Work Fellowship. In a related project, funded by the Wiedenbaum Center at WashU, we are also running a survey to understand to what extent workers are already perceiving the learning impacts AI has on their career growth.

Another set of projects I am working on concerns variance reduction in clinical trials. A large statistics and econometrics literature provides theoretical results about how to optimally leverage a *given* set of covariates to estimate average treatment effects. However, the practical question of which covariates to collect is more open-ended. Figuring out how to reduce variance is economically very valuable, because, for instance, a 10% variance reduction allows one to hold statistical power fixed while reducing required sample size by 10%. Given the high per-participant costs of running clinical trials, these cost savings can be substantial in dollar terms, yet very few trials actually seem to take advantage of covariates in treatment effect estimation. I am thus interested in searching for covariates that are scalable in the sense that they are broadly applicable across many clinical trial settings. A first paper is motivated by the fact that in about 20% of clinical trials, the primary endpoint is simply the change in some outcome measure, taken once at baseline and once at endline. Whenever this is the case, the baseline measure itself is a candidate covariate, and given typical autocorrelations in these medical outcomes, the variance reduction expected from using the baseline in covariate adjustment is roughly 25%. Despite this, 80% of these trials do not have any covariate adjustment in their statistical analysis plans. I am working on a paper that demonstrates just how economically costly not adjusting is for the medical research system at large. I am also exploring two more speculative ideas. First, I posit

that doctor predictions are in principle cheap to collect and in practice, seem quite predictive of patient outcomes. As a result, simply asking a doctor to predict patient health if they were in the control group seems like a promising way to introduce a highly predictive covariate. Second, I am interested in whether fitting an ML model on patient electronic health record data might also be a good way to get predictive power.

My advising committee consists of Brent Hickman, Barton Hamilton, and Bradley Larsen.