

Singapore Initiative on Digital Economics (SIDE) Workshop

7 March 2026, Saturday

SHHK-B1-11 (Seminar Room 9), at School of Social Sciences, NTU, 48 Nanyang Ave, Singapore 639818

Time	Program	Chair
9.00am-9.05am	Opening remarks By Yohanes Eko Riyanto – Chair of School of Social Sciences, NTU	Tat-How Teh (NTU)
9.05am-10.00am	Keynote: AI for Social Science: Generative Modeling Beyond LLMs By Will Cong - President’s Chair Professor in Finance, Computing and Data Science, NTU	
10.00am-10.45am	Reimagining the Firm with AI: Experimental Evidence from Startups Presenter: Hyunjin Kim - INSEAD Singapore Discussant: Pinchuan Ong - NUS	
Onsite Tea Break (30 minutes)		
11.15am-12:00pm	Ranking Algorithms, Learning, and Pricing Presenter: Sandro Shelegia - UPF and BSE Discussant: Matthew Leisten - FTC	Xijie Gao (NTU)
12:00pm-12:45pm	Measuring Welfare Gains from Digital Variety: Evidence from App-level Time Use in China Presenter: Xijie Gao - NTU Discussant: Wenxuan Xu - NUS	
Onsite Buffet Lunch (65 minutes)		
1.50pm-2.35pm	Open(?) AI Presenter: Matthew Leisten - Federal Trade Commission Discussant: Tat-How Teh - NTU	Kwok-Hao Lee (NUS)
2.35pm-3.20pm	GenAI Misinformation, Trust, and News Consumption: Evidence from a Field Experiment Presenter: Ruben Durante - NUS Discussant: Zhonglin Li - NUS	
Onsite Tea Break (30 minutes)		
3.50pm-4.35pm	Designing Consent: Choice Architecture and Consumer Welfare in Data Sharing Presenter: Chiara Farronato- Harvard Discussant: Tiffany Tsai - NUS	Tiffany Tsai (NUS)
4.35pm-5:20pm	Personalization, Data Portability and Competitive Dynamics Presenter: Julian Wright - NUS Discussant: Sandro Shelegia – UPF and BSE	
5.40pm	Shuttle Bus depart from NTU to the dinner venue	
6.30pm	Dinner reception New Ubin Seafood @ Chijmes	

- Keynote: 50-minute presentation + 5-minute Q&A
- All presentations: 40-minute presentation + 5-minute discussion and Q&A
- **Chair:** keep track of time and remind the speakers
- **Discussant:** prepare some brief comments / questions (up to 3 minutes), then open the floor for Q&A

- Shuttle van from **NTU** to the **dinner venue**, and from the dinner venue to **Park Avenue Rochester & NTU**
- Please let me (Tat-How) know if you wish to propose a change in the program / talk information.

Talk Information

AI for Social Science: Generative Modeling Beyond LLMs

By Will Cong

I characterize modern AI development as featuring two core themes: (i) goal-oriented end-to-end optimization in large modelling space, and (ii) generative pre-trained foundational models that enable economic world models and AI agents. Combining the insights from both, I identify promising directions in AI for Social Science Research.

To start, I introduce Goal-Oriented Algorithms in Large Space (GOALS) involving transformer-based reinforcement learning or panel trees, which are particularly suited for answering questions related to optimal decision-making and modelling grouped heterogeneity in social science research. In several specific financial applications, I show how GOALS can effectively and flexibly manage investment portfolios, generate test portfolios or latent factors for evaluating extant pricing models or more accurate pricing, and separate assets of higher and lower return predictabilities under different macroeconomic regimes. Among them, I highlight how GenAI based on GOALS can assist corporate decision-making that entails complex, high-dimensional, and non-linear stochastic control during which managers possessing various business objectives learn and adapt via dynamic interactions with the market environment.

With pre-trained foundational models and GOALS effectively capturing human agents' optimizing behavior in a given economy or market, or supplying the machine equivalent of that, I further introduce the concept of data-driven generative equilibrium for counterfactual analysis. Specifically, I show how one can take a data-driven approach to examine the counterfactual equilibrium in the online lending market when borrowers endogenously adopt LLMs to complete loan applications. Extending this further, I will discuss AI-Agent-Based modeling and how that can be combined with economic world models.

I conclude the talk with remarks on several caveats or unresolved issues when applying AI in social science research: (i) When using AI agents for experimentation or to generate counterfactual data, : we need to understand AI agents as a new species that potentially differ from humans, necessitating the new field of behavioral economics of AI; (ii) Large models and computations may not be the optimal path, and we need to effectively bridge theory with data; (iii) Before we can train AI social scientists, we need to build the ImageNet equivalent for empirical studies in the social sciences.

Reimagining the Firm with AI: Experimental Evidence from Startups

By Hyunjin Kim

AI has the potential to transform how firms are built and scaled, but its impact depends critically on whether they can identify where and how to apply this technology within the firm. Across 515 startups, we designed a field experiment that encouraged treated firms to explore how AI tools could be used to improve their ventures. We find that treated firms discover more uses for AI and, in doing so, build firms that are significantly more productive given the same labor and capital inputs: they complete more tasks, are more likely to launch products and acquire paying customers, and generate higher revenues conditional on earning. Despite faster growth, these firms do not scale inputs proportionally: their

demand for external investment falls by just over \$220,000 relative to the control group, while their demand for labor remains unchanged. Analyses of weekly progress reports indicate that these effects are not driven by simple awareness frictions or differences in technical capability. Instead, they arise from challenges entrepreneurs face in discovering how their products and processes can be reorganized around AI. Our findings suggest that while AI vastly expands what firms can do, realizing that potential depends on founders' and managers' ability to reimagine their firms around this general purpose technology.

Ranking Algorithms, Learning, and Pricing

By Heski Bar-Isaac and Sandro Shelegia

We study ranking algorithms that determine product exposure when sellers can set different on- and off-platform prices, raising the prospect of misallocation across sales channels. Initial uncertainty about product desirability is resolved through sales experience, leading to sellers having incentives akin to career concerns. Ranking design that relies on historical sales can replicate contractual price parity clauses (PPCs) but, with uncertainty, neglecting off-platform prices in ranking decisions can lead to distorted promotion of sellers who aim to boost on-platform sales and exposure. Explicitly conditioning on off-platform prices (including through PPCs) can raise both sales-channel allocation efficiency and promotion efficiency.

Measuring Welfare Gains from Digital Variety: Evidence from App-level Time Use in China

By Erik Brynjolfsson, Xijie Gao, and David Nguyen

This paper measures the welfare gains from the rapid expansion of digital apps using comprehensive app-level time-use data from China. We develop a framework in which apps are offered at zero monetary price but compete for users' time, allowing the welfare effects of new digital varieties to be inferred from observed reallocations of consumer attention. Applying the framework to data covering the universe of major digital apps between 2018 and 2025, we quantify the welfare impact of expanding digital variety. We find that the expansion of digital variety raises consumer welfare by approximately 8% over this period, an amount equivalent to roughly 20% of income growth. These findings highlight the importance of accounting for free digital goods in welfare analysis and demonstrate how time-use data can improve the measurement of consumer welfare in the digital economy.

Open(?) AI

By Matthew Leisten

I compare "open" AI model development to that of "closed" models. In my setting, a "founder" builds a general-purpose foundation model of endogenous quality. In the open paradigm, each member of a continuum of developers decides whether to fine-tune their own specialized application or to use a neighboring developer's application. In the closed paradigm, the founder centrally builds the specialized applications and sells access to them. I compare both to the social planner's optimum in terms of the number of applications, the degree of specialization of these applications, and the investment in the overall quality of the foundation model. The open paradigm features countervailing distortions that may lead to too little or too much specialization and too little or too much quality. The closed model yields the optimal quality investment but too many applications that are not specialized enough. In an extension, I consider the case where developers fine-tune open models primarily to gain clout.

GenAI Misinformation, Trust, and News Consumption: Evidence from a Field Experiment

By Filipe Campante, Ruben Durante, Ruben Durante, Ananya Sen

We study how AI-generated misinformation affects demand for trustworthy news, using data from a field experiment by a major German outlet, Süddeutsche Zeitung (SZ). Readers were randomly assigned to a treatment highlighting the challenge of distinguishing real from AI-generated images. The treatment raised concern with misinformation (0.3 s.d.) and reduced trust in news (0.1 s.d.), including SZ. Importantly, it affected post-survey browsing behavior: daily visits to SZ digital content rose by 2.5% in the immediate aftermath of the treatment. Moreover, we find that subscriber retention increased by 1.1% after five months, corresponding to about a one-third drop in attrition rate. Results are consistent with a model where the relative value of trustworthy news sources increases with the prevalence of misinformation, which may thus boost engagement with those sources even while lowering trust in news content.

Designing Consent: Choice Architecture and Consumer Welfare in Data Sharing

By Chiara Farronato, Andrey Fradkin, and Tesary Lin

We study the welfare consequences of choice architecture for online privacy using a field experiment that randomizes cookie consent banners. We study three ways in which firms or policymakers can influence choices: (1) nudging users through banner design to encourage acceptance of cookie tracking; (2) setting defaults when users dismiss banners; and (3) implementing consent decisions at the website versus browser level. Absent design manipulation, users accept all cookies more than half of the time. Placing cookie options behind extra clicks strongly influences choices, shifting users toward more easily accessible alternatives. Many users dismiss banners without making an explicit choice, underscoring the importance of default settings. Survey evidence further reveals substantial confusion about default settings. Our structural estimates show that among consent policies requiring site-specific decisions, consumer surplus is maximized when consent interfaces clearly display all options and default to acceptance in the absence of an explicit choice. However, the welfare gains from optimizing banner design are smaller than those from adopting browser-level consent, which eliminates the time costs of repeated decisions.

Personalization, data portability and competitive dynamics

By Andrei Hagiu and Julian Wright

We study data portability in a model of dynamic price competition where product quality improves for each customer as personal data builds up with usage. This can reflect AI-driven personalisation based on customer data or customers that embed their own data into the firm's product via their usage. We show that when consumers are forward-looking and firms can personalize their prices, each firm will offer data portability in equilibrium so that their customers can freely take their data to a rival. This ensures the rival no longer needs to initially subsidize consumers in an attempt to attract them so as to benefit in the future from their personal data. By contrast, when consumers are myopic, data portability still has a competition-softening effect, but firms no longer always benefit by offering it, so there are equilibria where neither firm offers data portability. A data portability policy is thus redundant when consumers are forward looking, but becomes potentially relevant when consumers are myopic. However, with myopic consumers a data portability requirement can also cause a complete market breakdown if consumers' willingness to pay for firms' products is initially (i.e., before their data has been used to improve the products) below cost. We extend these results to a setting without personalized prices.