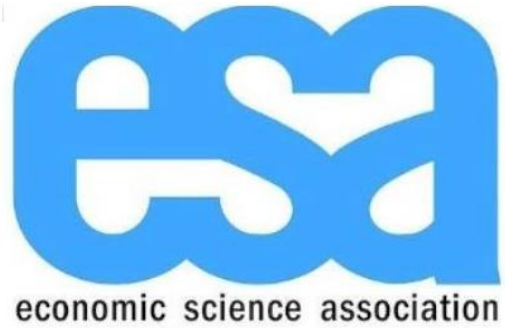

Final Program (with abstracts)

Asia-Pacific Meeting of the Economic Science Association Melbourne 2026

9 – 12 April 2026

The University of Melbourne

Venue Location

The Conference will take place at Melbourne Law School, The University of Melbourne, 185 Pelham Street, Carlton.

The conference will take place on both Ground Mezzanine and Level 1.

Parallel Sessions: Will be held in GM17, GM16, 0104, 0106, 0108, 0109

Plenary: Will be held in David P. Derham Theatre (GM15)

Registration: Level 1 Lobby

Catering: Level 1 Lobby

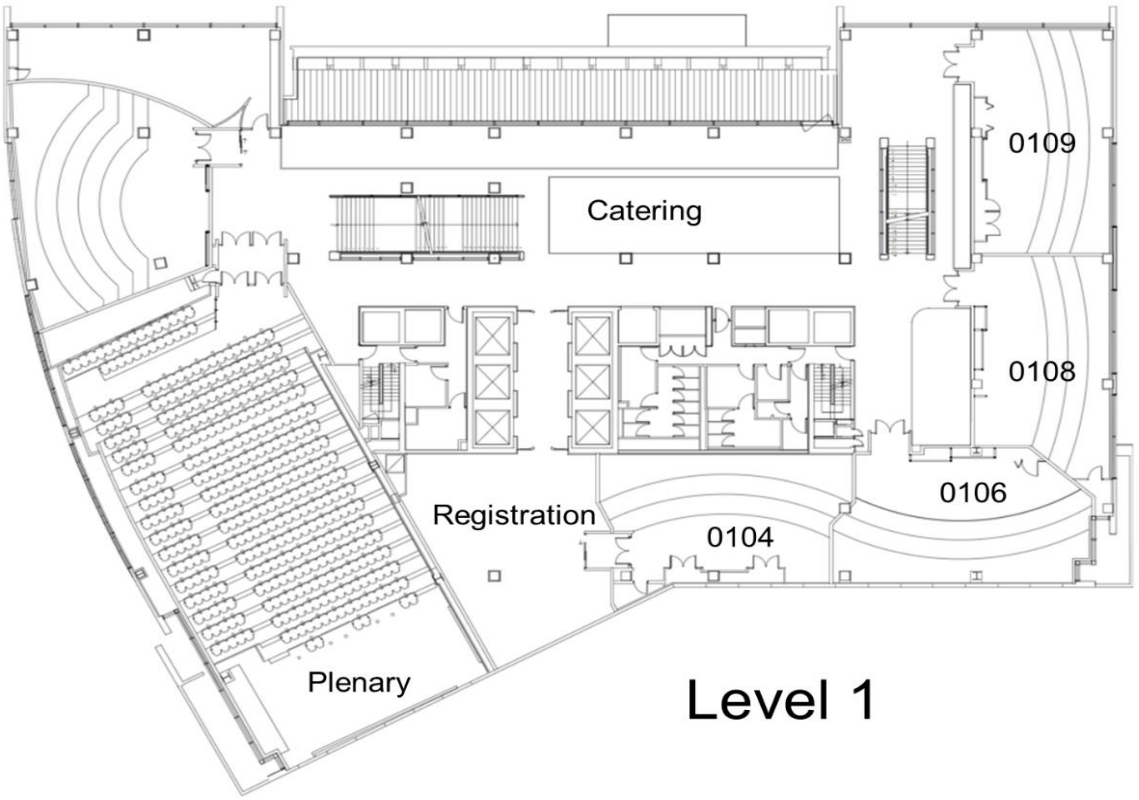

Level 1

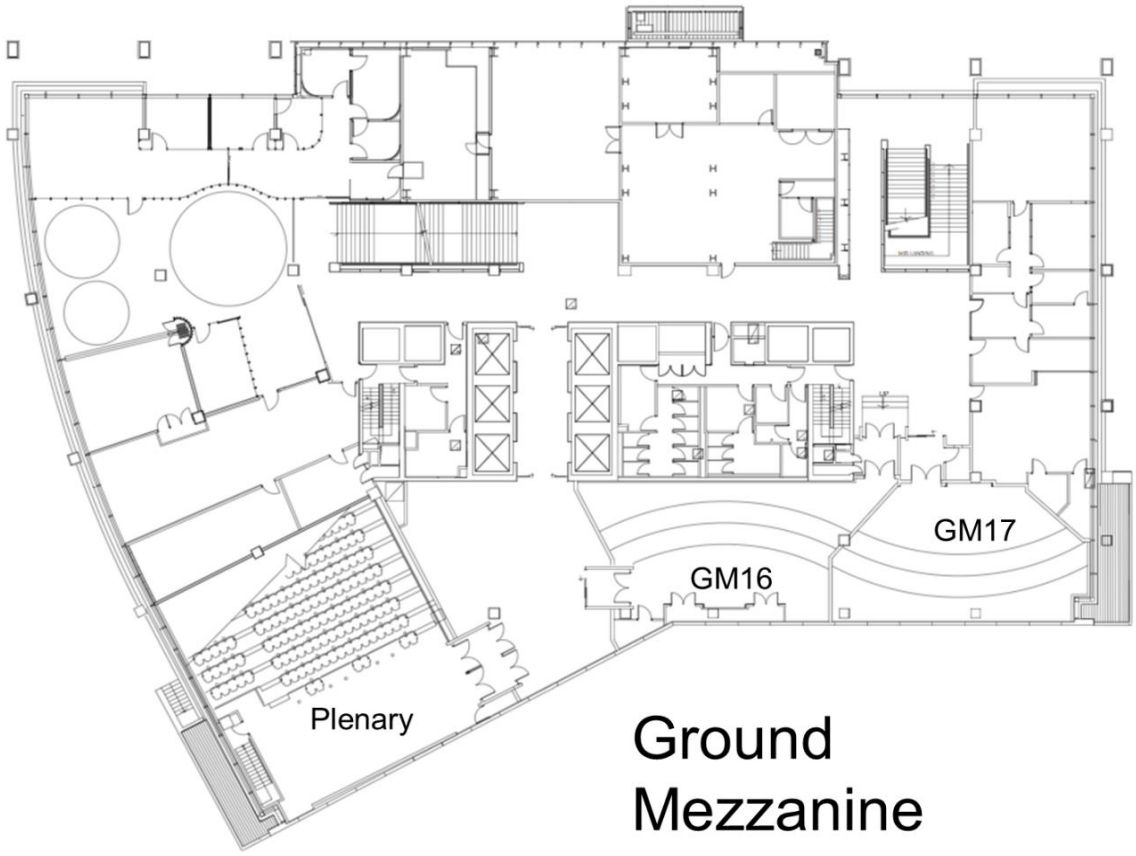

Ground Mezzanine

General Information

Registration

Please note that conference bags and name tags can be collected from the registration desk on Level 1 from 1:30pm on Thursday 9 April. Name tags will be organized alphabetically by last name. Your name tag serves as your admission ticket for all sessions and Welcome Reception and Conference Dinner.

Sessions

1. **Parallel sessions** are 100 minutes long and accommodate up to four presenters. Each presenter is allocated 25 minutes for the presentation, including time for Q&A and preparation time. (20 minutes presentation + 5 minutes for Q&A and switching presentation)

2. If there are fewer than four presenters in your session, please do not extend presentations to fill the extra time.

3. **The session chair**, who is the last speaker of the session, is responsible for time keeping and ensuring that each presenter stays within their allocated time.

4. **Before the session begins**, all presenters must upload their presentation slides (PDF or PowerPoint format) to the computer in the session room.

5. **Wi-Fi**: Conference Wi-Fi is available. Please use the below details:
Username: apesa1
Password: 1Z@ykP

6. **Program** subject to change. Please look at the date stamp or go online for the most current version.

ESA Code of Conduct

The ESA condemns harassment, abuse of power and all forms of discrimination based on gender, ethnicity, race, sexual orientation, religious affiliation, age, national origin, disability status, or any criteria unrelated to scientific matters. It does not tolerate them from the authorities of the association, from its members or from anyone attending its conferences and workshops or engaged in the publication process in the journals of the society.

As defined by the American Economic Association, unacceptable behavior includes, but is not limited to:

- solicitation of emotional or physical intimacy despite expressions or indications that it is unwelcome;
- solicitation of emotional or physical intimacy accompanied by real or implied threat of personal or professional harm;
- intentionally intimidating, threatening, harassing, or abusive actions or remarks (both spoken and in other media);
- prejudicial actions or comments that undermine the principles of equal opportunity, fair treatment, or free academic exchange.

None of these restrictions are intended to stifle the open exchange and discussion of ideas; instead we aim to ensure that all members of the association can participate fully in such discussion.

When members of ESA become aware of professional misconduct, harassment, discrimination, lewdness, or any form of unacceptable behavior in the context of events or activities organized or endorsed by the ESA, they may report it to any ESA points of contact.

https://www.economicscience.org

Plenary Talks

Self-stereotyping: Measurement, Impacts, and New Directions

Katherine B. Coffman
Piramal Professor of Business Administration
Harvard Business School, Harvard University

Date: Thursday, 9 April
Time: 14:00-15:00
Location: David P. Derham Theatre (GM15)

Rationally Selected Utility – A New Theory of Choice

Agnieszka Tymula
Professor
School of Economics, University of Sydney

Date: Friday, 10 April
Time: 16:30-17:30
Location: David P. Derham Theatre (GM15)

Noisy Cognition and Experimental Economics

Ryan Oprea
Professor
Haas School of Business, University of California, Berkeley

Date: Saturday, 11 April
Time: 16:30-17:30
Location: David P. Derham Theatre (GM15)

ESA Mentoring Panels

Mentoring Panel 1: Experimental Economics in the Age of AI

Date and time: Friday 10 April from 13:00-14:00 (during lunch)
Location: David P. Derham Theatre (GM15)

Panelists:

- Katherine B. Coffman (Harvard University)
- David Cooper (University of Iowa)
- Juanjuan Meng (Peking University)

Chair: Keaton Ellis (Monash)

Abstract: This panel will discuss how AI is reshaping behavioral and experimental economics. The goal is to highlight methodological challenges and opportunities created by AI tools for experimental economics. Guiding questions for discussion include: (1) How is AI changing experimental design, data collection, and analysis; and (2) What are the opportunities and challenges for maintaining experimental control and validity with the advent of easily accessible AI tools.

Mentoring Panel 2: Pre-registration, Pre-analysis Plans, and Replicability

Date and time: Saturday 11 April from 13:00-14:00 (during lunch)
Location: David P. Derham Theatre (GM15)

Panelists:

- Tim Cason (Purdue)
- Lucas Coffman (Boston College)
- Ryan Oprea (UC Berkeley)
- Agnieszka Tymula (U of Sydney)

Chair: Maria Recalde (U of Melbourne)

Abstract: This panel will share views on pre-registration and pre-analysis plans in the current research landscape. Guiding questions for discussion include: (1) How norms in the field are changing; (2) How to deal with deviations from pre-registered plans, and (3) What best practices panelists would recommend young scholars including the level of detail needed in pre-registrations and PAPs.

Program Overview

	Time	Event
Day 1: Thursday 9 April	13:30-14:00	Registration
	14:00-15:00	**Plenary: Katherine B. Coffman**
	15:00-15:30	Afternoon Break
	15:30-17:10	Parallel Sessions 1
	17:10-17:30	Mini Break
	17:30-19:30	Welcome Reception
Day 2: Friday 10 April	08:30-09:00	Arrival Tea & Coffee
	09:00-10:40	Parallel Sessions 2
	10:40-11:10	Morning Break
	11:10-12:50	Parallel Sessions 3
	12:50-14:20	Lunch Break
	13:00-14:00	**Mentoring Panel 1**
	14:20-16:00	Parallel Sessions 4
	16:00-16:30	Afternoon Break
	16:30-17:30	**Plenary: Agnieszka Tymula**
Day 3: Saturday 11 April	08:30-09:00	Arrival Tea & Coffee
	09:00-10:40	Parallel Sessions 5
	10:40-11:10	Morning Break
	11:10-12:50	Parallel Sessions 6
	12:50-14:20	Lunch Break
	13:00-14:00	**Mentoring Panel 2**
	14:20-16:00	Parallel Sessions 7
	16:00-16:30	Afternoon Break
	16:30-17:30	**Plenary: Ryan Oprea**
	17:30-19:00	Break
	19:00-22:30	Conference Dinner: State Library

Overview Day 1

	Time	Session Name		Room
Day 1: Thursday 09 April	13:30-14:00	Registration		Level 1 Lobby
	14:00-15:00	**Plenary: Katherine B. Coffman** Harvard Business School, Harvard University *Self-stereotyping: Measurement, Impacts, and New Directions*		David P. Derham Theatre (GM15)
	15:00-15:30	Afternoon Break		Level 1 Lobby
	15:30-17:10	Parallel Sessions 1	Decision-making under Risk	GM16
			AI and Human Behavior 1	GM17
			Belief Updating and Information Processing 1	0104
			Social Preferences 1	0106
			Environmental Economics 1	0108
			Complexity, Attention, and Bounded Rationality 1	0109
	17:10-17:30	Mini Break		
	17:30-19:30	Welcome Reception		Level 10

Overview Day 2

	Time	Session Name		Room
Day 2: Friday 10 April	08:30-09:00	Registration and Arrival Tea & Coffee		Level 1 Lobby
	09:00-10:40	Parallel Sessions 2	Gender Differences and Discrimination 1	GM16
			Market Design 1	GM17
			Behavioral Interventions 1	0104
			Experimental Methodology 1	0106
			Ethics and Decision-Making 1	0108
			Feedback and Behavioral Dynamics	0109
	10:40-11:10	Morning Break		Level 1 Lobby
	11:10-12:50	Parallel Sessions 3	AI and Human Behavior 2	GM16
			Trust, Reciprocity, and Cooperation Dynamics 1	GM17
			Time Preferences and Self-Control	0104
			Education	0106
			Complexity, Attention, and Bounded Rationality 2	0108
			Voting	0109
	12:50-14:20	Lunch Break		Level 1 Lobby
	13:00-14:00	**Mentoring Panel 1: Experimental Economics in the Age of AI** Panelists: Katherine B. Coffman, David Cooper, Juanjuan Meng		David P. Derham Theatre (GM15)
	14:20-16:00	Parallel Sessions 4	Gender Differences and Discrimination 2	GM16
			Financial Markets	GM17
			Social Preferences 2	0104
			Networks, and Peer Effects	0106
			Social Norms	0108
			Information Disclosure and Bayesian Persuasion	0109
	16:00-16:30	Afternoon Break		Level 1 Lobby
	16:30-17:30	**Plenary: Agnieszka Tymula** School of Economics, University of Sydney *Rationally Selected Utility – A New Theory of Choice*		David P. Derham Theatre

Overview Day 3

	Time	Session Name		Room
Day 3: Saturday 11 April	08:30-09:00	Registration and Arrival Tea & Coffee		Level 1 Lobby
	09:00-10:40	Parallel Sessions 5	Trust, Reciprocity, and Cooperation Dynamics 2	GM16
			Gender Differences and Discrimination 3	GM17
			Auctions	0104
			Environmental Economics 2	0106
			Contests	0108
			Experimental Methodology 2	0109
	10:40-11:10	Morning Break		Level 1 Lobby
	11:10-12:50	Parallel Sessions 6	Gender Differences and Discrimination 4	GM16
			Social Preferences 3	GM17
			Market Design 2	0104
			Public Goods	0106
			Behavioral Interventions 2	0108
			Belief Updating and Information Processing 2	0109
	12:50-14:20	Lunch Break		Level 1 Lobby
	13:00-14:00	**Mentoring Panel 2: Pre-registration, Pre-analysis Plans, and Replicability** Panelists: Tim Cason, Lucas Coffman, Ryan Oprea, Agnieszka Tymula		David P. Derham Theatre (GM15)
	14:20-16:00	Parallel Sessions 7	AI and Human Behavior 3	GM16
			Complexity, Attention, and Bounded Rationality 3	GM17
			Environmental Economics 3	0104
			Ethics and Decision-Making 2	0106
			Repeated Games and Cooperation	0108
			Strategic Interaction	0109
	16:00-16:30	Afternoon Break		Level 1 Lobby
	16:30-17:30	**Plenary: Ryan Oprea** Haas School of Business, University of California, Berkeley *Noisy Cognition and Experimental Economics*		David P. Derham Theatre (GM15)
	19:00-22:30	**Conference Dinner**		State Library

Thursday 9th April 2026

Parallel Sessions 1 | 3:30 PM – 5:10 PM

Decision-making under Risk — Room GM16

- **Fear on the Plank? Virtual but more Real!**
  Zhongwen Chen, Monash University
- **Multi-Period Risky Choice with Boundary Conditions**
  Sean Collins, Fordham University
- **Do People Gamble More if They Can Cash-out? Evidence from a Laboratory Experiment**
  Rubayat Sarwar, Adelaide University
- **Should I State or Should I Show? Aligning AI with Human Preferences**
  Keaton Ellis, Monash University (Session Chair)

AI and Human Behavior 1 — Room GM17

- **Bargaining with ChatGPT in Alternating-Offer Games**
  Yuhao Fu, Graduate School of Economic, The University of Osaka
- **The Impact of Generative AI on Economic Persuasion**
  Kristian Lopez Vargas, University of California, Santa Cruz
- **Does Generative AI enhance group efficiency: Evidence from a threshold public good game**
  Yuwen Zhou, Nanyang Technological University
- **Can AI Advisors Solve Social Dilemmas? Evidence from Threshold Public Goods Games**
  Tim Cason, Purdue University (Session Chair)

Belief Updating and Information Processing 1 — Room 0104

- **How do Individuals React to Information in Naturalistic Settings? Rational Decision Making in Elite Sports**
  Scott Dickinson, University of Exeter
- **Uncertainty in Beliefs: an Experiment**
  Wanying Huang, Monash University
- **The Information-Computation Gap: Computational Complexity Predicts Performance in Beliefs Updating**
  Zhongyu Xu, The University of Melbourne
- **Experiencing carbon pricing**
  Stefano Carattini, Georgia State University (Session Chair)

Social Preferences 1 — Room 0106

- **Social Consideration in Simple and Complex Decision Problems**
  Sen Geng, Xiamen University
- **Cultural Norms and Economic Decision-Making: An Experimental Study on Saving and Giving Behaviors in Thailand**
  Lerson Kirasamuthranon, King Mongkut's University of Technology North Bangkok
- **Multi-Step Solicitations for Charitable Donations**
  Maroš Servátka, Macquarie Business School
- **Real-time social information and charitable giving: An experimental investigation**
  Ronald Peeters, University of Otago (Session Chair)

Environmental Economics 1 — Room 0108

- **Information, Incentives, and Goal-Setting: A Field Experiment in Water Usage**
  Anthony van den Berg, University of Melbourne
- **Attention and Green Delivery**
  Ruge Zhang, Nanyang Technological University
- **Deep Industrial Decarbonization: Theory and Evidence from South Korea**
  Emmanuel Murray Leclair, University of Melbourne (Session Chair)

Complexity, Attention, and Bounded Rationality 1 — Room 0109

- **Revealing Preferences and Choice Procedures**
  Ahrash Dianat, University of Essex
- **Healthy ageing in a complex world: decision complexity and basic cognition explain deficits in decision-making capacity**
  Karlo Doroc, Centre for Brain, Mind and Markets, University of Melbourne
- **Are Decisions Under Risk Really Decisions Under Complexity?**
  Zenghui Liu, Shandong University
- **Costly Information and Beliefs: An Experiment**
  Jason Tayawa, UNSW (Session Chair)

Friday 10th April 2026

Parallel Sessions 2 | 9:00 AM – 10:40 AM

Gender Differences and Discrimination 1 — Room GM16

- **A game of status: gender asymmetry in loss aversion**
  Yaxin Liu, University of New South Wales
- **Strategic Misreporting in Wage Negotiations: Cheap Talk and Outside Obligations**
  Jennifer Pate, Loyola Marymount University
- **Differential listening in diverse teams**
  Maria Recalde, The University of Melbourne (Session Chair)

Market Design 1 — Room GM17

- **Signalling in Matching Markets**
  Gloria Chen, The University of Melbourne
- **Matching with Batches**
  Pablo Guillen, The University of Sydney
- **'Knowing more' and 'Step by step': Correlation neglect under preference uncertainty in school choice**
  Ming Jiang, Shanghai Jiao Tong University
- **Matching by Characteristics**
  Siqi Pan, University of Melbourne (Session Chair)

Behavioral Interventions 1 — Room 0104

- **How Information Shapes Retirement Saving: Experimental Evidence from China's Private Pension Account**
  Hanlin Lou, UNSW Sydney
- **Attention and Personalization in Promoting Retirement Savings: Evidence from Two Field Experiments**
  Zeyang Chen, Renmin University of China
- **Behavioral Personalization and Algorithmic Explainability: Evidence from a Field Experiment on Robo-Advising**
  Juanjuan Meng, Peking University (Session Chair)

Experimental Methodology 1 — Room 0106

- **Choice Bracketing and Hedging under Random Incentive System**
  Nobuyuki Hanaki, Institute of Social and Economic Research, Osaka University
- **Obvious Monotonicity as a Test of Behavioral Incentive Compatibility**
Jin-yeong Sohn, Jeonbuk National University
- **Dismantling regret: testing the links between self-reported, behavioural & post-decisional measures of regret**
  Alexander Svenson, The School of Economics, The University of Sydney
- **Identifying Nontransitive Preferences**
  Michele Garagnani, University of Melbourne (Session Chair)

Ethics and Decision-Making 1 — Room 0108

- **Intuition, Reflection, and the Roots of Utilitarian Decision-Making**
  Ozan Isler, The University of Queensland
- **Income and Intrinsic Honesty around the World**
  Moritz Janas, University of Gothenburg
- **I would not lie for you**
  Tanvir Nabi Khan, North South University (Session Chair)

Feedback and Behavioral Dynamics — Room 0109

- **An Experiment on Reflective Prompts and Self-Efficacy in Risky Decision Making to Reduce Dynamic Inconsistency**
  Sol Chung, The University of Sydney
- **Housing Conditions and Social Mobility Expectations**
  Wenyi Zhang, Nanyang Technological University
- **Dynamic Feedback and Incentive Design: Theory and Experiment**
  Ying Wang, HSE University (Session Chair)

Parallel Sessions 3 | 11:10 AM – 12:50 PM

AI and Human Behavior 2 — Room GM16

- **Delegating in the Age of AI: Preferences for Decision Autonomy**
  Radosveta Ivanova-Stenzel, TU Berlin
- **How Information Shapes the Expression of Moral Judgments: Evidence from AI Use in Higher Education**
  Yujiao Li, University of Technology Sydney
- **Does New Technology Induce Task Mismatch?: Experimental Evidence from Generative AI**
  Junghyun Park, Seoul National University
- **The Elusive Returns to AI Skills: Evidence from a Field Experiment**
  Anastasia Danilov, Humboldt-Universität zu Berlin (Session Chair)

Trust, Reciprocity, and Cooperation Dynamics 1 — Room GM17

- **Once a Competitor, Always a Competitor?**
  Adelson Teh, Cornell University
- **Do Campaign Promises Affect Elected Leaders' Behavior?**
  Muhammad Arslan Iqbal, University of Melbourne
- **Too Big to Innovate**
  Mike Zhiren Wu, Monash University
- **Economic Inequality, Social Capital and the Role of Merit and Luck: An Experimental Study**
  Ananish Chaudhuri, University of Auckland (Session Chair)

Time Preferences and Self-Control — Room 0104

- **I Will Pay Later: Time Inconsistency, Self-Control, and Loan Delinquency**
  Leo Bao, Monash University
- **Time and State in Commitment: Experimental Evidence from Crop Insurance in Uganda**
  Sili Zhang, LMU Munich, Department of Economics
- **When Do You Expect to Be Paid? Expectations Shape Patience**
  Elif Incekara-Hafalir, UTS (Session Chair)

Education — Room 0106

- **Learning from Peers: The Impact of Ability Grouping and Incentives**
  Bhagya Gunawardena, RMIT University
- **AI Policy in Education: A Randomized Controlled Trial**
  Yanlin Wan, Hong Kong University of Science and Technology
- **Group Announcements vs. Individual Messages: Experimental Evidence on Digital School-Home Communication in Rural China**
  Jun Zhao, South China Normal University
- **A Mindset Intervention in Rural China**
  Juliana Goncalves, University of Sydney (Session Chair)

Complexity, Attention, and Bounded Rationality 2 — Room 0108

- **Smart Choice Procedures: how to overcome choice overload with simple choice procedures**
  Jerome Campos, The University of Sydney
- **Gaze and Attribute Weighting in Carbon-Labeled Food Choice: A MaaDDM Approach**
  Ying-Chen Lin, National Taiwan University
- **Heterogeneity and Flexibility in Reference Points**
  Nanyin Yang, University of Sydney
- **Winning and Losing: Sleep and Circadian Rhythm Effects on Risky Decisions in a Gain-Loss Framing Paradigm**
  David Dickinson, Appalachian State University (Session Chair)

Voting — Room 0109

- **Giving a Voice: Increasing Individual Self-Expression to Enhance the Resilience to System Discontent**
  Mathilde Bechdolf, University of Magdeburg
- **Rationality-based Preference Aggregation**
  Syngjoo Choi, Seoul National University
- **Linear and Quadratic Voting in Strata Assembly**
  Panharidh Kun, University of Adelaide
- **The Impact of Prejudice and Stereotypes on Electoral Efficiency: Experimental Evidence**
  Ralph-Christopher Bayer, University of Adelaide (Session Chair)

Parallel Sessions 4 | 2:20 PM – 4:00 PM

Gender Differences and Discrimination 2 — Room GM16

- **Speaking with Confidence: Who Benefits and When? Evidence from Physicians**
  Nisvan Erkal, University of Melbourne
- **Empowering or Entrenching? AI Assistance and Gender Gaps in Competition**
  Xu Zhang, Hong Kong University of Science and Technology (Guangzhou)
- **AI, Gender and Fairness**
  Lingbo Huang, Shandong University
- **Who wants to lead in crises: an experimental study of the Glass Cliff**
  Rebecca Heath, University of Cambridge (Session Chair)

Financial Markets — Room GM17

- **The Impact of Trader Incentives on Stability in Option Markets: A Controlled Market Experiment**
  Hoang Long Nguyen, The University of Melbourne
- **Behavioral Dimensions of IPO Mispricing: Evidence from the Laboratory**
  Jingru Wang, Waseda University
- **Equilibrium Play in Experimental Parimutuel Betting Markets**
  Joshua Miller, University of Melbourne
- **The Impact of Attention Allocation Bias on Stock Market Overreaction: An Experimental Investigation**
  Xiaoxue Zhao, Shanghai International Studies University (Session Chair)

Social Preferences 2 — Room 0104

- **Competing Motivations: When More Incentives Lead To Less Effort**
  Kieran Gibson, University of Queensland
- **AI Helps Until You Say It Does**
  Yue Wang, Shandong university
- **Ethical Bandits: An experimental study on the multi-armed bandit with an externality**
  Daniel Woods, Macquarie University
- **Poking Holes and Adding Points in Dictator Games**
  Cary Deck, University of Alabama (Session Chair)

Networks, and Peer Effects — Room 0106

- **Pricing for Product Awareness in Social Network: An Experimental Study**
  King King Li, The Hang Seng University of Hong Kong
- **Social Learning or Social Pressure? Peer Effects in Parental Investment: Field Experimental Evidence from China**
  Le (Lyla) Zhang, Macquarie Business School
- **Social Networks and the Emergence of Social Norms: An Experiment**
  Yohanes Riyanto, Nanyang Technological University (Session Chair)

Social Norms — Room 0108

- **Norm Nudging Under Delegation**
  Sumaiya Bhura, Monash University
- **Norms Make Investors Socially Responsible**
  Peiran Jiao, Maastricht University
- **Information Pooling in the Household: The Role of Expertise**
  Priyoma Mustafi, Ahmedabad University
- **Gender, Culture, and Norms of Equality**
  Pushkar Maitra, Monash University (Session Chair)

Information Disclosure and Bayesian Persuasion — Room 0109

- **Party-influenced media, independent news and information traps**
  Ravindra Agrawal, Indian Institute of Technology Bombay
- **Dr. Cutoff: How I Learned to Love Preregistration**
  Wei-Cheng Chen, Academia Sinica
- **Credibility in Persuasion: A Laboratory Experiment**
  Yuki Shiomi, The University of Osaka
- **Informed Information Design: An Experiment**
  Ninghua Du, Shanghai University of Finance and Economics (Session Chair)

Saturday 11th April 2026

Parallel Sessions 5 | 9:00 AM – 10:40 AM

Trust, Reciprocity, and Cooperation Dynamics 2 — Room GM16

- **Sequential Investment Mechanism and Gradual Trust: An Experimental Study**
  Yunwen He, Central University of Finance and Economics
- **Communication and the Polarization of Trust Attitudes**
  Zhuokun Liu, Goethe University Frankfurt
- **Reciprocal Trust: Lead with Trust to Inspire Trust**
  Yun Zhang, University of Sydney
- **Personality as a Strategic Signal: Experimental Evidence from Trust and Dictator Games**
  Helena Luo, The University of Melbourne (Session Chair)

Gender Differences and Discrimination 3 — Room GM17

- **Bound by Norms, Seen Through Stereotypes: Meta-Beliefs and Minority Women's Labour Outcomes**
  Elif Incekara Hafalir, University of Technology Sydney
- **Misperception of Gender gaps in Performance and Persistence: The Role of Task Stereotype**
  Emily Tran, The University of Melbourne
- **Experiencing unfair chances reduces investments in ambiguous assets**
  Xueting Wang, RMIT University (Session Chair)

Auctions — Room 0104

- **How to Boost Revenues in First-Price Auctions? The Magic of Disclosing Only Winning Bids from Past Auctions**
  Peter Katuscak, RWTH Aachen University
- **Deterring Collusion with Fines in Auction Experiments**
  Tiffany Tsz Kwan TSE, University of Osaka
- **Information Provision in Private-value Auctions**
  Xinghao Yao, The University of Queensland
- **The Missing 'Loser's Curse': Experimental Evidence on Belief-Based Models in Common-Value Auctions**
  Antonio Rosato, Deakin University (Session Chair)

Environmental Economics 2 — Room 0106

- **Input Competition vs. Pollution: The Effect of Mining on Agriculture in Africa**
  Hosam Ibrahim, University of Minnesota
- **Allocation Rules and Market Power in Carbon Trading: Evidence from Laboratory Experiments**
  Wei James Chen, National Taiwan University
- **Coordinating on an environmental path: theory and experiment**
  Inigo DeJuan-Razquin, University of Technology Sydney
- **An Experimental Comparison of Cap- and Intensity-based Pollution Markets**
  Lana Friesen, The University of Queensland (Session Chair)

Contests — Room 0108

- **Impact of prize structure on strategies in all pay contests: An Experiment**
  Yang Liu, RMIT University
- **Handicapping in Two-Period Contests**
  Qin Wu, RMIT University
- **(Sub)optimality and Remedies of the Majority Rule in Team Contests: Theory and Experiment**
  Changxia Ke, Queensland University of Technology (Session Chair)

Experimental Methodology 2 — Room 0109

- **Quasi-exponential discounting**
  Stephen Cheung, The University of Sydney
- **Efficient and Valid Within-Subject Designs**
  Gwen-Jiro Clochard, ISER - Institute of Social and Economic Research
- **uproot: A Software Framework for Behavioral Experiments**
  Max R. P. Grossmann, The University of Melbourne (Session Chair)

Parallel Sessions 6 | 11:10 AM – 12:50 PM

Gender Differences and Discrimination 4 — Room GM16

- **Revisiting Gender Differences in Volunteering for Non-Promotable Tasks**
  Kalyani Chaudhuri, Ashoka University
- **When Gender Is Seen: How Visibility Shapes Collaboration**
  Sukran Dinc, York University
- **Experience in the Same-Gender Environments and Low-Promotability Tasks**
  Duk Gyoo Kim, Yonsei University (Session Chair)

Social Preferences 3 — Room GM17

- **Political Polarization, Wage Inequality and Preferences for Redistribution**
  Lionel Page, The University of Queensland
- **A Bliss Point Model of Distributional Preferences**
  Moira Gidseg, University of Arizona
- **Zero-Sum Views Reduce Support for Redistribution Across Borders**
  Diego Marino Fages, Durham University
- **Preferences over the Timing of Redistribution Policies**
  Mallory Avery, Monash University (Session Chair)

Market Design 2 — Room 0104

- **Style over Substance in Market Design**
  Lucas Coffman, Boston College
- **Complexity Beyond Incentives: The Critical Role of Reporting Language**
  Manshu Khanna, Peking University
- **Consolidating land through rental swaps: Experimental evidence from Uganda**
  Tom Wilkening, The University of Melbourne (Session Chair)

Public Goods — Room 0106

- **Marital Arrangement and Spousal Cooperation**
  Uzma Afzal, RMIT University
- **When Democratic Choice Meets Identity: Cooperation in a Public Goods Game with Third-Party Punishment**
  Hyoji Kwon, University of Hyogo
- **Government Subsidies under Climate Risk: Provision of Rural Public Goods in the Field**
  Shuwen Li, WU Vienna University of Economics and Business
- **Solving normative conflicts in collective action by promoting redistribution**
  Lata Gangadharan, Monash University (Session Chair)

Behavioral Interventions 2 — Room 0108

- **Informational and Financial Nudges for Sugar Reduction in Non-Standardized SSBs: Field Experimental Evidence from Taiwan**
  Chiu-Lin Huang, National Chi Nan University
- **Does the informational promotion nudge credit card repayment?**
  Simin Tao, University of Technology Sydney
- **Rights and Wrongs: The Effect of Labor Rights Information on Workplace Decisions among Temporary Migrants**
  Monica Beeder, University of Southampton (Session Chair)

Belief Updating and Information Processing 2 — Room 0109

- **Intrinsic preferences in prediction markets: An experiment**
  Senran Lin, Southwestern University of Finance and Economics
- **An Axiomatic Model and Test of Grether (1980) and Bayes' Rule**
  Kenneth Chan, National University of Singapore
- **Information representation and meta-cognition of informativeness**
  ChienHsun Lin, National Taipei University
- **The effect of price signals on effort and productivity in a hard task**
  Hassan Andrabi, Centre for Brain, Mind and Markets (Session Chair)

Parallel Sessions 7 | 2:20 PM – 4:00 PM

AI and Human Behavior 3 — Room GM16

- **Augmenting Collective Intelligence: The Effects of AI Cognitive and Coordination Support on Team Performance**
  Zebang Deng, Nanyang Technological University
- **Testing Digital Health Assistants: An Online Experiment**
  Gyula Seres, National University of Singapore
- **Understanding Effects of Decision Recommendation: Evidence from Large Language Models**
  Fengfei Sun, NUS
- **Automating Job Interviews and Interview Evaluations: A Field Experiment**
  Edwin Ip, University of Exeter (Session Chair)

Complexity, Attention, and Bounded Rationality 3 — Room GM17

- **Expected Effort Allocation Under Task Difficulty and Reward Uncertainty**
  Michelle Lee, The University of Melbourne
- **Complexity and the Use of 'Supposedly Irrelevant' Factors in Decision-Making**
  Ao Wang, National University of Singapore
- **Minimal Intelligence in the Double Auction: Logit-Choice and Reservation Utility**
  Brett Williams, University of New South Wales (Session Chair)

Environmental Economics 3 — Room 0104

- **Framing Carbon Offsets: Moral Spillovers and Decision Context in Sustainable Supply Chains**
  Danielle Kent, The University of Sydney
- **Optional Enforcement Backfires: Experimental Evidence on the Regulation of Negative Externalities**
  Erte Xiao, Monash University
- **Community-Based Solutions to Overcome Collective Action Problems: An Experiment in Rural Vietnam**
  Joseph Vecci, University of Gothenburg
- **Fairness views about the International Distribution of Climate Change Costs**
  Davide Pace, LMU Munich (Session Chair)

Ethics and Decision-Making 2 — Room 0106

- **Paying to Avoid the Spotlight**
  Te Bao, Nanyang Technological University
- **The King's Banquet: Altruistic Wilful Ignorance in the Laboratory**
  Xiaomeng Chen, Monash University
- **Using fNIRS hyperscanning to test neurocognitive mechanisms of financial advisor misconduct in face-to-face interactions**
  Xiaofei Niu, Shandong University
- **Playing Dumb to Look Green?**
  Alice Solda, Emlyon Business School (Session Chair)

Repeated Games and Cooperation — Room 0108

- **Comparative learnings**
  Mia Tam, The University of Melbourne
- **Uncertain Others: How Ambiguity Attitudes Shape Cooperation Across Social Dilemmas**
  Tse-Min Wang, National Taipei University
- **Modeling Strategies in Indefinitely Repeated Prisoner's Dilemma Games**
  David Cooper, University of Iowa (Session Chair)

Strategic Interaction — Room 0109

- **Strategic sophistication and mixed strategies**
  Evan Calford, Australian National University
- **Psychological Pressure: The Case of Penalty Kicks**
  Romain Gauriot, Deakin University
- **Natural Experiment in Fair Division: Armageddon Chess**
  John Wooders, NYU Abu Dhabi (Session Chair)

ABSTRACTS

Fear on the Plank? Virtual but more Real!

Zhongwen Chen, Monash University

This paper uses virtual reality (VR) technology to conduct a fully controlled economic experiment of human decisions such as risk-taking behaviours associated with life and death that could be difficult to achieve in both the standard fields and laboratories. In our experiment, participants virtually stand on the roof of a hundred-storey building. They need to build a plank bridge to reach a monetary reward suspended from a drone on the other side, and then walk through the plank to grab the reward, facing a probability of falling virtually. The planks vary in width and price, with wider planks offering greater security at a higher cost. We compare risk attitudes in VR to those revealed through existing methods such as (context-free) multiple price list, hypothetical questions, etc, and point out limitations of traditional risk preference elicitation methodologies used in economic experiments.

Multi-Period Risky Choice with Boundary Conditions

Sean Collins, Fordham University

We implement a multi-period individual choice problem, where the evolution of earnings across past rounds affects working capital available for the next round. In addition, we impose treatments which reinitialize working capital and/or hold out a continuation value for not exhausting working capital (i.e., going bankrupt). We incorporate the bankruptcy boundary into a dynamic program that allows us to parse the effect of the bankruptcy boundary from other potential predictors of risky choice.

Do People Gamble More if They Can Cash-out? Evidence from a Laboratory Experiment

Rubayat Sarwar

People often place bets based on their risk preferences and an overconfident belief that they can predict game outcomes better than bookmakers. If we eliminate this channel of overconfidence by using objective probabilities and observe different betting behaviors for two equivalent betting products, the underlying psychological factors behind the products will likely be driving these varying decisions. In our study, we conducted laboratory experiments to investigate the underlying behavioral factors that attract sports bettors towards innovative, trendy multi-bets over traditional, non-innovative single bets. Behaviorally, innovation often adds extra thrills for many people, which may lead people to behave differently than they would do in its absence. Thus, we hypothesized that innovative multi-bets might make people gamble more than equivalent simple ones. To examine this, we performed a cross-subject experiment involving three treatment conditions. We planned treatments in order to disentangle the bias effect (misjudgment of compound probability of multies) and the thrill effect between multi-bets and single bets. In Treatment 1, participants made choices to invest in single bets. In Treatment 2, participants made decisions for multi-bets. Here, we provided both the success probabilities for each individual bet in the multi and the compound probability for the entire multi. In Treatment 3, participants made choices for the same multies as in Treatment 2, but we provided only the individual probabilities and withheld the compound probability. Thus, any observed differences in betting stakes between Treatment 1 and 2 suggest that participants get extra thrills for the multi-bets over the single bets. Variations in betting stakes between Treatment 2 and 3 indicate whether participants misjudge compound probabilities when that information is not made explicit. Finally, comparing Treatments 1 and 3 captures the total effect of probability misjudgment and thrill for multi-bets. Each participant made ten investment decisions across ten games, choosing between $0 and $8 per round. We find that on average, individuals exposed to multi-bets without the compound

probability (Treatment 3) spent $0.75 more than the individuals in the equivalent single bets (Treatment 1), as estimated using a Tobit random-effects model (p = 0.06). However, differences between Treatments 1 and 2, as well as between Treatments 2 and 3, were not statistically significant. This result implies that the observed increase in spending on multi-bets cannot be attributed solely to thrill or solely to the probability of misjudgment. When examining mean stakes across treatments-$2.73 in Treatment 1, $2.96 in Treatment 2, and $3.39 in Treatment 3-a consistent pattern emerges. The results point toward a combined psychological mechanism, where both the additional thrill for innovative betting products and the misjudgment of compound probabilities likely contribute to higher spending in multi-bet environments. Thus, our results suggest that innovation in betting products exerts a stronger influence on choice behavior than its absence. Pre-registration: https://www.socialscienceregistry.org/trials/13852

Should I State or Should I Show? Aligning AI with Human Preferences

Keaton Ellis, Monash University

As AI agents become more autonomous, properly aligning their objectives with human preferences becomes increasingly important. We study how effectively an AI agent learns a human principal's preference in choice under risk via stated versus revealed preferences. We conduct an online experiment in which subjects state their preferences through written instructions ("prompts") and reveal them through choices in a series of binary lottery questions ("data"). We find that on average, an AI agent given revealed-preference data predicts subjects' choices more accurately than an AI agent given stated-preference prompts. Further analysis suggests that the gap is driven by subjects' difficulty in translating their own preferences into written instructions. When given a choice between which information source to give to an AI agent, a large portion of subjects fail to select the more informative one. Moreover, when predictions from the two sources conflict, we find that the AI agent aligns more frequently with the prompt, despite its lower accuracy. Overall, these results highlight the revealed preference approach as a powerful mechanism for communicating human preferences to AI agents, but its success depends on careful implementation.

Bargaining with ChatGPT in Alternating-Offer Games

Yuhao Fu, Graduate School of Economic, The University of Osaka

We conducted a lab experiment using an alternating-offer bargaining game to test (i) whether humans behaved differently when facing an GPT opponent versus when facing a human opponent, and (ii) whether human decisions changed when the AI's payoff was linked to the payoffs of other participants. We found that in Stage 1, human proposers' behaviors were not affected by opponent type. However, because the AI responders accepted unfair and even fair offers less often than human responders in Stage 1, human proposers increased their offer size to GPT opponents in the later rounds of Stage 1. Regarding the transfer condition, we found no clear evidence that linking the AI's payoff to other participants' payoffs affected human decisions in all stages. Overall, these findings improve our understanding of role and stage-dependent patterns in human-AI interaction and indicated when AI-related payoff externalities did (and did not) shape strategic choices.

The Impact of Generative AI on Economic Persuasion

Kristian Lopez Vargas, University of California, Santa Cruz

The use of generative AI tools to enhance communication is rapidly expanding globally across workplaces, markets, and everyday online interactions. Yet it remains unclear whether these tools alter the form of communication or also yield systematic material benefits for their users in

interpersonal interactions, and for whom. This question is particularly salient in settings where persuasive communication shapes distributive outcomes, such as bargaining, hiring, or charitable giving, and where AI may either level the playing field or widen existing disparities in communication skills. This paper examines how access to generative artificial intelligence (AI) affects economic outcomes by augmenting individuals' persuasive capabilities. We conduct a controlled laboratory experiment based on the classic dictator game, in which receivers can send short messages to dictators before allocation decisions. In a within-subject design, each participant is exposed, in random order, to three conditions: no communication, a human-written message, and a human+AI-assisted message. In the AI-assisted condition, receivers interact with a frontier language model through the experimental interface, subject to light message-length and content constraints. The experiment, implemented in oTree and deployed on Prolific, is designed to isolate the causal effect of AI access on material gains via persuasion. Building on recent evidence that large language models can match or exceed human persuasive performance and enable scalable personalization, this ongoing study quantifies the impact of AI-mediated persuasion on distributive behavior and investigates whether such assistance acts as an equalizer or instead amplifies pre-existing inequalities in communication-driven outcomes.

Does Generative AI enhance group efficiency: Evidence from a threshold public good game

Yuwen Zhou, Nanyang Technological University

Generative artificial intelligence (GenAI) is increasingly used in many areas, with people relying on its advice or even letting it make decisions for them. This study looks at how GenAI influences behavior in an infinitely repeated threshold public goods game. It focuses on whether GenAI is more effective than other coordination enhancing tools, like pre-game communication, in encouraging turn-taking-a form of cooperation over time. In the game, participants are given tokens with different values and must decide each round whether to contribute a token. Turn-taking can ensure fairness by sharing the burden, but it reduces efficiency because the most cost-effective strategy is for low-cost players to always contribute. This creates a trade-off between fairness and efficiency. The study first examines whether GenAI favors fairness over efficiency when giving advice or making decisions for players. It then investigates how human players respond to the advice or decisions provided by GenAI. By exploring these questions, we shed light on how GenAI can shape human decision-making and offer insights for designing policies that balance fairness and efficiency.

Can AI Advisors Solve Social Dilemmas? Evidence from Threshold Public Goods Games

Tim Cason, Purdue University

This lab experiment studies how human-AI interaction affects strategic behavior in a repeated social dilemma. Subjects in groups of four choose each round how much to allocate to a group project. The group project is funded if total allocations reach a specified threshold. When the group project is funded, subjects receive a heterogeneous payment (randomly drawn without replacement before their decision) that represents their public good return. When the threshold is not met, subjects receive their allocation back plus a potential refund bonus (a fraction of their allocation) in some treatments. Treatments also vary whether human subjects have access to a generative AI expert advisor (OpenAI's GPT-5), and outcomes are also compared to a benchmark where AI experts make allocations directly. Preliminary results indicate that AI experts' lack of concern for relative earnings in this heterogeneous environment promotes coordination and public goods provision.

How do Individuals React to Information in Naturalistic Settings? Rational Decision Making in Elite Sports

Scott Dickinson, University of Exeter

In many experimental settings, individuals appear to either underreact or overreact to information when making decisions about the future. Less is known about how individuals react to information when making decisions in naturalistic settings. In this paper, I study the ability of individual decision makers to incorporate information about past outcomes when making judgements about future outcomes in naturalistic settings using data on elite sports. In the setting under observation, fans make hundreds of decisions relating to the future performance of elite athletes. Using variation in outcomes between athletes who took comparable actions, I show that fans correctly incorporate persistence in the data generating process into their predictions. Not only do fans react to information in the correct direction, but they also appear to react with the correct magnitude. This is consistent with fans rationally reacting to information.

UNCERTAINTY IN BELIEFS: AN EXPERIMENT

Wanying Huang, Monash University

We experimentally study individuals' subjective uncertainties about their own beliefs. Subjects are incentivized to truthfully report their distribution of beliefs about the empirical frequency of an i.i.d. process. From a frequentist perspective, this empirical frequency represents a subject's first-order belief about an unknown state; hence, the elicited distribution can be interpreted as their second-order belief uncertainty over their own first-order beliefs. We vary both the prior and the signal environment to trace how belief uncertainty forms and updates. We find that (i) belief uncertainty is prevalent, even at the prior stage; (ii) belief uncertainty responds to information, though more precise signals do not eliminate it; and (iii) subjects with greater posterior uncertainty are less likely to choose a risky bet over a safe option, even after controlling for their risk and ambiguity attitudes. This implies that without accounting for belief uncertainty, standard economic models would overpredict risky behavior.

The Information-Computation Gap: Computational Complexity Predicts Performance in Beliefs Updating

Zhongyu Xu, The University of Melbourne

In belief updating, people generally exhibit behaviour that deviates from optimal inference. In this work, drawing on computational complexity theory from computer science, we offer a novel method to predict performance in beliefs updating using computational requirements. Specifically, we propose a task-independent metric that quantifies the computational resources needed for belief updating. We then validate this metric by showing that it makes unique predictions regarding both decision quality and the time participants spend on a belief-updating task. These findings provide a prescriptive theory that explains existing descriptive accounts of deviations from rationality in belief updating, underscoring the importance of incorporating computational constraints when evaluating decision-making performance.

Experiencing carbon pricing

Stefano Carattini, Georgia State University

Many socially desirable policies are not implemented because of their ex-ante unpopularity, but this unpopularity may be overcome through experience with the policy. In this paper, we examine how opposition to carbon pricing in the state of Washington turned into support after voters experienced a

cap-and-trade policy with revenues earmarked for environmental purposes -- "cap-and-invest." Analyzing voting behavior at the census block group level, we observe that support varies by political affiliation as expected, but experience consistently increases support across the board. Using a proprietary survey, we further show that the increase in support among voters in Washington state is specific to the cap-and-invest policy they experienced; support for carbon pricing or climate policies more generally remained unchanged.

Social Consideration in Simple and Complex Decision Problems

Sen Geng, Xiamen University

We develop a rational-choice framework that can explain why behavior appears selfish in some contexts and prosocial in others, even if the underlying preferences are stable. In our model, decision makers choose a decision procedure, which is formalized by the degree of their social consideration and the level of fine-tuning of the procedure. We show that as the complexity of the decision problem (either individual decision making problem or strategic game) increases, the expression of social preferences (active social consideration) is endogenously suppressed. When complexity is sufficiently high, decision makers revert to a simple (default) strategy and optimize only for self-interest. Our experimental results suggest that the disappearance of prosocial behavior in complex environments is not a change in preference, but a rational adaptation to the cognitive costs of deliberation.

Cultural Norms and Economic Decision-Making: An Experimental Study on Saving and Giving Behaviors in Thailand

Lerson Kirasamuthranon, King Mongkut s University of Technology North Bangkok

Thailand displays high levels of prosocial behaviour across different socioeconomic groups, yet the driving forces behind these behaviours remain insufficiently understood. While findings in behavioral economics demonstrate that people deviate from the rational choice theory, they often overlook culturally embedded motivations. Andreoni's "Warm Glow Theory" captures generic motivations for prosocial behavior but still fails to explain the distinct norms that shape them. As such, this study aims to investigate how two central Thai norms - namjai (heartfelt generosity) and karma - shape prosocial behavior, immediate reward preference, and future-oriented savings. An a priori power analysis (alpha = 0.05, power = 0.80, and d = 0.40) determined a minimum sample size of 300. Participants were recruited from a local university and nearby community clinics and were randomly assigned to one of three conditions: karma prime, namjai prime, or the neutral control. After completing a cultural salient prime, participants played a modified dictator game to allocate 1000 THB between charitable donation, immediate reward, or a delayed payment savings option that would yield 20% interest after two weeks. Demographic covariates include age, gender, income, religiosity, and past donations. A manipulation check asked participants to identify the priming condition and self-rate their post - prime willingness to help. Then, the data collected was analyzed through OLS Regression models, ANCOVA, and MANOVA to compare group means, interaction effects, and to determine each variable's significance. Results suggest that both experimental primes, namjai and karma, increased charitable giving relative to the control, with the karma prime displaying a greater effect. Participants in the namjai prime were shown to prefer savings over immediate reward, displaying reduced temporal discounting. This research provides causal evidence that norms actively shape economic preferences and prosocial behaviors. Findings are useful in creating culturally aligned financial literacy programs, designing targeted saving schemes, and creating charitable giving campaigns. Crucially, this study highlights the value of integrating cultural psychology into behavioral economics.

Multi-Step Solicitations for Charitable Donations

Maroš Servátka, Macquarie Business School

Solicitations for charitable donations are often framed as a multi-step ask; people are first asked whether they are willing to make a donation and, if yes, are then immediately provided with relevant details about the potential recipients and corresponding time commitment. We hypothesize that uncertainty about both recipient details and transaction costs, in addition to commitment aversion, makes multi-step charitable solicitations less effective than one-step asks. We test our hypotheses using an online experiment. Our results confirm that a multi-step ask elicits a lower donation rate than a one-step ask. Reducing uncertainty about transaction costs and controlling for commitment aversion are ineffective at increasing the multi-step ask donation rate. However, eliminating uncertainty about the potential recipients restores the donation rate to the one-step ask level. Implications are discussed.

Real-time social information and charitable giving: An experimental investigation

Ronald Peeters, University of Otago

We design a novel laboratory game to analyze how real-time information on others' donations affects charitable giving. Subjects interact in a crowdfunding platform in which social information is posted and continuously updated in real-time. Some groups receive information on individual-level donations by all group members, while others only know aggregate donations by the group. These information treatments are compared to one another and to a control with no social information. We find that both the distribution of donations across participants and the dynamics of donations differ when varying the social information that we provide. The more social information that is provided the lower the within-group dispersion becomes. Yet, interestingly, we find no differences at the aggregate level in the amounts raised between treatments. Our findings suggest that social information abets within-group coordination on donation amounts.

Information, Incentives, and Goal-Setting: A Field Experiment in Water Usage

Anthony van den Berg, University of Melbourne

Can gamification motivate resource conservation? We explore this question by encouraging households to download an app that uses smart meter data to set weekly water usage goals and provide daily feedback on water consumption. The intervention reduced average daily water use by 8.8%. We found no bunching around goals but observed distributional shifts in daily water usage with systemic goal overshooting. These results, combined with complementary survey data, illustrate how household optimization frictions influence behavior in goal-based programs. At scale, the intervention is cost-effective, complementing and potentially lessening the need for costly water capacity investments like desalination plants.

Attention and Green Delivery

Ruge Zhang, Nanyang Technological University

Household participation is essential for the low-carbon transition, yet everyday sustainable actions remain fragile and context-dependent. This paper studies how limited attention and cognitive scarcity shape a simple but meaningful pro-environmental choice: whether to select a consolidated 'green delivery' option in online grocery purchases. We combine transaction-level data from RedMart, Singapore's largest online supermarket, with high-frequency environmental and policy data to show that stress and salience jointly drive household contributions to climate mitigation. We use transaction-

level data between March 2023 and March 2025, in which at least one 'green delivery' slot was available. At checkout, customers see a menu of delivery time slots, some of which are labelled as green because they allow route consolidation and lower emissions. These slots carry no price discount or loyalty incentive and differ only by an environmental label and, in some cases, lower scheduling flexibility. For each order, we observe the full menu, the chosen slot, and detailed basket characteristics, which we merge with district-level wet-bulb temperature, air-quality (PSI), electricity temporary price-cap events (TPC), and a hand-coded daily panel of nationwide environmental campaigns. We find that households are less likely to adopt green delivery under environmental pressures: higher air pollution, elevated wet-bulb temperatures, and electricity price-cap events all reduce uptake. By contrast, nationwide sustainability campaigns sharply increase participation, functioning as salience shocks that redirect attention toward environmental goals. However, this salience effect is significantly attenuated under high environmental pressures. To interpret the evidence, we develop a simple behavioral inattention model in which households allocate scarce attention between immediate convenience/cost and diffuse environmental benefits. Physical and financial stress reduces the attention weight on future environmental payoffs, while campaigns increase it, with salience effects attenuated when stress is high. Our results highlight attention scarcity as a central constraint on sustainable consumption and uncover an additional cost of environmental degradation: heat and pollution not only damage welfare directly but also weaken behavioral support for climate policy.

Deep Industrial Decarbonization: Theory and Evidence from South Korea

Emmanuel Murray Leclair, University of Melbourne

We examine deep industrial decarbonization, the process of altering industrial production processes to remove fossil fuels, using new theory and empirical evidence from South Korea. South Korea has instituted a series of national and district-specific policies aimed at curbing industrial emissions. These include cap-and-trade systems, various command-and-control regulations, and subsidies for the adoption of green technologies. These policies correlate with substantial green technology adoption in energy-intensive industries like cement and steel. Drawing on unique establishment-level data on energy use and production, we document a dramatic rise in electrification: between 2016 and 2022, the share of plants relying solely on electricity almost doubled (from 35% to 65%). We develop and estimate a rich dynamic spatial general equilibrium model where firms can pay fixed switching costs to adopt or drop fuels. In our model, national and district-specific regulations affect local fuel prices while subsidies reduce the fixed costs of adopting cleaner fuels. We find that subsidies targeting small and medium-sized enterprises drive most electrification, rather than environmental regulations. We then quantify the welfare effects of these policies across space and over time. A key finding is that, despite widespread electrification, greenhouse gas reductions remain limited because the largest polluters, who account for most emissions, are less likely to become fully electrified.

Revealing Preferences and Choice Procedures

Ahrash Dianat, University of Essex

Classical economic models assume that the decision-maker (DM) acts in a way to maximize her preferences. However, the economics literature is largely silent about the procedure that the DM uses to maximize her preferences. We argue that observed inconsistencies can be attributed to one of two sources: either a failure to perform the preference maximization procedure or a failure to hold classically rational preferences. In this paper, we disentangle observed choices along the procedural and preference dimensions. To do so, we conduct a laboratory experiment to collect data on the potential procedures that DMs use to implement the maximization of their preferences. We find that,

for around half of subjects who fail to satisfy rationality, the failure stems from the choice procedure and not from their ability to perform binary comparisons. We also present evidence that the heuristics used by subjects are distinct from standard search models.

Healthy ageing in a complex world: decision complexity and basic cognition explain deficits in decision-making capacity

Karlo Doroc, Centre for Brain, Mind and Markets, University of Melbourne

Later life involves demanding and consequential decisions, such as providing consent to risky medical treatments or deciding how to manage one's retirement income. Existing research has documented differences in decision outcomes and basic cognitive abilities between younger and older adults, but these studies struggle to disentangle whether these differences stem from motivational changes or a decline in decision-making capacity. For instance, it is often not possible to determine whether changes in choice reflect a change in underlying preferences or instead a change in the capacity to identify the preferred option. To address this gap, we conducted a large (n=357), pre-registered, and incentive compatible experiment administering a battery of decision-making and cognitive tasks to younger (18-30) and older (65+) adults. Decision-making capacity was measured with the knapsack optimisation task (KOT), an objective resource allocation task, which allows us to precisely modulate the computational demands of individual decisions (decision complexity). With theoretical ties to the standard consumer choice problem, the KOT also allows for an objective, preference-free measure of decision quality. In addition, five tasks assessed basic cognitive abilities, including measures of executive functioning, fluid reasoning, and processing speed. On average, older adults made significantly lower quality decisions in the KOT and were more overconfident, despite spending more time. These differences in decision-making capacity were primarily explained by performance on tasks of basic cognitive abilities, but not age, education, or motivation. Surprisingly, the largest differences in decision time and decision quality occurred on less computationally demanding decisions - suggesting that the performance gap between younger and older adults narrows as complexity / computational demands increase. Importantly, we observe significant overlap in cognitive abilities and decision-making capacities between age groups. These results challenge age-based assumptions by showing that decision-making capacity is not inherently determined by years lived, but by cognitive health - a finding with direct implications for policy frameworks that currently rely on chronological age to govern financial, employment, and medical decision rights.

Are Decisions Under Risk Really Decisions Under Complexity?

Zenghui Liu, Shandong University

We investigate the replicability of Oprea's (2024) findings, with particular emphasis on a controlled laboratory setting. We find that his main findings do not replicate in the laboratory, in part due to the procedure of reading instructions aloud. In the online setting, implementing this procedure yields divergent results, whereas removing it reproduces his main results. A plausible explanation for the emergence of Oprea's (2024) findings in the online environment is experimental confusion. We argue that there is no compelling evidence that classical lottery anomalies such as probability weighting and loss aversion are not special phenomena of risk.

Costly Information and Beliefs: An Experiment

Jason Tayawa, UNSW

We design an experiment to document the validity of rational inattention models. Our design evaluates these models in direct (choosing posterior beliefs) and natural (choosing sample size) frames. The former frame implements a decision variable closest to the model's primitives, while the latter is more natural for experimental settings. By observing the acquired information in either frame, we can determine which class of rational inattention models aligns with the decision-making processes in the lab. These frames provide insights into the validity of these classes in rational inattention models, which are often tested solely based on choice data without observing the chosen level of informativeness. Assessing the consistency of their behavior between the two frames allows us to know whether it is without loss of generality that experimental designs may adopt an indirect yet natural decision variable and still make use of models that are ultimately defined on random posterior distributions.

A game of status: gender asymmetry in loss aversion

Yaxin Liu, University of New South Wales

This study examines gender differences in loss aversion, focusing on the fear of losing status. Using data from more than 14,000 Grand Slam matches in professional tennis, I introduce a novel measure-the Net Tournament Balance (NTB), which captures the number of points a player must earn to maintain their current total ranking points. To identify causal effects of placement in the “loss” domain on performance, I employ a regression discontinuity design. The results reveal a striking gender difference: male players exhibit strong loss aversion to status. They are significantly more likely to win matches and adopt a more aggressive playing style when they face the threat of falling in ranking points, consistent with the predictions of prospect theory. In contrast, female players show no comparable behavioral adjustment, highlighting a clear gender asymmetry in responses to potential status loss. Additional analysis suggests that culture, social norms, and identity may help explain why male players are more sensitive to status loss.

Strategic Misreporting in Wage Negotiations: Cheap Talk and Outside Obligations

Jennifer Pate, Loyola Marymount University

We examine wage negotiations using an ultimatum game with a real-effort task, introducing an "outside obligation" where employees must allocate 0 to 50 percent of their earnings to a third party. Employees can report their obligation via cheap talk and propose a wage split, while employers make counteroffers. If accepted, the employee completes the task of generating the surplus; if rejected, both parties complete the task for a small payment. We find frequent misreporting, with 80 percent of employees over-reporting when their obligation is zero. Men exaggerate their obligations more often than women, who also overstate, but to a lesser extent. Employers increase offers when they believe the reported obligation, but when skeptical, they make offers similar to those in settings with no information. However, employers fail to anticipate men's greater misreporting, leading them to respond in ways that unintentionally contribute to a gender wage gap, despite exhibiting no explicit gender bias.

Differential listening in diverse teams

Maria Recalde, The University of Melbourne

Conflict and communication difficulties are commonly cited as reasons for why assembling diverse teams does not generate its expected benefits. This paper studies a possible mechanism underlying this result: the voices of women and minoritized racial/ethnic groups are disproportionately discounted. In an online experiment, we study differential listening in diverse teams in a hiring context,

where committee leaders aggregate input from committee members when selecting a wage to offer a job applicant and no differences in quality of input exist. We find clear evidence of differential listening by race. Committee leaders are less influenced by input from Black committee members than they are by white committee members. The impact of gender is mixed: white men are the most influential and Black men the least; yet the input of all women and Black men is discounted relative to that of white men. We also find that only some leaders engage in differential listening; specifically, white leaders, both women and men. When examining mechanisms, we find that gender differences in confidence explain the lower influence of women, while beliefs about the quality of committee member input in part drive the differential listening of Black men.

Signalling in Matching Markets

Gloria Chen, The University of Melbourne

In many real-life matching markets, such as residency matches and job applications, employers often send signals of interest to candidates, and these signals frequently influence outcomes. However, under the Deferred Acceptance mechanism, standard matching theory predicts that informing a candidate they are favoured should not affect how they submit their preferences, since truth-telling remains a weakly dominant strategy. To address this discrepancy between theoretical predictions and observed behaviour, we develop a behavioural model in which candidates exhibit rank-dependent preferences and derive additional utility from receiving a positive signal. This framework explains why candidates may respond to signals and how firms can strategically use signalling to influence preference submissions in their favour, thereby improving their own match outcomes. We test the predictions of our behavioural model through a laboratory experiment. Our results suggest that firms do send signals strategically to enhance their match outcomes. Candidates, in turn, respond by ranking the signalling firm higher upon receiving the signal. The firm's strategic signalling, coupled with candidates' behavioural responses, causes the market outcome to shift from one that favours firms to one that benefits candidates.

Matching with Batches

Pablo Guillen, The University of Sydney

Most real-life centralised university admissions systems ask applicants to submit a list shorter than the number of options available to them, as lists would be unwieldy long otherwise. Problematically, constrains on list length results on any matching mechanism to become non-strategy-proof. That is, the applicants' incentive to rank their courses truthfully is lost. Under these circumstances, it becomes impossible for the clearinghouse to provide simple and sensible advice to applicants. We propose the use of batches, asking participants to send short lists sequentially until the market clears, to recover strategy-proofness. Our approach is particularly useful when applied to systems based on cut-off entry marks rather than quotas, such as the one used by the Universities Admissions Centre (UAC) in Australia. We test the efficiency of batching in an individual decision-making experiment, comparing the batched mechanism to both the status quo (a constrained list) and the use of an unconstrained list, with and without advice. We find a significant and sizable increase in realised payoffs in the batched treatments relative to the status quo, but no effect of advice.

'Knowing more' and 'Step by step': Correlation neglect under preference uncertainty in school choice

Ming Jiang, Shanghai Jiao Tong University

We conduct a laboratory school-choice experiment in which admission cut-offs arise endogenously, testing correlation neglect (CN) in school choice settings. Relative to Rees-Jones et al. (2024), which fixes thresholds e-ante, our experiment confirms that CN persists under complex environments We evaluate three debiasing tools: (i) a reminder that acceptances are correlated, (ii) personalized admission probabilities, and (iii) an Iterative Deferred Acceptance (IDA) mechanism that reveals first-round results before the next choice. The reminder has no effect; the treatment providing admission probabilities temper overall risk-taking yet leaves CN unchanged; in contrast, switching to the IDA mechanism reduces the impact of correlation neglect. Our findings suggest that dynamic mechanisms are potentially more effective tha informational interventions in addressing correlation neglect in school choice.

Matching by Characteristics

Siqi Pan, The University of Melbourne

Much of the matching literature has focused on algorithms that incentivize participants to submit complete rank-order lists of all acceptable options-often numbering in the hundreds-and relies on these exhaustive submissions to produce efficient and fair allocations. However, the practical implementation of such theoretically desirable mechanisms faces significant challenges. Forming preferences over a large number of options is costly and cognitively demanding, requiring extensive information acquisition. Moreover, ranking all acceptable options can be overwhelming, leading to suboptimal decision-making and exacerbating existing inequalities. We propose an alternative approach that elicits preferences over characteristics rather than over individual objects, thereby reducing the complexity of both preference formation and preference expression. Our study asks: given a fixed matching mechanism, which elicitation method-traditional object-based ranking or characteristics-based elicitation-yields better outcomes? We examine this question through a controlled laboratory experiment framed as an object-allocation problem, in which participants receive objects with varying characteristics based on their submitted preferences. Our results show that characteristics-based elicitation generates higher welfare and is preferred by participants, underscoring its potential as a practical and scalable solution for implementing matching in real-world settings.

How Information Shapes Retirement Saving: Experimental Evidence from China’s Private Pension Account

Hanlin Lou, UNSW Sydney

This study explores how upfront rewards and peer group composition jointly shape motivation, performance, and belief updating in rank-based contests. We conducted a field experiment with 1,835 students in a Chinese high school, assigning participants to one of six experimental groups. Participants were placed in groups of five and completed a two-day vocabulary memorization task, with rewards determined by their rank within the group. We also elicited beliefs about expected rankings before and after preparation. The experimental groups varied along two dimensions: (1) whether participants received upfront tokens (none, the average prize, or the top prize) and (2) whether five-person groups were formed randomly or based on prior performance. Results show that even under identical reward rules, differences in reference points and belief formation significantly influenced outcomes. Notably, upfront rewards were most effective when participants were grouped with peers of similar ability, suggesting an important interaction between incentive framing and social context. This research advances understanding of how behavioral and structural factors combine to drive contest outcomes, offering practical insights for designing effective incentive systems in educational and organizational settings.

Attention and Personalization in Promoting Retirement Savings: Evidence from Two Field Experiments

Zeyang Chen, Renmin University of China

We collaborate with a bank to conduct two large-scale field experiments (N = 36,116) promoting China's newly launched first-ever voluntary personal retirement savings plan through automated calls. We find that financial interventions highlighting tax benefits and compound interest outperform the industry-standard bonus by about 30% in improving call listening behavior. Behavioral nudges targeting time misperception, bequest motives, and procrastination also yield significant gains between 4%-8%. Investigating subsequent app browsing dynamic patterns uncovers novel evidence of an attention-memory-action process, triggered by time-sensitive cues such as tax deadlines and traditional holidays, highlighting the need to evaluate nudge effects beyond immediate responses. Algorithmically personalized scripts further increase listening by 6.2% and browsing by 17.4%, showcasing the value of personalized nudges. Our findings contribute to the literature on retirement savings, attention, and nudges by showing how scalable, low-cost, and personalized interventions can address the global challenge of low participation in voluntary retirement saving schemes.

Behavioral Personalization and Algorithmic Explainability: Evidence from a Field Experiment on Robo-Advising

Juanjuan Meng, Peking University

Algorithmic financial advice has the potential to improve investment performance but often suffers from low adoption. We conduct a large-scale field experiment in partnership with a major commercial bank in China to provide the first causal evidence on how behaviorally informed personalization and algorithmic explainability affect investor take-up of robo-advice. We find that a simple rule-based algorithm that tailors recommendations based on elicited behavioral traits including risk and loss aversion significantly increases the likelihood of investing in the recommended asset category by 8.11% relative to a standard Sharpe ratio-maximizing baseline. Providing investors with feedback on their elicited preferences, along with an explanation of how these inform the recommendation, also increases uptake by 14.24%. Similar significant increases are observed in both the amount held in the recommended asset type and its share in the overall portfolio. These effects are shown to be persistent over time. Our findings highlight the value of incorporating behavioral preferences and explainability in promoting more inclusive adoption of robot advice.

Choice Bracketing and Hedging under Random Incentive System

Nobuyuki Hanaki, Institute of Social and Economic Research, Osaka University

We investigate extent to which participants' tendency to bracket their choices under risk and under ambiguity to better understand possible reasons behind the difference between the experimental results of Baillon et al. (2022) and Aoyama and Hanaki (2025) regarding the effect of random incentive system in distorting participants' choice under ambiguity. The experimental results show suggestive evidence that those who fail to profit from the opportunity given by the random incentive system to hedge their choices under ambiguity is also less likely bracket their risky choices broadly (and more likely to bracket their risky choices narrowly).

Obvious Monotonicity as a Test of Behavioral Incentive Compatibility

Jin-yeong Sohn, Jeonbuk National University

We propose a test of behavioral incentive compatibility (BIC) of preference elicitation mechanisms, obvious monotonicity (OM). Applying this to Oprea's (2024a) "mirror" experiment-intended to challenge the behavioral foundations of risk preferences-we show that the elicitation method itself violates BIC. First, a large proportion of subjects fails OM: e.g., 47% of subjects value 90 boxes with $25 less than 10 boxes with $25. Second, the mirror anomalies are mainly driven by random choices of OM violators. Third, among those satisfying OM, we found the fourfold pattern and loss aversion that cannot be attributed to complexity. Thus, simplicity equivalents offer a practical benchmark for assessing BIC in future experiments.

Dismantling regret: testing the links between self-reported, behavioural & post-decisional measures of regret

Alexander Svenson, The School of Economics, The University of Sydney

Regret is an important concept for understanding decision-making. A broad literature has developed eliciting regret with various methodologies, however, the correlation across methods is less clear. This paper examines the consistency of regret as a phenomenon: regret was elicited in three different ways across different stages of decision-making: i) self-reported regret proneness before decision-making, ii) behavioural regret aversion during decision-making, iii) and post-decisional affect ratings following feedback on choices. Participants generally exhibited regret aversion, although there was individual variation in whether behaviour agreed with Regret theory. Post-decisional affect ratings were negative following feedback designed to elicit regret, and this pattern was stronger for those who self-report regret proneness. The link between self-reported regret proneness and behavioural regret aversion was weak and task dependent. There was also mixed evidence for a temporal link between exposure to regret and subsequent regret aversion. Overall, findings suggest regret may not be a unidimensional construct, rather, self-reported regret proneness and the tendency to experience regret affectively may overlap and be separate phenomena from behavioural regret aversion.

Identifying Nontransitive Preferences

Michele Garagnani, University of Melbourne

Transitivity is perhaps the most fundamental axiom in economic models of choice. The empirical literature has regularly documented violations of transitivity, but these violations pose little problem if they are simply a result of somewhat-noisy decision making and not a reflection of the deterministic part of individuals' preferences. However, what if transitivity violations reflect genuinely nontransitive preferences? And how can we separate nontransitive preferences from noise-generated transitivity violations-a problem that so far appears unresolved? To tackle these fundamental questions, we develop a theoretical framework which allows for nontransitive choices and behavioral noise. We then derive a non-parametric method which uses response times and choice frequencies to distinguish genuine (and potentially nontransitive) preferences from noise. We apply this method to two different datasets, demonstrating that a substantial proportion of transitivity violations reflect genuinely nontransitive preferences. These violations cannot be accounted for by any model using transitive preferences and noisy choices.

Intuition, Reflection, and the Roots of Utilitarian Decision-Making

Ozan Isler, The University of Queensland

We report two incentivized experiments investigating how intuitive versus reflective thinking affects moral decision-making in real monetary dilemmas-an incentivized version of the classic trolley problem, where participants decide whether to harm one person to benefit several others. Study 1

found that Greene's dual-process model was not supported experimentally: induced thinking style (emotion induction, debiasing, or time pressure) did not significantly alter utilitarian behavior. However, correlational evidence indicated that utilitarian choices were positively associated with cognitive reflection and most frequent in reflection conditions. Study 2 builds on these results by introducing a self-relevant cost manipulation to test Capraro's Generalized Social Heuristics hypothesis, which posits that intuitive moral decision-making emerges primarily when self-preservation or retaliation-avoidance motives are engaged. Participants decide under conditions varying in potential retaliation risk (no, low, or high personal cost). Our approach allows a direct empirical contrast between reflection-based and intuition-based accounts of utilitarian behavior, offering new insights into when moral reasoning or self-protection drives utilitarian choices.

Income and Intrinsic Honesty around the World

Moritz Janas, University of Gothenburg

Understanding the ethical behavior of economic elites is critical in an era where high-income individuals wield growing influence over legal and political systems. Drawing on incentivized behavioral data from over 65,000 nationally representative respondents across 70 countries, this study examines the relationship between income and intrinsic honesty around the world. Participants take part in the "Mind Game," which enables them to lie for money without risk of being detected, while allowing researchers to measure dishonesty at the group level. We uncover substantial cross-country variation: in some countries, income positively correlates with greater honesty, whereas in others with greater dishonesty. The variation in relative dishonesty is associated with institutional quality in general, and the relative treatment of high-income individuals by the judicial system in specific. While the effectiveness of the criminal justice system decreases dishonesty throughout the income distribution, high-income individuals in countries with a more impartial systems are more honest, while honesty levels of low-income individuals remain constant. Our findings highlight the role of institutions in shaping elite morality.

I would not lie for you

Tanvir Nabi Khan, North South University

In our experimental study (N = 292), subjects reported privately observed information on behalf of a group of two individuals. There are two treatment variations: the religion of the passive partner (same vs. different) and the payoff structure of the decision maker (variable vs. fixed). The primary objective of the study is to investigate the impact of religious connection (with a passive partner) on cheating behavior. We recruit students from a large university in Bangladesh, a highly religious, predominantly Muslim country. Using the Fischbacher and FÃ¶llmi-Heusi (2013) die-roll task, the study examines whether subjects exhibit a higher level of aversion to lying when a passive dependent is from the same religion as themselves, rather than a different religion, in the presence and absence of monetary incentives. Our experimental design is a 2Ã2 between-subjects design. There are four treatments: 1) subjects report for themselves and a passive member from the same religion, 2) subjects report for themselves and a passive member from a different religion, 3) subjects receive a fixed payment to perform reporting task on behalf of a passive member from the same religion, and 4) subjects receive a fixed payment to perform reporting task on behalf of a passive member from a different religion. In the first two treatments, subjects can get a bigger payoff for themselves by lying. And in the other treatments, lying is altruistic, i.e., lying only benefits a passive partner. Subjects lie more in the presence of monetary incentives to lie. We do not observe in-group favoritism in the presence of a monetary incentive or in its absence; i.e., subjects do not lie more to help a person from the same in-group. Comparing Treatment 1 and 2 reveals that, in the presence of a monetary incentive to lie,

subjects in fact lie significantly less when the passive partner is from the same religion (p-value < 0.001). We suggest that subjects lie less when lying is economically beneficial for both parties, possibly because the internal cost of lying is higher when the passive recipient shares the same religion. It appears that the information on the religious connection serves as a moral reminder, making the subject less inclined to lie. A comparison between Treatment 3 and 4 shows that subjects do not engage in altruistic lying to help a partner from the same religion. Consistent with the literature, we do not find a strong association between the trait religiosity and the tendency to cheat. Our findings suggest that subtle religious reminders can promote honest behavior in countries like Bangladesh, which exhibit high levels of religiosity concurrent with widespread corruption. The study contributes to the literature of honesty and decision making for others by documenting the influence of partners' religious identity on cheating behavior.

An Experiment on Reflective Prompts and Self-Efficacy in Risky Decision Making to Reduce Dynamic Inconsistency

Sol Chung, The University of Sydney

People often make inconsistent choices under risk, preferring safer options for the future but acting more impulsively when decisions involve immediate outcomes. This dynamic inconsistency is especially evident in gambling, where short-term rewards can dominate long-term financial goals. Understanding how to support more consistent and deliberate decision making is important for promoting financial wellbeing and responsible gambling. Previous research has shown that feedback frequency can shape risk-taking through mechanisms such as myopic loss aversion, yet little is known about whether cognitive interventions designed to reflective prompts can counteract these effects. This study examines whether reflective prompts, defined as short and timely messages that encourage individuals to pause, reflect, and recognise their control over decisions, can help reduce dynamic inconsistency in gambling behaviour. Building on theories of self-efficacy and cognitive reflection, we propose that reinforcing a sense of agency can improve the alignment between immediate and long-term preferences. We will conduct a 2 x 2 between-subjects experiment that varies the presence of reflective prompts (prompt versus control) and the frequency of outcome evaluation (frequent versus infrequent). A total of 1,000 regular online bettors, recruited via industry partners (Racing NSW and Entain Group) and commercial online research panels, will complete incentivised risky-choice tasks to measure changes in preference stability across repeated rounds. In the prompt condition, participants will receive brief self-efficacy messages that encourage reflection before or during betting, such as "You're in charge of when this session ends, is now a good time to call it?". The manipulation of evaluation frequency compares feedback provided after each round with feedback after multiple rounds. We test three hypotheses. First, reflective prompts will reduce time-inconsistent risk-taking by promoting greater reflection and self-control. Second, infrequent evaluations will support more stable preferences by reducing attention to short-term losses. Third, combining reflective prompts with infrequent evaluations will lead to the most consistent risk preferences over time. The expected outcomes include behavioural evidence on how self-efficacy influences risky decision making and practical insights for designing digital tools that promote responsible gambling. By integrating behavioural economics with psychological theories of control and reflection, the study contributes to understanding how simple and well-timed interventions can encourage more deliberate financial decisions. The findings will inform both policy and industry discussions on reducing gambling-related harm and contribute to a broader understanding of how brief reflection can reduce dynamic inconsistency, a central issue in intertemporal decision research.

Housing Conditions and Social Mobility Expectations

Wenyi Zhang, Nanyang Technological University

This paper studies how information about housing market conditions and government intervention shapes beliefs about social mobility in Singapore. Housing is the central asset class for Singaporean households: more than 80 percent of residents live in public Housing and Development Board (HDB) flats, over 90 percent of whom are homeowners. At the same time, rapid increases in resale HDB flat prices have elevated concerns about inequality, and affordability of upward mobility for younger generations. These features make Singapore an ideal setting to examine how macroeconomic signals from the housing market and policy signals from the state shift expectations about opportunity and fairness. We implement a large, preregistered online randomized controlled trial (RCT) on a nationally representative sample of adult Singaporeans and Permanent Residents. Before receiving any information, respondents report baseline beliefs about their own mobility prospects, intergenerational mobility for children from the bottom and top quintiles. Respondents then receive one of four informational treatments or no information: (i) a rising-housing-price signal, (ii) a cooling-price signal, (iii) an increase in government subsidies for homebuyers (Enhanced CPF Housing Grant), or (iv) an increase in property taxes for homeowners. The design uses active control contrasts within each information dimension, allowing us to cleanly isolate belief updating to opposing but equally credible signals rather than to an absence of information. Post-treatment expectations are elicited using parallel formats, including a repeated mobility-ladder allocation task (for 300 children instead of 100 in pre-treatment). We analyze belief updating using pre and post differences, supplemented by heterogeneity tests based on homeownership, income, education level and housing-cost burden. Baseline results indicate that individuals interpret rising housing prices through a wealth-effect lens: higher prices are associated with greater optimism about one's own upward mobility. In contrast, the policy treatments reveal a salience-driven reinterpretation of the housing environment. When exposed to the subsidy information, respondents update pessimistically for bottom 20% children and more optimistically for top 20% children. We interpret this as a salience effect: because many respondents were previously unaware of the policy change, learning that the government must intervene with sizeable additional support is itself taken as a negative signal. The need for expanded subsidies becomes salient and is interpreted as evidence that underlying housing affordability or mobility conditions are worse than expected. As a result, mobility expectations shift downward for the bottom and upward for the top, consistent with respondents inferring heightened underlying inequality rather than focusing on the policy's redistributive intent.

Dynamic Feedback and Incentive Design: Theory and Experiment

Ying Wang, HSE University

When should a firm inform employees about the outcomes of their performance evaluations, especially regarding their status for bonuses or promotions? Moreover, what type of reward scheme should accompany the timing of information disclosure to effectively motivate effort? Continuous feedback may cause high performers to slack off once they realize they have already cleared the bar for rewards or promotion, suggesting that delayed disclosure can sustain effort. However, delaying feedback can also demotivate underperformers, who, in the absence of updates, might assume they have already met the criteria. Firms therefore face a trade-off when deciding when to disclose performance information and how to structure monetary rewards. Motivated by this trade-off, I develop a continuous-time principal-agent model in which the agent decides at each instant whether to exert costly effort to generate a success. Success arrives stochastically at a constant rate when effort is exerted and can occur at most once. The principal values effort but cannot observe it; instead, the principal privately observes whether success has occurred and determines both an information disclosure policy and a reward scheme. The optimal contract follows a cutoff design: information about success is withheld until a predetermined time threshold, after which it is disclosed immediately. Withholding information preserves incentives for continued effort after success has occurred but

becomes increasingly costly as the agent's belief that success has already arrived grows. To compensate for this lack of information, the principal must offer a higher reward for early success, leading to an optimally decreasing reward schedule. The effectiveness of these optimal contractual features critically depends on the underlying assumptions of the theoretical model, both implicitly and explicitly. To test the theory by evaluating these assumptions and identifying potential behavioral deviations from theoretical predictions, I design a clear and engaging "Gold Prospecting Game" that captures the model's core trade-offs. Information treatments vary the timing of disclosure about whether gold has been struck (full information, no information, and cutoff disclosure), while reward treatments compare a decreasing reward schedule to a constant reward schedule. A within-subject design enables clean pairwise comparisons while controlling for individual heterogeneity. The experimental results strongly support the main hypotheses. Withholding information increases effort when success has already occurred, whereas revealing information increases effort when success has not yet occurred. The decreasing reward scheme effectively deters shirking and encourages early-stage effort. Individual characteristics, such as risk preferences and gender, significantly predict effort responses across treatments. (I plan to run additional experiments in the coming months and expect to have new data to present by April.) Overall, this paper provides a unified theoretical and experimental analysis of dynamic information provision and incentive design. Whereas prior studies typically treat one dimension as exogenous, this paper analyzes how their interaction shapes effort over time. It further contributes to the experimental economics literature by developing a controlled laboratory game that directly tests these mechanisms, offering novel evidence on how information disclosure and reward structure jointly influence motivation and effort dynamics.

Delegating in the Age of AI: Preferences for Decision Autonomy

Radosveta Ivanova-Stenzel, TU Berlin

To what extent are the aversion to and distrust of algorithms rooted in a fundamental discomfort with giving up decision authority? Using two incentivized experiments across distinct decision domains, hiring (social decision-making) and forecasting (analytical decision-making), and decision architecture (nature and number of decisions), we elicit participants' willingness to delegate decisions separately to an AI agent and a human agent. The within-subject design enables a direct comparison of delegation preferences across different agent types. Our results suggest that algorithm aversion may be driven less by distrust in AI and more by a general preference for decision autonomy.

How Information Shapes the Expression of Moral Judgments: Evidence from AI Use in Higher Education

Yujiao Li, University of Technology Sydney

This study examines whether misperceptions about others influence university students' adoption of generative AI. Research indicates that generative AI adoption is lower among vulnerable populations in developed countries. Prior research also shows that women tend to use generative AI tools less than men, but the underlying causes of this gap remain unclear. One possible factor is that certain subgroups may underestimate how widely generative AI is accepted or valued by others. We implement a two-stage online study with university students in Australia. The first-stage survey measures students' generative AI use, personal attitudes, and second-order beliefs about peers and employers. Results show that generative AI is widely adopted among university students. Male students and non-native English speakers report higher frequency and intensity of use. Students' own usage is correlated with their perceptions of others' usage and attitudes. A notable pattern is that students systematically underestimate how much employers value generative AI skills in hiring decisions. Informed by the first-stage results, the second stage of the project tests whether providing

accurate information about employers' demand for generative AI skills can update students' directly and indirectly elicited beliefs and influence subsequent behavior. Together, the study aims to deepen our understanding of how misperceptions affect technology use and to identify low-cost information interventions that may promote more equitable engagement.

Does New Technology Induce Task Mismatch?: Experimental Evidence from Generative AI

Junghyun Park, Seoul National University

While the adoption of new technologies such as generative AI promises productivity gains, its effects on individual decision-making and labor allocation remain unclear. We run a laboratory experiment in which subjects choose between a technology-complementary (TC) task and a non-complementary (TNC) task, with randomized access to AI. AI access increases selection of the TC task but also raises task mismatch-subjects abandon the task of comparative advantage, lowering earnings and worsening allocation. The evidence indicates a mis-calibration mechanism: participants systematically overestimate their performance with AI and make inefficient task choices. We then evaluate two light-touch interventions. Providing social information and allowing repeated task selection both significantly reduce mismatch. Our findings reveal a behavioral friction in early technology adoption: without corrective feedback, overestimation reallocates workers away from their comparative advantage and erodes the productivity gains of new technologies.

The Elusive Returns to AI Skills: Evidence from a Field Experiment

Anastasia Danilov, Humboldt-Universität zu Berlin

As firms increasingly adopt Artificial Intelligence (AI) technologies, how they adjust hiring practices for skilled workers remains unclear. This paper investigates whether AI-related skills are rewarded in talent recruitment by conducting a large-scale correspondence study in the United Kingdom. We submit 1,185 resumes to vacancies across a range of occupations, randomly assigning the presence or absence of advanced AI-related qualifications. These AI qualifications are added to resumes as voluntary signals and not explicitly requested in the job postings. We find no statistically significant effect of listing AI qualifications in resumes on interview callback rates. However, a heterogeneity analysis reveals some positive and significant effects for positions in Engineering and Marketing. These results are robust to controls for the total number of skills listed in job ads, the degree of match between resumes and job descriptions, and the level of expertise required. In an exploratory analysis, we find stronger employer responses to AI-related skills in industries with lower exposure to AI technologies. These findings suggest that the labor market valuation of AI-related qualifications is context-dependent and shaped by sectoral innovation dynamics.

Once a Competitor, Always a Competitor?

Adelson Teh, Cornell University

Using two experiments, this paper investigates whether the nature of a prior interaction between two individuals, specifically, whether they competed or collaborated, affects their subsequent interactions with each other. Both experiments have the same two-part structure. Participants are paired up randomly and anonymously, and interact with the same matched participant throughout the experiment. In Part 1, participants either competed with their matched participant or collaborated with their matched participant. In Part 2, participants interact with that same matched participant in a series of scenarios unrelated to their prior competition or collaboration. On the whole, compared to people they had collaborated with, I find that individuals are less prosocial and more antisocial towards people

they had competed with in their subsequent interactions. Consistent with previous studies, I also find that individuals have more negative beliefs about the behaviour of a past competitor.

Do Campaign Promises Affect Elected Leaders' Behavior?

Muhammad Arslan Iqbal, University of Melbourne

Do campaign promises constrain elected leaders when accountability institutions are weak? Experimental literature on guilt aversion suggests promises create psychological costs of reneging, yet elected officials often break their campaign promises. We investigate this using one-shot laboratory elections in which candidates campaign on binding tax for voters and non-binding promise about their own contributions, before deciding whether to actually contribute to or embezzle from the public fund. We vary whether leaders control taxation, whether they can embezzle, and whether they make contribution promises. Our key finding is that while over 80 percent of candidates promise to contribute at least as much as the tax paid by citizens, more than 60 percent subsequently break their promise, and do so maximally, leaving voter payoffs no better than when contribution promises are absent. Our results do not support guilt aversion among the elected leaders as, despite believing voters trust their promises, their behavior is driven by who they are rather than what voters expect from them. Moreover, making taxation endogenous does not improve promise-keeping or voter outcomes. Finally, removing embezzlement opportunities does not lead to more responsible behavior, it merely constrains the choice set. Interestingly, elected leaders without embezzlement option exploit a greater share of their available power selfishly. Overall, our results show that non-binding campaign promises alone cannot discipline leader behavior.

Too Big to Innovate

Mike Zhiren Wu, Monash University

Empirical research suggests that radical R&D activities decline as organizational size increases. This study investigates a potential explanation for this size disadvantage: ex post shading by agents. We design an experiment in which the principal repeatedly chooses between the flat payment scheme and the lottery-based payment scheme (with an option to switch back to the flat payment after observing each agent's realized outcome) for groups of different size. We find that under the lottery-based payment scheme, agents in larger groups engage in more shading against both the principal and other agents, which discourages principals from adopting this payment scheme. To mitigate this size disadvantage, we examine the effectiveness of two team management practices: screening and team building. While screening out agents with stronger negative reciprocity ex ante has no effect on agents' shading or principals' payment scheme choice, a principal-led team building activity significantly reduces agents' shading against the principal and thereby modestly increases principals' adoption of the lottery-based payment.

Economic Inequality, Social Capital and the Role of Merit and Luck: An Experimental Study

Ananish Chaudhuri, University of Auckland

Research consistently links inequality to lower social capital, the networks, norms, and trust that facilitate cooperation. Yet not all inequality is perceived the same: people are generally more accepting of inequality arising from merit than from luck. This study investigates whether the negative effect of unequal environments on social capital depends on the source of inequality: merit or luck. Using trust, prisoner's dilemma, and stag-hunt games, we find that inequality, regardless of its source, reduces trustworthiness and under-mines coordination to the payoff-dominant outcome in the stag-

hunt game. Notably, luck-based inequality results in higher trust compared to merit-based inequality, even though participants perceive meritocratic allocations as fairer.

I Will Pay Later: Time Inconsistency, Self-Control, and Loan Delinquency

Leo Bao, *Monash University*

Time consistency and self-control are central in achieving long-term economic objectives, yet their role in debt repayment remains understudied. We develop a theoretical framework demonstrating how time inconsistency and limited self-control can lead to unintended loan delinquency. We empirically test the model using a unique dataset of 58,345 fintech loan records. Consistent with the theory, time consistency and self-control measures significantly predict delinquency after controlling for traditional risk factors. These findings suggest that self-control, often overlooked in credit scoring, supplements evaluation of creditworthiness. Our results highlight the potential benefits of incorporating psychological metrics into lending decisions.

Time and State in Commitment: Experimental Evidence from Crop Insurance in Uganda

Sili Zhang, LMU Munich, Department of Economics

Many economic decisions are state-dependent, leading individuals to value flexibility. However, future state dependence may be viewed as undesirable and hence can induce-rather than reduce-demand for commitment, contrary to standard economic intuition. In a field experiment in Uganda, we offer farmers pay-at-harvest crop insurance: the demand for this product depends heavily on harvest states, and premiums are deducted post-harvest, thereby removing the role of present bias. Farmers choose insurance for two seasons and can commit by forgoing the option to cancel after the first season. We find a substantial demand for commitment: 40 percent of farmers choose commitment over flexibility. We randomize an intervention prompting farmers to think through their insurance decisions after a good and bad harvest. This intervention further increases farmers' insurance commitment by an additional 11 percentage points, consistent with the presence of undesirable state dependence, such as internal conflicts and state-dependent mistakes, as well as more negative attitudes toward state-dependence after the intervention.

When Do You Expect to Be Paid? Expectations Shape Patience

Elif Incekara-Hafalir, UTS

We investigate the effect of the expected timing of payments on elicited time preferences within an experimental framework. The key aspect of our experimental design is that we manipulate the subjects' expectations regarding when they will receive their payments. One group of participants expects their payment in 8 weeks (Future Expectation), while another expects to receive their payments on the same day as the experiment (Today Expectation). We did not manipulate the expectations of the Control group participants, who are likely to have an implicit expectation of receiving their reward on the same day as it is standard to pay subjects in economics experiments just after the participation. We then test the time preference measurements of these three groups. We find that the Future Expectation group is more patient than the other two groups, and the Today Expectation group is not statistically different from the Control group. These results align with the predictions of a simple theoretical model we developed, which illustrates how expectations over the timing of payment might affect elicited time preferences.

Learning from Peers: The Impact of Ability Grouping and Incentives

Bhagya Gunawardena, RMIT University

The effects of ability composition within groups or classrooms on learning have long been studied, yet there is mixed evidence on optimal ability grouping. In this study, we investigate the effects of different ability grouping and incentive mechanisms that promote peer-learning and teaching using a series of unfamiliar logical puzzles in a lab experiment. The findings support mixed ability grouping where high-ability participants learn more when paired with lower-ability partners, whereas medium and low-ability participants are indifferent to their partner's ability. The improved learning among high-ability participants with lower-ability partners is primarily observed under individual performance-based incentives. The findings suggest that high-ability participants learn more by teaching than by learning from other high-ability participants. Even though group incentives promote the likelihood of effective communication among participants, especially when they have relatively lower-ability partners, incentives alone do not improve learning.

AI Policy in Education: A Randomized Controlled Trial

Yanlin Wan, Hong Kong University of Science and Technology

We investigate the optimal strategies for integrating Artificial Intelligence (AI) into middle school education. Using a Randomized Controlled Trial (RCT) conducted across two urban middle schools in China, we examine the causal effects of two distinct interventions: AI guidance and restrictive AI use. While the interventions yield a statistically insignificant overall impact on academic performance, the AI guidance treatment significantly increases students' exposure to AI use, proxied by increased AI use frequency. Crucially, our results reveal a critical trade-off: the interventions, particularly AI guidance, have a negative impact on students' learning behaviors and attitudes (e.g., organizational behavior and goal orientation). However, we find significant gender heterogeneity. The combined AI use guideline and restrictive AI use intervention significantly improves math scores for girls by approximately four percentage points more relative to their male counterparts. Concurrently, the negative impact on learning behaviors and attitudes is significantly smaller for girls. We attribute this gender difference to the possibility that girls in East Asian settings, being generally more self-disciplined, utilize the provided AI guidelines more effectively. This differential utilization allows them to maximize the benefits of increased AI exposure while mitigating the associated harms to learning behaviors and attitudes.

Group Announcements vs. Individual Messages: Experimental Evidence on Digital School-Home Communication in Rural China

Jun Zhao, South China Normal University

Rapid progress in digital infrastructure and the widespread adoption of the internet in rural China have substantially reduced information costs, relaxed geographic constraints, and made communication more efficient. Leveraging digital technologies to improve educational outcomes and narrow the urban-rural achievement gap presents both an important opportunity and a major policy challenge. In rural China, left-behind children often experience severe information asymmetry due to limited parental presence; parents typically lack accurate information about their children's academic performance, which may distort their educational investment decisions. We conduct a randomized field experiment in rural China (N = 2,056) that delivers academic performance information to parents via a low-cost WeChat-based system, and examine how different modes of digital communication affect student achievement and underlying mechanisms. Both "group announcements" and "individual messages" significantly increase the frequency of parent-teacher interactions. The group-announcement intervention improves students' academic performance on average, while the positive effects of individual messages are concentrated among children from lower socioeconomic status

families. Mechanism analyses suggest that group announcements operate primarily by increasing parents' time and financial investments in education, whereas individual messages appear to weaken parent-child relationship quality to some extent. Our findings provide new evidence on effective school-home communication strategies and offer policy insights for strengthening collaborative education, promoting educational equity, and reducing the urban-rural education gap.

A Mindset Intervention in Rural China

Juliana Goncalves, University of Sydney

We evaluate a growth mindset intervention implemented in disadvantaged schools in rural China, where parental migration and socio-economic disadvantage are widespread. The study involves early adolescents (grades 7-8) across 18 middle schools in Jiangxi Province. Classes were randomized to a six-week curriculum designed to foster growth mindset and related socio-emotional skills. Using survey and administrative data collected before and after the intervention, we find significant short-term improvements in growth mindset (0.27 SD), with effects persisting although attenuated after six months. Maternal involvement amplifies these benefits, underscoring its complementarity with school-based programs. Despite socio-emotional gains, we find no impact on academic achievement or aspirations.

Smart Choice Procedures: how to overcome choice overload with simple choice procedures

Jerome Campos, The University of Sydney

Contrary to traditional economics assumption that having more options for choice is better, many papers in behavioural economics, marketing, and psychology have shown that having more alternatives to choose from can deteriorate choice quality. As a solution to this choice overload problem, behavioural economists have proposed reducing the number of alternatives available to a chooser, or to use nudges that guide people to select presumably "better" alternatives. In this paper, we take a different approach and suggest how people can make better choices autonomously. Using an understanding of how value representations in the brain determine choice accuracy, we design three choice procedures to improve choice quality without reducing the range of options available. Using a within-subject experiment, we test how our procedures compare to the most commonly used procedure, simply picking the preferred option. We find that each of our choice procedures increases the probability that a participant will correctly choose their independently identified favourite from the choice set. We show that this leads to a significant decrease in the monetary cost of mistakes. Our paper provides simple alternative choice procedures which overcome choice overload and maintain the autonomy of choice without reducing the benefit of having a variety of options to choose from.

Gaze and Attribute Weighting in Carbon-Labeled Food Choice: A MaaDDM Approach

Ying-Chen Lin, National Taiwan University

This study applies the Multi-attribute Attentional Drift Diffusion Model (MaaDDM) and eye-tracking to assess how carbon footprint labels influence food choices. Participants rated 100 snacks for taste and then made binary choices between pairs shown with carbon emission data. We found that non-fixated options and attributes were down-weighted by 40% and 20%, respectively, confirming multiplicative attentional discounting. Carbon labels significantly increased gaze time on carbon, and visual attention allocation closely mirrored model-estimated attribute weights. When taste and carbon scores were similar, increased gaze on carbon led to a higher likelihood of choosing low-carbon options, but with longer response time (RT) and more gaze shifts, consistent with U-shaped RT patterns from drift-diffusion model. In contrast, when the differences between taste and carbon were

large, choices were rapid, taste-driven, and minimally influenced by carbon attention. The full MaaDDM, incorporating both option- and attribute-level discounting, outperformed simpler variants and produced discount parameters in line with those previously reported. These findings demonstrate that carbon footprint labels exert substantial influence by directing attention to carbon information, dynamically reweighting attributes, and affecting both choice outcomes and RT.

Heterogeneity and Flexibility in Reference Points

Nanyin Yang, University of Sydney

Reference points are among the most influential yet mysterious concepts in decision-making, with continuing discussions on what they actually are and where they come from. We design a foraging task in which individuals make choices between continuing exploiting a diminishing patch or incurring costs to explore new patches, and we manipulate the exploration cost and the resource variance between patches. This game allows us to investigate how environments and experiences shape reference points and therefore influence people's exploration-exploitation strategies. Furthermore, by engaging a socioeconomically representative sample, we examine how the reference point adjustment patterns differ by socioeconomic, demographic, and psychological factors. This project will not only offer novel insights into the origin and dynamics of reference points, but also shed lights on the heterogeneity of financial behaviors from a reference-point perspective.

Winning and Losing: Sleep and Circadian Rhythm Effects on Risky Decisions in a Gain-Loss Framing Paradigm

David Dickinson, Appalachian State University

Research and anecdotal evidence indicate that lack of sleep and circadian factors impact how one makes decisions involving monetary risk. And, because risky choices have been shown to differ over gains versus losses, whether a risky choice is framed against a loss or gain sure alternative may play an important role in understanding these sleep effects. In this three-week within-subjects randomized crossover design, n=149 participants were exogenously assigned to one-week of sleep-restriction (SR) and one-week of well-rested (WR) sleep levels in their naturalistic at-home environment--sleep treatment order was randomly assigned and a wash-out week of ad lib sleep levels occurred between treatment weeks. After each treatment week, participants were administered an incentivized, framed monetary risk choice task at a time-of-day that was more or less aligned with a validated measure of their diurnal preference. The findings showed that SR and suboptimal circadian timing (i.e., circadian mismatch) are both associated with increased risk-taking behavior, while being both sleep-restricted and circadian mismatched mutes some of the individual effects. Furthermore, the study revealed that sleep and circadian mismatch both interact with characteristics of the risky choice task in ways that underscore the complex interplay between sleep, circadian rhythms, and other task attributes in risky decision-making.

Giving a Voice: Increasing Individual Self-Expression to Enhance the Resilience to System Discontent

Mathilde Bechdolf, University of Magdeburg

Individuals in a group who repeatedly experience that their group's policy selection system does not decide in their favor may develop system discontent and system disbelief. System discontent reflects individual dissatisfaction with the decision-making process, while system disbelief captures the perception that the system does not benefit the group as a whole. Both may be detrimental to collective outcomes, affecting group members' psychological well-being and their willingness to

provide work effort, make financial contributions, or participate in cooperative coordination. In this experimental study, I investigate whether allowing individuals to express and explain their preferences, affects the development of system discontent and system disbelief. I examine three different group policy selection mechanisms, each combined with two communication modes: with and without voice. Decisions are made either by a single decision maker (dictator), by AI (ChatGPT), or by Borda Count (automated). In the latter two treatments an independent observer is added to keep the overall number of players constant across treatments. The study adds to the knowledge on the drivers of discontent with policy selection systems and provides insights for managers and policymakers concerning the design of such systems, especially with respect to building resilience towards system discontent and system disbelief.

Rationality-based Preference Aggregation

Syngjoo Choi, Seoul National University

This paper examines how individual rationality shapes collective decision-making and how differences in rationality translate into bargaining power within groups. The literature on collective choice recognizes that aggregating rational individual preferences can yield collectively inconsistent outcomes (Condorcet 1785). Yet empirical evidence on how the degree of individual rationality influences group outcomes remains limited, partly due to measurement and identification challenges. We address these challenges through a large-scale randomized field experiment that directly links individual rationality to group decision quality and individual influence. We conducted two experimental waves with 1,573 middle school students from 12 schools in South Korea. Each participant made 18 portfolio choices under risk individually and another 18 jointly with a randomly matched classmate. Each choice represented an optimal allocation of Arrow securities within a linear budget set, following Choi et al. (2007). Group formation was randomized within classrooms and maintained over time, producing a balanced panel of 652 pairs (1,304 individuals). Complementary surveys elicited cognitive and non-cognitive traits and mapped friendship networks, enabling us to control for personality and social proximity. We measure rationality using the Critical Cost Efficiency Index (CCEI; Afriat 1972). Our first finding is that group rationality increases with individual rationality. Groups achieve higher CCEI scores when at least one member displays greater individual consistency with utility maximization. A larger within-pair CCEI gap, by contrast, reduces collective rationality. These relationships are stable across two waves and robust to controls for gender composition, math ability, and friendship links. The Shapley decomposition analysis confirms that individual CCEIs explain the largest share of the variance in group CCEI, far exceeding other observable traits. We then propose a new nonparametric revealed-preference index to measure individual bargaining power in group choices. The index evaluates how much the group's consistency (measured by the CCEI) declines when combining an individual's data with the group dataset, normalized by the joint decline when both members' data are added. This revealed bargaining-power index is interpretable as a Shapley allocation in a cost-sharing game and satisfies efficiency, symmetry, and monotonicity properties. Validation exercises show that the index closely tracks both participants' self-reported perceptions of influence and true bargaining weights in simulated parametric models. We find that individuals with higher CCEI scores exert greater revealed bargaining power: their choices are systematically closer to group decisions, and their influence increases over time. On average, the bargaining index equals 0.7 for the more rational member versus 0.3 for the less rational one, and the gap widens in the endline, suggesting persistence and reinforcement of influence. The Shapley decomposition analysis confirms that individual CCEIs explain the largest share of the variance in bargaining weights, far exceeding other observable traits. Overall, the results show that individual rationality not only enhances the quality of collective decision-making but also determines relative influence within groups. The paper develops a general revealed-preference

framework for identifying how rationality drives both collective rationality and bargaining power in preference aggregation.

Linear and Quadratic Voting in Strata Assembly

Panharidh Kun, University of Adelaide

In many jurisdictions, a collective sale of multiple strata units can only be implemented if a supermajority of sellers- often 75-80 per cent- agree. These rules were meant to relax unanimity and unlock redevelopment, but has limitations; no matter their individual preferences, they treat every "yes" and "no" the same. They can also let a small minority prevent a highly valuable sale, and can just as easily let a slim majority force out deeply undesirable outcomes, fueling debates over fairness, compensation, and even calls for eminent domain. Quadratic voting (QV) rules are generally considered a more refined alternative, which encourages people to buy votes to reflect the intensity of their preferences. In large groups, this has been shown to be a theoretically efficient pricing rule. However, QV is intricate and unfamiliar in real life situations- owners must understand a curved pricing schedule, so its appeal on paper may not translate easily into everyday decisions. A much simpler rule is a linear voting (LV) price mechanism, which is more straightforward, but standard theory does not yield clean predictions on efficiency. Yet, if sellers simply follow rules of thumb in practice, this plain pricing rule may still give rise to broadly efficient outcomes. This paper investigates the relative performance of LV and QV price mechanisms in the context of small-scale strata assemblies, and asks whether either improves on a standard 75% sales rule. To test the efficiency of these voting mechanisms, we run experiments where participants play as sellers of individual units of a collective parcel. In each market, a buyer simultaneously gives each seller an individual offer to buy their unit. In contrast to standard one-person-one-vote rules, each seller, when presented with an offer, has the chance to buy votes either for or against the sale of their unit. Our treatments differ in how votes are priced: sellers buy votes under a linear and quadratic cost function for our LV and QV treatments respectively. In both voting treatments, the collective sale is implemented if and only if total votes in favour exceed votes against. We find that both LV and QV rules substantially reduce holdouts and expropriation-type outcomes relative to the baseline standard 75% rule. QV performs marginally better than LV in terms of efficient outcomes and surplus creation, but these differences are not significant. Rather, the magnitudes of the gains (or loss) from trade- the stakes at hand- emerges as the main determinant of whether efficient outcomes occur. Our results suggest that, despite its theoretical inefficiency, a simple linear pricing rule for votes can perform comparably well to more complex quadratic voting in small-scale strata assembly environments. This opens up the possibility of jurisdictions currently utilising 75% rules in moving toward simpler costly voting mechanisms that better respect minority sellers while still allowing valuable redevelopments to proceed.

The Impact of Prejudice and Stereotypes on Electoral Efficiency: Experimental Evidence

Ralph-Christopher Bayer, University of Adelaide

We conduct a large-scale online experiment to examine how prejudice and stereotypes affect the efficiency of elections, where voters can choose to acquire information, vote uninformed, or abstain. Surprisingly, neither prejudice nor stereotypes increase uninformed voting, nor do they influence the likelihood of abstaining or purchasing information. However, both introduce bias among uninformed voters and shape the decisions of those who acquire information, reinforcing confirmation bias. As a result, prejudice and stereotypes distort electoral outcomes and reduce efficiency.

Speaking with Confidence: Who Benefits and When? Evidence from Physicians

Nisvan Erkal, University of Melbourne

Gender biases can shape career trajectories, contributing to persistent gender gaps in labour market outcomes. In many professions, senior workers assess the progress of junior colleagues, and these evaluations may be influenced by gender-based stereotypes. In collaboration with Regional Anaesthesia United Kingdom, we study whether medical consultants assess trainees' competency and readiness for independent practice differently based on gender. To investigate this, we conduct a randomized vignette experiment in which consultants evaluate trainees under varying conditions, including confidence level, prior experience, procedural outcome, and gender. After observing an interaction between a trainee and a consultant in the vignette, consultants provide their assessments of the trainee's performance. Our results reveal a confidence premium, such that trainees who come across as more confident receive more favourable treatment, specifically, less direct supervision, holding everything else constant. Confidence also affects how outcomes are evaluated. We observe a gender difference with confident women being penalized less than unconfident women for failures and confident men being rewarded more than unconfident men for successes. Thus, confidence benefits women when they fail and men when they succeed.

Empowering or Entrenching? AI Assistance and Gender Gaps in Competition

Xu Zhang, Hong Kong University of Science and Technology (Guangzhou)

We conducted a laboratory experiment (N=240) to test how access to generative AI affects the well-documented gender gap in competitiveness. Building on the Niederle and Vesterlund (2007) paradigm centered on the choice between piece-rate and tournament incentives, all participants completed three rounds of numerical and logic questions, while half were randomly assigned to optional GPT assistance. Our results confirm a significant baseline gender gap in the control group, which is fully explained by confidence, risk aversion, and ability. Access to GPT raised performance equally across genders but had no significant net effect on the competition gap. A mechanism analysis reveals that GPT influences entry through two offsetting pathways: it raises performance, which encourages entry, but simultaneously lowers confidence, which discourages it. Our findings suggest that achieving equity in the AI-augmented labor market requires pairing access with complementary interventions like confidence-building support, rather than relying on technology alone.

AI, Gender and Fairness

Lingbo Huang, Shandong University

How does access to assistive AI affect judgments of desert? We run a two-stage online experiment with mixed-gender pairs. In a symmetric benchmark, neither worker has access to Smart Tools (ST); in an asymmetric condition, one worker is randomly given ST with known accuracy. Workers first work on a task with or without access to ST and are assigned an initial bonus according to their relative performance; they then request a split of bonus. Independent spectators then reallocate the bonus, before and after observing requests. When technology is symmetric, male losers are penalized (by receiving lower transfers) relative to female losers. When ST is salient, this gender gap collapses and is replaced by a robust technology discount against ST users. Requests attenuate but do not eliminate the penalty. The results bridge gendered evaluation and fairness under unequal opportunities created by differential AI access, and motivate technology-neutral evaluation rules and equalized or randomized access within cohorts.

Who wants to lead in crises: an experimental study of the Glass Cliff

Rebecca Heath, University of Cambridge

Women often take leadership positions with a high risk of failure. This phenomenon has been labelled as the Glass Cliff. While existing literature has focused on demand-side motivations for nominating women for precarious positions, limited attention has been paid to the supply side. In this study, we investigate gender differences in the willingness to lead financially successful versus financially failing entities under two different leadership selection mechanisms. Under the Opt-in mechanism, candidates must actively express interest in the leadership position. Under the Opt-out mechanism, all qualified individuals are automatically considered for the leadership position, unless they actively choose to decline. Under the Opt-in mechanism, we find that men are less likely to apply for leadership positions in failing entities than in successful entities. In comparison, an entity's financial status does not affect women's willingness to lead. Consequently, we find a gender gap in leadership in successful entities but not in failing entities. In contrast, under the Opt-out mechanism, this gap disappears, suggesting the leadership selection mechanism plays an important role in shaping the conditions under which the Glass Cliff arises.

The Impact of Trader Incentives on Stability in Option Markets: A Controlled Market Experiment

Hoang Long Nguyen, The University of Melbourne

Traditionally, options served producers and consumers of the underlying goods seeking to hedge inherent risks. However, the potential for profits has attracted institutional traders whose incentives are often tied to profit-based performance bonuses. In this project, we explore how the tension between different incentives influence price stability and traders' ability to fully hedge their position. Specifically, we address three core research questions: (1) Does supply and demand alone generate no-arbitrage pricing in an option market where the standard arbitrage channel is disabled? (2) Is this pricing mechanism stable and self-correcting? (3) Does changing the trader's incentive structure fundamentally impact this stability? We adopt the Binomial Option Pricing Model (BOPM) framework, where the market features a stock whose prices follows a binomial process and a European call option whose payoff depends on the stock. We assume equal numbers of traders who are endowed with either long or short stock position that is non-tradable, shutting down the standard arbitrage channel. We model traders with a homogeneous, piece-wise linear utility function on their portfolio's value change. Our price model is based on option supply and demand generated by these utility-maximizing agents. We predict a unique, stable market-clearing price at the BOPM level and that traders can hedge optimally if loss-averse, while with gain-seeking agents, prices deviate and traders end up with unhedged risk. In a controlled experiment, we test this setting, implementing the market as an online continuous double auction, with a 2 (Incentive) by 3 (Strike Price) within-subject design. Our primary manipulation is participants' performance (utility) function so they either are all loss-averse or gain-seeking. We orthogonally vary strike price to ensure our results are not tied to a certain price level. This design allows us to cleanly investigate how incentives impact price convergence, market stability, and traders' ability to hedge. Our pilot study (N=30) confirms the feasibility of this design and provides strong preliminary support for our baseline. In the single-period with loss-averse setting, option prices converged to the BOPM no-arbitrage level in under two minutes, and the majority of traders successfully hedged their risk. Confirming that supply and demand from pure hedgers are a sufficient mechanism for stable price discovery, this provides a strong foundation for our full experiment, which will generate the first behavioral evidence on how incentive structures impact option market stability. Our work provides a novel test of the foundations of option pricing. By testing whether hedging-based supply and demand provide a sufficient, independent channel for price convergence, we examine the necessity of the no-arbitrage assumption on pricing. Our findings are directly relevant to market participants like FX managers operating in incomplete markets, offering crucial insights into price reliability. For policy makers, our manipulation of incentives speaks to debates on trader

compensation. If gain-seeking incentives are shown to fundamentally destabilize markets, it provides concrete evidence that compensation structures can be a source of systemic risk.

Behavioral Dimensions of IPO Mispricing: Evidence from the Laboratory

Jingru Wang, Waseda University

IPO underpricing and long-term secondary market performance are jointly shaped by how information is aggregated in the primary market and how investors trade once secondary markets open. Although theory typically assumes incentive-compatible information reveal in book-building, and empirical work often attributes IPO anomalies to asymmetric information, we have limited evidence on how mispricing and sentiment can arise when information is simple and symmetric, but strategic uncertainty and heterogeneous risk attitudes remain. Furthermore, our understanding of how various market entry patterns influence the level of secondary market sentiment is limited. This paper addresses these questions with a two-stage laboratory experiment that links an IPO book-building mechanism to a dynamic secondary market. In the first stage, investors receive private signals and submit reports in a theory-based book-building environment in which complete information aggregation is an equilibrium, but requires correct beliefs about the reporting strategies of others. In the second stage, participants trade the allocated shares over five call-auction periods, updating beliefs from prices and order flow. Across treatments, we vary the effective entry conditions into the secondary market to study how the composition of active traders shapes the emergence of sentiment. The data reveal a robust mispricing path: IPO prices are systematically below the fundamentals implied by the private signals, opening prices in the secondary market overshoot, and later prices partially but not fully reverse. Markets with more selective entry exhibit stronger upward pressure and more persistent overvaluation. At the individual level, risk aversion and prudence reduce participation in the IPO, prudence further lowers the likelihood of entering the secondary market and increases holding duration, and longer holding is associated with lower profits as prices drift back toward fundamentals. These findings show that even with simple and symmetric information, mispricing and sentiment can arise from strategic uncertainty in the primary market, updating of beliefs in the secondary market, and the composition of traders excluded at the IPO stage. The experiment thus provides a behavioral and strategic foundation for understanding IPO underpricing, first-day returns, and long-term secondary market dynamics.

The Impact of Attention Allocation Bias on Stock Market Overreaction: An Experimental Investigation

Xiaoxue Zhao, Shanghai International Studies University

This study investigates the impact of attention allocation bias on investors' decision-making by integrating eye-tracking technology with real-world data from China's A-share market. In the experiment, participants were instructed to evaluate historical price charts of stocks and make subsequent investment decisions based on these charts. The charts were standardized to accentuate salient features, including price trends and the timing of specific price shocks. To simulate a realistic investment environment and measure willingness to pay, we employed the Becker-DeGroot-Marschak (BDM) mechanism. At the outset, each participant was endowed with a fixed amount of experimental tokens. They were then presented with the price fluctuations of a specific stock over a preceding 20-week period. Following this observation phase, participants engaged in a simulated bidding exercise against a computer algorithm. The final payoff for each participant was determined by the combination of the bidding outcome and the subsequent price movement in the 21st week. During the experiment, we employed the EyeLink 1000 Plus eye-tracking device to record participants' eye movements, gaze patterns, and pupil dilation in real time throughout the decision-making process.

The study sample comprised 121 non-students with financial backgrounds and 129 students with finance-related coursework. The experimental results demonstrated a clear association between the degree of participants' overreaction to stock price shocks and specific characteristics of the historical price charts, particularly the fixation duration on those charts. Participants exhibiting stronger overreaction spent significantly more time fixating on shock points. While both student and non-student groups showed a significant overreaction tendency, their eye-movement patterns differed. Additionally, factors such as cognitive ability and financial literacy also exert an impact on the degree of overreaction.

Equilibrium Play in Experimental Parimutuel Betting Markets

Joshua Miller, University of Melbourne

Parimutuel wagering markets have attractive features, both as test beds for financial market theory, and as mechanisms for aggregating information. We design a laboratory experiment to test theoretical predictions of behavior in parimutuel betting markets. In our baseline game, market outcomes align with Bayesian Nash equilibrium predictions and shed light on the origin of the well-known favorite-longshot bias. In games with slightly richer information structures theoretical predictions depart from intuition and market outcomes depart from theoretical predictions.

Competing Motivations: When More Incentives Lead To Less Effort

Kieran Gibson, University of Queensland

I propose a framework of competing motivations in which different motives to exert effort can interfere with one another. Consequently, combining a strong incentive with a weaker one may be less motivating than offering the strong incentive alone. I test this hypothesis in three pre-registered experiments on Amazon MTurk via CloudResearch (N = 3,412). In the first study, participants completed real-effort tasks for monetary rewards, charitable rewards, or both; those working solely for money completed 23 percent more tasks than participants offered the same monetary reward combined with a charitable incentive. In the second study, participants worked for monetary rewards, a lottery reward, or both. While increasing the monetary piece rate raised effort, combining money with a lottery reward did not. The third study isolates the weaker incentives previously paired with money and shows that both the charity and lottery rewards independently increased effort relative to a no-incentive control. These findings support the competing motivations framework and challenge the assumption that motivations from different incentives combine additively.

AI Helps Until You Say It Does

Yue Wang, Shandong university

Policies related to AI often mandate transparency, particularly the disclosure of AI use in content generation. We evaluate the impact of such disclosure with a charitable fundraising experiment, in which participants receive AI-generated fundraising letters with or without disclosure of their origin. We find that AI-generated letters, whether edited by humans or not, significantly increase donations. However, this positive effect diminishes when the use of AI is disclosed. A similar pattern emerges when human-written letters are edited by AI. These findings suggest that, while AI improves the effectiveness of fundraising, disclosing its use offsets this benefit.

Ethical Bandits: An experimental study on the multi-armed bandit with an externality

Daniel Woods, Macquarie University

In this project, we theoretically incorporate and experimentally investigate adding an externality into a multi-armed bandit problem. The externality also evolves through a bandit process, in that it is initially unknown but learnable through experience. In a typical multi-armed bandit problem, an individual trades off immediate reward against information about future rewards, where rewards directly benefit the individual. By adding an externality, the individual must conduct a separate trade-off for the affected party, alongside their own rewards according to their other-regarding preferences. We vary the type of externality by treatment: it either affects the payoff of another experimental subject, or the amount of carbon offsetting credits purchased through a certified retailer. We also include treatments with no externality, as well as treatments where the externality affects the individual's own payoffs, in order to investigate the effects of additional complexity from the externality bandit process.

Poking Holes and Adding Points in Dictator Games

Cary Deck, University of Alabama

Deviations from choices predicted by self-regarding preferences have regularly been observed in standard dictator games. Such behavior is not inconsistent with conventional preference theory or revealed preference theory, which accommodate other-regarding preferences. By contrast, experiments in which giving nothing is not the least generous feasible act produce data that is inconsistent with conventional preference theory including social preference models and suggest the possible relevance of reference point models. Two such models are the reference-dependent theory of riskless choice with loss aversion and choice monotonicity in moral reference points. Our experiment includes novel treatments designed to challenge both theoretical models of reference dependence and conventional rational choice theory by poking holes in or adding to the dictator's feasible set along with changes to the initial endowment of the players. Our design creates tests that at most one of these models can pass. However, we do not find that any of these models fully capture behavior. In part this result is due to our observing behavior in some treatments that differs from previous experiments for reasons attributable to implementation differences across studies.

Pricing for Product Awareness in Social Network: An Experimental Study

King King Li, The Hang Seng University of Hong Kong

We experimentally investigate how social network affects firms' pricing decision. We compare pricing decision of firms when there is social network vs. without social network. In the experiment, some consumers are aware of the product and some are not. When there is social network, consumers who are not aware of the product can learn about it from their friends who purchased it earlier. In such case, firms may be motivated to set lower price to induce consumers to purchase so that their unaware friends can be aware of the product in later stage and potentially purchase the product. The experimental result supports this hypothesis. The results suggest that price in the social network was lower than the price when there was no social network when the proportion of initial proportion of consumers aware of the product was large. There was a non-monotone u-shape relationship between size of social network and price.

Social Learning or Social Pressure? Peer Effects in Parental Investment: Field Experimental Evidence from China

Le (Lyla) Zhang, Macquarie Business School

How do parents respond to their peers' choices when deciding whether and how much to invest in their children's extracurricular education? Using a field experiment, we identify distinct mechanisms through which peer information shapes investment decisions: positive and negative social learning, as well as social utility. Consistent with positive social learning, we find that peers' willingness to purchase a course increases both parents' willingness to purchase and their level of investment. Negative social learning effects, however, depend on the nature of the signal. Explicit peer criticism significantly reduces both willingness to invest and spending, whereas peer non-purchase without negative evaluation has no effect. Social utility effects also vary by peer purchase price: high-price peer purchases increase spending, whereas low-price peer purchases raise a parent's willingness to invest but do not lead to higher expenditures. We further document heterogeneity in responses to negative peer signals across socioeconomic and residential contexts. Using an enhanced causal forest approach to complement our regression analyses, we find that peer non-purchase without explicit criticism increases purchase preference among parents in less central areas with smaller socioeconomic gaps, but reduces it among parents in central areas with larger gaps.

Social Networks and the Emergence of Social Norms: An Experiment

Yohanes Riyanto, Nanyang Technological University

Although much research has examined how established social norms influence behavior, less is understood about how these norms emerge and what factors shape this process. We experimentally investigate how network structure shapes the emergence of social norms. Subjects with heterogeneous preferences are randomly assigned to one of four network topologies: Random, Segmented, Centralized, or Core-Periphery. Participants face both local incentives to match the choices of their immediate neighbors and global incentives to align with the majority choice in the broader population. Relative to random connections, organizing individuals into small communities in Segmented networks fosters local coordination but sustains persistent segregation. By contrast, the presence of a central agent in Centralized networks promotes global convergence through a top-down process, whereas the lower connectivity of Core-Periphery networks hinders such convergence. These differences can be partly explained by variation in belief formation, as more sophisticated reasoning emerges only in Centralized networks.

Norm Nudging Under Delegation

Sumaiya Bhura, Monash University

Social norms have powerful influences on individual decision-making. However, in many real-world contexts, decisions can be delegated to others, allowing individuals to distance themselves from the blame associated with their choices. This study examines the effectiveness of norm nudges in allocation decisions when individuals have the option to delegate those decisions to others. Our theoretical framework predicts that if delegation reduces the decision maker's disutility from norm violations, norm messages intended to promote fair allocations will be less effective when delegation is possible. We test these predictions in an online experiment. The findings highlight a potentially important limitation of norm-based interventions when individuals can shift responsibility to others.

Norms Make Investors Socially Responsible

Peiran Jiao, Maastricht University

Socially responsible investment (SRI) has been extremely popular among investors in the past decade. However, the reasons for this interest are not completely clear. The prevailing theoretical explanations rely on assumptions of some form of altruistic preferences that nevertheless fall short of

accounting for various known phenomena related to SRI. Using carefully designed experimental tasks, we test the alternative hypothesis that SRI behavior is motivated by the desire to comply with social norms. In our experiment, the behavior of the participants contradicts the predictions of pure social-preferences models. However, it is consistent with the view that SRI choices are driven by perceived social appropriateness. We provide multiple lines of evidence that corroborate this hypothesis. Our results suggest that promoting norms related to SRI is key to sustaining high interest in SRI.

Information Pooling in the Household: The Role of Expertise

Priyoma Mustafi, Ahmedabad University

A dismal view of household decision-making arises from evidence that individuals fail to incorporate information their spouses hold. I explore whether failure to pool in- formation in the household arises from a reluctance to learn from one's spouse. Using a lab-in-field experiment with 400 married couples in Kolkata, India, I examine information pooling across two domains: (i) a gender-neutral ball-in-urn domain where neither spouse is better informed, and (ii) a novel gendered pricing domain, where individuals have to price a basket of either male or female products, and each spouse is better informed in their gender-congruent domain. In the gender-neutral domain, I replicate the finding that households fail to pool information. However, in the gendered domain, I find that households pool information successfully. The results point to common knowledge of expertise as facilitating communication, effectively incentivizing spouses to both talk and listen.

Gender, Culture, and Norms of Equality

Pushkar Maitra, Monash University

Are women punished differently than men for violating social norms, and does punishment vary across cultures with different levels of gender equity and autonomy? We study these questions using experiments conducted in two Indian cultures at opposite ends of gender equity: matrilineal societies in Meghalaya and patriarchal societies in Haryana. We focus on the norm of equal distribution of resources, considered a fundamental social norm. In Meghalaya, we observe strong adherence to the norm of equality and high punishment of transgressions, with men being punished more often than women. In Haryana, we find frequent violations of this norm and low punishments among both men and women. Our findings are potentially explained by differences in norms across societies: norms of equality are strong in Meghalaya, but weak for both genders in Haryana.

Party-influenced media, independent news and information traps

Ravindra Agrawal, Indian Institute of Technology Bombay

Independent news is a cornerstone of liberal democracy. Among its many roles, free and fearless journalism keeps the voters informed and enables them to choose good leaders through elections. On the other hand, almost all major democracies are experiencing direct or indirect takeover of an increasing number of media outlets by political parties, enabling politicians to influence news and endanger press freedom. Despite the co-existence of party-influenced media alternatives where politicians undertake hidden but costly investments, which in combination with unobserved quality produce electorally favorable signals for them, we show, using a `noisy signaling model', that any improvement in independent news continues to enhance voter welfare other than in situations which we call `information traps'. When democracies fall in an information trap, stronger independent news asymmetrically affects media investments (and thereby media control) across high- and low-quality politicians, reduces the reliability of party-influenced media further, overpowers the direct benefits of

independent news, and leads to an ultimate decline in the expected quality of the elected leader. While exiting from information traps requires large improvements in independent news, allowing for some degree of media influence may enhance voter welfare when policy changes for large improvements in independent news are hard to come by in the short run. Despite information traps, we then demonstrate that ex-ante voter favoritism towards any politician can increase equilibrium information, enhance expected leadership quality, and even curb information traps. The results obtained in this paper provide novel implications for talent-spotting in labor markets.

Dr. Cutoff: How I Learned to Love Preregistration

Wei-Cheng Chen, Academia Sinica

We study how peer-review policies shape researchers' incentives to invest in research design versus post hoc persuasion. Scientific publication is modeled as a game between a scholar and a reviewer. Before seeing data, the scholar chooses design, which raises expected true quality and increases the marginal cost of later persuasion. After observing data, the scholar allocates effort between analysis, which improves informational content, and persuasion, which only changes how results are presented. The reviewer observes different subsets of these choices under three regimes: traditional review (results only), binding preregistration (design and results with verifiable deviations), and results-blind review (design only). We characterize Perfect Bayesian Equilibria under a constructive belief refinement derived from a full-support perturbation. Design acts as a commitment device: by making persuasion more costly, stronger design shifts effort from persuasion to analysis and alters the reviewer's optimal cutoff. Under log-concave priors (``incremental'' environments), higher design FOSD-shifts posterior quality and strictly lowers the optimal cutoff, which underpins a Pareto improvement of preregistration over traditional review. In contrast, with bimodal priors (``disruptive'' environments), higher design can move weight toward extreme unfavorable states and raise the optimal cutoff. In such cases, preregistration and results-blind review need not improve welfare and can tighten acceptance standards despite higher design effort.

Credibility in Persuasion: A Laboratory Experiment

Yuki Shiomi, The University of Osaka

We experimentally test how the interaction between the sender's commitment ability and the players' payoff structure affects information transmission in a Bayesian persuasion framework. The theory predicts that even when the sender's commitment ability is limited, meaningful information can still be transmitted (i.e., persuasion is credible) when the marginal incentives between the sender and the receiver are aligned, whereas persuasion is not credible when their incentives conflict. Consistent with these predictions, our experimental results show that the presence of both limited commitment and conflicting incentives almost completely suppresses information transmission between players. Contrary to the theoretical prediction, however, we also find that limited commitment and conflicting incentives independently exert negative effects on information transmission.

Informed Information Design: An Experiment

Ninghua Du, Shanghai University of Finance and Economics

Since it is common in real business practice for senders (e.g., sellers) to acquire information about states (e.g., product quality) before choosing disclosure strategies (e.g., advertising), we experimentally investigate the choice of information mechanisms by informed senders in a seller-buyer context. Theory predicts that an informed seller commits to the same (interim-optimal) noisy information mechanism in all states. However, our experimental results show that the informed seller

commits to full disclosure in a favorable state but to noisy disclosure in an unfavorable state, as naïve buyers do not glean information about states from the information mechanisms. When deciding whether to follow a noisy signal, the results show that the buyer's risk attitude plays a significant role.

Sequential Investment Mechanism and Gradual Trust: An Experimental Study

Yunwen He, Central University of Finance and Economics

This study investigates the dynamics of sequential investment mechanisms and the development of trust in multi-stage interactions using an investment game experiment. By comparing pre-committed and discretionary mechanisms, the study uncovers distinct behavioral patterns and efficiency outcomes. Pre-committed mechanisms promote steady trust development, balanced fund allocation, and high efficiency in two-stage investments, making them suitable for stable and well-defined projects. In contrast, discretionary mechanisms emphasize flexibility and adaptability, achieving higher efficiency in three-stage investments but exhibiting lower efficiency in two-stage investments compared to pre-committed mechanisms. Trustors in pre-committed settings allocate funds evenly across stages, while those in discretionary settings follow progressive strategies, reserving larger portions for later stages. These allocation behaviors likely arise from a heuristic approach, where trustors simplify their expectations of trustees' actions. The findings highlight the trade-offs between commitment and flexibility in trust-based systems, offering practical insights for designing institutions that balance these elements to foster trust, enhance social efficiency, and address diverse applications.

Communication and the Polarization of Trust Attitudes

Zhuokun Liu, Goethe University Frankfurt

Given the critical role of generalized trust in the functioning of both public and private organizations, it is important to better understand the factors driving the substantial differences in trust across states, regions, and organizations. Group polarization theory suggests that communication might foster consensus within, while increasing differences across social groups. To test this argument, we conduct an online experiment in which some participants communicate before playing separate trust games. Participants only communicate with other participants in the same role -- trustors talk with other trustors. Our results show that communication reduces the variation in trust within, while increasing trust differences between the chat groups, thus leading to a polarization of trust cultures. We further find that groups that develop strong positive trust cultures focus in their discussions on both moral principles and self-interest.

Reciprocal Trust: Lead with Trust to Inspire Trust

Yun Zhang, University of Sydney

This study explores the concept of reciprocal trust, which refers to a causal relationship in interactions between two parties, where one party's trust influences the other's trust in return. Through a randomized trust game experiment, we investigate the presence of reciprocal trust across different treatment conditions, specifically, whether participants interact with a new or previous partner, and whether information is available regarding the partner's trustworthiness. Our findings indicate that reciprocal trust generally exists across all treatment groups, though its magnitude varies. These differences suggest the influence of underlying mechanisms such as perceived trustworthiness and intentions for direct exchange. We further examine whether reciprocal trust constitutes a profitable strategy in terms of material payoff. Our results show that it is indeed profitable when participants interact with a known partner and no negative signal about the partner's trustworthiness is present.

These findings suggest that initiating trust can effectively inspire trust in others. Moreover, trusting a previous partner at a level consistent with their own trust level is materially beneficial, unless there is evidence indicating that the partner's trustworthiness falls below the median.

Personality as a Strategic Signal: Experimental Evidence from Trust and Dictator Games

Helena Luo, University of Melbourne

Why do people choose to cooperate in some interactions but not others, even when the objective incentives are identical? Economic models typically explain cooperation through structural features of the game, whereas psychological models emphasise stable traits that predict cooperative preferences. We integrate these perspectives by conceptualising personality expression as a strategic signalling process in which individuals selectively express trait-relevant states to influence others' beliefs and behaviour. Drawing on signalling theory, we model cooperative interaction as a sender-receiver game where one party holds private information about their prosociality and may choose to express personality cues as signals; the receiver observes these cues, forms beliefs about the sender's intentions, and chooses whether to cooperate. Strategic signalling arises when expressions are selected to shape counterpart beliefs and behaviours in beneficial ways. We test this framework in two studies using economic games. Study 1 (N = 6 experts; N = 320 general-population participants; online experiment) identified personality items seen as informative signals in cooperative contexts. Experts and laypeople evaluated the usefulness of partner personality items for determining their own behaviour in a Trust Game and a Dictator Game. Both groups judged knowing their partner's scores on prosocial items that reflect care for others (e.g., "I try to give generously to those in need") as more useful than Extraversion items, especially in the Trust Game. In hypothetical choice tasks, participants adjusted their giving and expectations about partners based on these scores, treating higher prosocial scores as evidence of cooperative intent. Study 2 (N = 315; incentive-compatible pre-registered lab experiment) tested whether people incur costs to strategically signal personality, and whether such signals shift beliefs and behaviour. Participants played both games in both roles under three signalling conditions (Prosocial, Non-prosocial, No-info). Before each game, senders of the signal could pay to change their personality score using a Becker-DeGroot-Marschak mechanism. Receivers then viewed the personality signal from the senders and made transfer decisions. As hypothesised, participants bid more to adjust prosocial than non-prosocial items, believing that higher prosocial scores would lead partners to give more. However, contrary to our hypothesis, this willingness to signal prosociality did not differ significantly between the Trust and Dictator Games. We found that receivers acted on these signals, giving more to partners with higher prosocial scores. Relative to the non-prosocial condition, prosocial signals modestly increased efficiency in the Trust Game and reduced inequality in the Dictator Game. Taken together, these findings show that personality information is treated as an economically meaningful signal as people believe it forecasts cooperative behaviour, will sometimes incur a cost to manage how it is presented, and adjust their own behaviour in response to it. Our framework and results bridge personality psychology and experimental economics by showing that cooperation depends not only on traits or incentives, but also on beliefs about a partner's personality. In this view, cooperation emerges from the strategic exchange of personality signals, the beliefs they generate, and the behaviours they induce.

Bound by Norms, Seen Through Stereotypes: Meta-Beliefs and Minority Women's Labour Outcomes

Elif Incekara Hafalir, University of Technology Sydney

Individuals' work decisions reflect not only their own preferences but also what they think others believe, within their communities and in the wider society. We study two belief layers: expected

stereotypes (third-order beliefs about how the majority perceives one's group) and perceived community norms (second-order beliefs), and how they relate to female labour outcomes at the gender-minority intersection. We implement an incentivised design that separately elicits first-, second-, and third-order beliefs in a labour-market context, introducing a novel third-order (meta-stereotype) measure: how minority participants believe the general public perceives their community's norms. Our approach complements evidence that misperceptions about others' views are large and malleable [Bursztyn and Yang, 2022], and provides the measurement needed to test whether meta-stereotypes correlate with labour outcomes. The setting is Muslim women's labour participation in Australia, where many navigate a dual constraint: gendered home norms within community and anticipated stereotypes in mainstream workplaces [Syed and Pio, 2010; Foroutan, 2008; Hassan, 2015; Yasmeen, 2007]. We study three samples in Australia: Muslim university students, working professionals, and general university students. We elicit beliefs on two statements: (i) women's right to work and (ii) whether a young child suffers if the mother works. Three results stand out. First, support for women's right to work is high in all groups (slightly lower among Muslim men), yet all groups markedly underestimate Muslim men's support. Especially professionals severely underestimate, consistent with a salient majority-held stereotype. Muslim students also underestimate male peers' support and correctly anticipate outsiders' underestimation, indicating relatively accurate meta-beliefs about public misperceptions. Second, support for maternal employment is lower across the board. A majority of Muslim students (62%) agree that a young child is likely to suffer if the mother works, compared with 40% among all students, signalling stronger gendered household roles. Misperceptions on this statement are less pronounced. Third, Muslim women's labour outcomes correlate with each belief layer. Muslim women report that employment is more likely with a preschooler when: (1) they personally disagree that a child suffers if the mother works (first-order, preference channel); (2) they believe other Muslim women also disagree (second-order, community-norms channel); and (3) they expect the general public sees Muslim men as supportive of women's work and Muslim women as not believing that children suffer (third-order, smaller expected stereotypes). These patterns do not hold for Muslim men, underscoring intersectionality. Overall, anticipated views-within community and from the majority-systematically relate to Muslim women's labour choices. Beyond standard anti-discrimination tools, information and signalling interventions that correct public misperceptions and make intra-community support visible can reduce meta-stereotype-induced avoidance. Shifting perceived community norms (e.g., role-model visibility) targets the second-order channel. Addressing these layered belief distortions can improve female labour outcomes in Muslim and other minority communities bound by norms and seen through stereotypes.

Misperception of Gender gaps in Performance and Persistence: The Role of Task Stereotype

Emily Tran, University of Melbourne

Misperceptions- beliefs that are misaligned with reality- play a significant role in shaping and perpetuating distortions in labor market outcomes (Bursztyn and Yang, 2022). This paper investigates how such misperceptions about performance and persistence emerge in gender-stereotypical domains and whether they are more pronounced in some domains than others. Across two incentivized experimental studies using real-effort tasks, we examine whether individuals hold gendered misperceptions in their own beliefs (first order beliefs) and in their beliefs about others' perceptions (second-order beliefs), and how these patterns vary across tasks. Study 1 focuses on beliefs about performance. We find no gender differences in actual performance in most tasks, yet individuals hold gendered beliefs that align with stereotypes: they perceive the female-typed task as favoring women and the male-typed task as favoring men. These misperceptions vary with both gender and task stereotype: men and women each tend to overstate their own gender's advantage in

gender-congruent tasks. Individuals who misperceive in their first-order beliefs are also more likely to misperceive in their second-order beliefs, typically in the same direction. Study 2 examines beliefs about persistence following positive and negative feedback. Despite no actual gender gaps in persistence, many participants expect such gaps across both tasks and feedback conditions. These misperceptions are not driven by stereotypes alone- conditional on being incorrect, participants are about equally likely to favor men or women- but they do depend on the interaction between task stereotype and the decision maker's gender: individuals tend to over-favor their own gender in gender incongruent tasks. Together, the studies provide evidence of misperceptions in both first and second-order beliefs that vary with both gender and task stereotypes. These findings highlight that efforts to reduce gender bias in gender-stereotypical domains must address both individual and social layers of misperception.

Experiencing unfair chances reduces investments in ambiguous assets

Xueting Wang, RMIT University

Discrimination leads to disadvantageous financial, labor market, academic, and social outcomes for those affected, driven by the decisions of those who discriminate. Inherent in the experience of discrimination is unfair treatment but little is known about whether experiencing unfair treatment subsequently alters own decision-making in unrelated contexts. Using a novel experiment, we examine how experiencing unfair treatment affects investment decisions. We introduced unfair treatment by varying the probability of receiving a small performance bonus (5% vs 95% chance) that determines participants' relative earning ranking. We found that after experiencing unfair treatment (without explicit discrimination), participants invested less in ambiguous assets but their investments in risky assets remained unchanged. We further explored whether wealth, gender, and previous experience of discrimination mitigate this reduction in investment. Our findings highlight the broader economic consequences of unfair treatment, emphasising its potential to shape financial decision-making.

How to Boost Revenues in First-Price Auctions? The Magic of Disclosing Only Winning Bids from Past Auctions

Peter Katuscak, RWTH Aachen University

Consider a long-term auctioneer who repeatedly sells identical or similar items and who might disclose selective information about past bidding. We present a theory that yields different predictions about bidding behavior depending on the information bidders are provided with, and then test it using a lab experiment. We focus on the disclosure of all bids and of winning bids only. Our theory is based on a selection bias: some of the bidders who are presented with historical winning bids mistakenly best-respond to that distribution, failing to realize that winning bids are not representative of all bids. In the steady state, this bias results in higher bids and auction revenue in comparison to the case when all bids are presented. Our experimental test confirms the qualitative predictions of the theory. On the theory side, our findings challenge the predictive power of Bayesian Nash Equilibrium based on rational bidders. On the market design side, they underline the role of historical market information as a key design choice.

Deterring Collusion with Fines in Auction Experiments

Tiffany Tsz Kwan TSE, University of Osaka

This study investigates the effectiveness of fines in deterring collusion in one-shot first-price and second-price sealed-bid auction experiments. Building on the experimental framework of Agranov

and Yariv (2018), which showed that cheap-talk communication and side payments significantly increase collusion, we examine how fines influence collusion when pre-auction communication and post-auction transfers are allowed. Our experimental design introduces treatments that vary the presence of fines and the auction format. Specifically, we design four treatments: (1) a baseline with no fines in the first-price auction, (2) fines with a 25% detection probability in the first-price auction, (3) a baseline with no fines in the second-price auction, and (4) fines with a 25% detection probability in the second-price auction. The results show that introducing fines reduces collusion. Notably, fines effectively deter collusion in both weak and strong cartels. These findings enhance our understanding of how fines can be used to deter collusion in auction settings where pre-auction communication and post-auction transfers are possible.

Information Provision in Private-value Auctions

Xinghao Yao, The University of Queensland

In many auctions, from real estate to online marketplaces, people form valuations from incomplete or selectively disclosed information. We experimentally study how information provision shapes bidding and market outcomes in second-price sealed-bid (SPSB) private-value auctions where each item's value equals the average of five attributes, but participants observe only a subset. A substantial share of participants place bids above their own private values across all information conditions. Relative to full information, selectively revealing attributes induces more conservative bids, yielding higher allocative efficiency and lower revenue. By contrast, revealing attributes at random leads participants, especially those who initially underestimate, to raise bids above their own estimates, producing the lowest efficiency. These results show that bidding bias persists even with rich disclosure and that the form of disclosure systematically trades off revenue and efficiency, offering practical guidance for auction and platform design under incomplete information.

The Missing 'Loser's Curse': Experimental Evidence on Belief-Based Models in Common-Value Auctions

Antonio Rosato, Deakin University

Models based on mistaken beliefs, such as Cursed Equilibrium and Level-k Thinking, are among the leading explanations for the Winner's Curse observed in common-value auctions. We argue that these models predict a Loser's Curse in other auctions formats. Using an experiment with a within-subject design, we test for the presence of both curses in uniform-price auctions with common values. At the aggregate level, we find evidence of a strong Winner's Curse, but no evidence of a Loser's Curse. These aggregate findings cast some doubts on the ability of belief-based models to fully explain the Winner's Curse. Indeed, at the individual level, the behavior of most subjects is better described by Joy of Winning and Quantal Response Equilibrium. We also find suggestive evidence of failures of contingent thinking: subjects behave closer to the rational benchmark in a non-strategic task when the relevant contingency is made more salient.

Input Competition vs. Pollution: The Effect of Mining on Agriculture in Africa

Hosam Ibrahim, University of Minnesota

I study how nearby mining affects agriculture through labor-market and environmental shocks using two complementary designs. First, a 2x2 two-way fixed effects difference-in-differences around the 2012 opening of a large gold mine in Tanzania shows that farms within 15 km reduce family labor and planted area, without changing hired labor. Labor productivity increases though, and yields do not decline, consistent with input reallocation rather than a fall in total factor productivity. Second, pooling

repeated cross-sections from Ethiopia, Tanzania, and Uganda, I estimate a staggered event study for different distance bins from a mine. Near-mine farms exhibit persistent post-opening reductions in family labor beginning about two years after opening; more distant farms show little change. In both designs, output losses are explained by lower inputs, not lower efficiency, and suggestive evidence on soil and rooting conditions show no short-run deterioration due to mine openings. Together, the evidence points to labor-market competition as opposed to pollution as the dominant short-run mechanism through which mining affects nearby farms, with effects concentrated closer to the mine.

Allocation Rules and Market Power in Carbon Trading: Evidence from Laboratory Experiments

Wei James Chen, National Taiwan University

This paper presents laboratory experimental evidence comparing carbon taxes and emissions trading schemes under conditions of market power. We implement a multi-stage design with 15 firms (3 dominant, 12 competitive), real emission accounting, and a continuous double auction market for permits. The treatments vary by policy regime (tax vs. trading), stringency (low, medium, high levels), and allocation rules (equal vs. grandfathering). Our results show that both policies reduce emissions, but trading achieves larger reductions while taxation preserves greater welfare. Equal allocation of permits improves both surplus and emission outcomes relative to grandfathering, mitigating monopsony distortions in price formation. These findings highlight the critical role of market structure and allocation rules in determining the effectiveness of carbon pricing instruments, with implications for the ongoing global expansion of emissions trading and tax systems.

Coordinating on an environmental path: theory and experiment

Inigo DeJuan-Razquin, University of Technology Sydney

In this paper, I present the results of an original laboratory experiment in which subjects make investment in two types of capital (k1 - lower emissions but less productive - and k3 - higher emissions but more productive -) and consumption in a model of economic growth with and without dynamic environmental externalities. Preliminary results reveal that both capital levels stabilize or converge to significantly above steady state levels, even at the expense of decreased consumption (i.e. lower experiment compensation for participants). Despite the highly complex nature of the experiment, participants perform better than expected in the degree to which they should internalize the emissions externality (i.e. relative allocations between k1 and k3 capital types). This paper contributes to the highly scarce research in experimental macro-growth. Lei and Noussair (2002) found that market feedback in Decentralized Markets (DM) helps people solve the Ramsey growth model in experimental laboratory settings more effectively - when compared to the Social Planner SP (i.e. each subject managing the entire economy by themselves). In my paper I expand the simplest version of the Ramsey model that Lei and Noussair (2001) used: (i) making it an environmental model with dynamic emissions externalities (using a core environmental damage factor from the Nordhaus DICE model), (ii) using two types of capital (instead of one) with productivity/emissions trade-off , (iii) changing the variable space from discrete to continuous in the lab, and (iv) reducing the capital depreciation rate from 100% to 10%, among others. Changing the depreciation rate to the more realistic 10%, by enhancing the inter-temporal complexity of investment, makes the experiment highly challenging for participants on the one hand, but generates rich data to test the hypothesis under a realistic model. I study in the lab whether price signals in decentralized markets help agents converge closer and faster to the theoretical steady states of the system, with respect to participants acting as the social planner of the economy. I implement a between-within -subject design where participants make decisions over (quasi) infinite horizons with a random termination rule of 7% after each period,

in the SP and DM treatments (within dimension, in random order), and with/without environmental impacts in the system (between dimension): whether temperature stock reduces productivity through the Nordhaus Omega factor.

An Experimental Comparison of Cap- and Intensity-based Pollution Markets

Lana Friesen, The University of Queensland

Markets are an increasingly popular regulatory choice to cost effectively control negative externalities. Traditionally, market designs have employed a cap-and-trade format that places an absolute limit on the quantity of emissions. In contrast, many new schemes—including the world's largest in China—limit the aggregate emissions intensity of production. This article theoretically and experimentally compares the relative performance of these two regulations. We design a novel laboratory experiment, where firms choose both output and allowance exchange. Consistent with our theoretical predictions, we find that employing an intensity-based market rather than an equivalent cap-and-trade scheme significantly increases aggregate output, average allowance prices, aggregate abatement, and decreases industry profits. Contrasting sharply with predictions, however, both the volume of allowance trades and final allowance holdings are virtually indistinguishable between the two scheme types at both the aggregate and individual firm type levels. This implies that the insufficiency of trades has unexpected negative implications for the relative cost-effectiveness of an intensity-based market compared to a standard cap-and-trade approach.

Impact of prize structure on strategies in all pay contests: An Experiment

Yang Liu, RMIT University

In many settings, such as bonus and promotion tournaments within firms, R&D races, political campaigns, and sports, agents compete for prizes by exerting costly effort. A central design question is how prizes should be allocated to best incentivize effort. Theory predicts that with concave effort costs, a winner-take-all (WTA) contest maximizes effort, whereas with convex costs, the optimal rule is the "punish-the-bottom" (PTB) scheme, where all but the lowest-ranked agent receives an equal prize. We test these predictions experimentally using a 2 Ã 2 between-subject design, where we orthogonally vary the allocation of prizes (WTA vs. PTB) and the cost function (concave vs. convex). Additionally, because the unique symmetric equilibrium in these all-pay contests entails continuous mixed strategies (Barut & Kovenock, 1998), standard elicitation of point-valued effort choices obscures equilibrium mixing. To address this, we introduce a within-subjects variation: participants choose a single effort level from an interval in the first part and subsequently specify the entire effort distribution in the second part using a distribution builder interface. This design enables direct observation of endogenous mixed strategies and provides a stringent test of theoretical predictions under alternative prize structures and cost curvatures.

Handicapping in Two-Period Contests

Qin Wu, RMIT University

We study how handicapping schemes in dynamic contests affect effort provision when workers differ in ability and managers cannot perfectly identify the strong and the weak. Using the framework of Ridlon and Shin's (2013) two-period Tullock contest with heterogeneous players, we examine the optimal handicapping rules when the second-period handicap must depend on the outcome of the first-period contest. Their model predicts that favouring the loser from the first-period contest maximises total effort when the ability gap is large, whereas favouring the winner is optimal when abilities are similar. We test these predictions in a laboratory experiment with a 2Ã2 between-subject

design that varies (i) the ability gap between a high- and a low-ability player (small vs large) and (ii) whether there is no handicap or a theoretically optimal handicap that is committed ex ante. The data broadly align with the theoretical benchmark. In both ability environments, optimal handicaps raise total effort relative to the corresponding no-handicap baselines. With a small ability gap, favouring the winner increases first-period effort and reduces second-period effort, as predicted. With a large gap, favouring the loser raises total effort through a stronger second-period response. However, both types reduce effort in the first-period contest, producing a race-to-the-bottom effect not anticipated by the model. Conditional analyses further show that effort concentrates on the recipient of the handicap and often collapses for the non-recipient when the handicap is mis-targeted.

(Sub)optimality and Remedies of the Majority Rule in Team Contests: Theory and Experiment

Changxia Ke, Queensland University of Technology

This paper studies how prize design can be used to maximize total effort in team contests composed of multiple pairwise battles among homogeneous players within teams. The contest organizer allocates rewards based on battle outcomes, subject to a budget-balance constraint. We show that the simple majority rule is generally suboptimal when teams differ sufficiently in strength, but optimality can be restored by incorporating additional design instruments. When the prize allocation rule must be identity independent, optimality is achieved by introducing a tie margin; when it may depend on team identity, a head start for the weak team serves the same purpose. Although the proportional rule is never theoretically optimal, it can elicit higher total effort than the majority rule when asymmetry between teams is large. We test these predictions in a controlled laboratory experiment and find strong empirical support for our main theoretical predictions. Notably, we also find that under both the proportional rule and the majority rule with a tie margin, total effort is higher when the prize is indivisible rather than divisible, echoing the optimality of winner-take-all schemes documented in the literature.

Quasi-exponential discounting

Stephen Cheung, The University of Sydney

Alternatives to the standard model of time preference typically relax the assumption of an exponential discount function while retaining the framework of discounted utility. We report novel behavioural data inconsistent with this approach. Illustrating this, we estimate highly significant "present bias", despite our data exhibiting stationarity. The paradox is resolved by relaxing discounted utility itself to allow discounting to be context dependent. We propose quasi-exponential discounting (QED), a fixed penalty applied to all episodes of delay, as a particularly simple model of this type and show that it provides an excellent approximation to the best fit to our data.

Efficient and Valid Within-Subject Designs

Gwen-Jiro Clochard, ISER - Institute of Social and Economic Research

Within-subject designs - where participants are exposed to both treatment and control conditions - offer significant advantages over standard between-subject designs. They allow researchers to observe the full joint distribution of outcomes and substantially increase statistical power. However, such designs are used far less frequently than between-subject designs in the experimental economics literature. One possible explanation for this limited adoption is that within-subject designs rely on stronger assumptions. In particular, they require temporal stability - the assumption that potential outcomes do not depend on the period in which treatment is assigned - and causal transience - the assumption that past or future treatment assignments do not affect current potential

outcomes. In standard within-subject designs, these assumptions cannot be tested directly, making it difficult to interpret the results. In this paper, we argue that introducing a pure control group - a subset of participants who are never treated - can both test the validity of these assumptions and correct for potential violations, thereby allowing recovery of the full distribution of treatment effects. We also discuss the trade-off introduced by including a pure control group: while increasing its size strengthens confidence in the findings, it simultaneously reduces the within-subject design's advantages in terms of statistical power and distributional analysis.

uproot: A Software Framework for Behavioral Experiments

Max R. P. Grossmann, The University of Melbourne

Existing popular software frameworks for conducting behavioral experiments-both online and in the lab-suffer from limitations including architectural inflexibility and limited feature sets that restrict their utility. They are closed source or do not encourage outside contributions. To overcome these restrictions, we introduce "uproot": a scalable, genuinely open-source framework for developing and conducting behavioral experiments. uproot revolutionizes data storage with an append-only log, ensuring data persistence and allowing the use of arbitrary data types. uproot provides many additional important features for modern experiments. For example, Likert scales are provided out of the box, file uploads are possible, and internal pages can be customized easily. With capabilities such as page repetition, random page orderings, indefinitely repeated sets of pages, management of individual subjects, dynamically sized sessions, and convenient async-first APIs, uproot overcomes many of the limitations of existing frameworks. Our framework is suited for both small-scale and large-scale studies, including surveys and experiments in which participants interact with each other. We encourage community-driven development.

Revisiting Gender Differences in Volunteering for Non-Promotable Tasks

Kalyani Chaudhuri, Ashoka University

Recent evidence shows men are less likely than women to volunteer for non-promotable tasks in organizations, which perpetuates gender gaps in career advancement. We design a series of lab experiments to identify the drivers of this gender difference in strategic volunteering. Specifically, we control whether participants receive real-time feedback about others' decisions in the volunteer's dilemma game. We find that men and women volunteer at equal rates in the absence of feedback, suggesting that they have similar prior beliefs about others' likelihood of volunteering. However, feedback that other players are free-riding causes women to volunteer more often than men, which suggests gender differences in the updating of beliefs. We also collect survey measures of volunteering that mimic non-promotable tasks in the workplace, and find that women volunteer more than men only in female-stereotyped tasks due to beliefs that others will not volunteer. Overall, our results point to beliefs as a crucial driver of gender differences in volunteering for non-promotable tasks.

When Gender Is Seen: How Visibility Shapes Collaboration

Sukran Dinc, York University

Collaboration drives productivity in firms and research, but also creates uncertainty about individual contributions. We study an experimental labor market where candidates choose whether to collaborate or work alone before being evaluated for hiring. In a gender-blind baseline, women and men are equally likely to collaborate, despite widely held beliefs that men collaborate less. When gender becomes visible to evaluators, men become significantly less willing to collaborate, while

women's behavior remains unchanged. The pattern is consistent with social norms rather than strategic incentives: men adjust their behavior to conform to gendered expectations about collaboration. Our findings overturn common intuition about gender visibility and show that gender fundamentally shapes collaboration and labor-market outcomes.

Experience in the Same-Gender Environments and Low-Promotability Tasks

Duk Gyoo Kim, Yonsei University

It is well documented that women are more likely to voluntarily take on, and are often expected to handle, tasks that are essential for group functioning but unlikely to lead to promotions (Babcock et al., 2017). These low-promotability tasks (hereafter, LPTs) may contribute to the gender pay gap in mixed-gender environments. At the same time, prior work has shown that the overall tendency to take on LPTs in female-only groups does not differ significantly from that in male-only groups. Together, these findings motivate our study. Our primary research question is whether women's behavior regarding LPTs in same-gender environments persists when they transition to mixed-gender settings. We address this question to shed light on the potential long-term impact of same-gender college education, an area that has received relatively little attention. In a situation requiring a group member to take on an LPT, a woman from a same-gender environment (S) may feel no social obligation to volunteer, whereas a woman from a mixed-gender environment (M) may either feel compelled or be expected to do so. Over time, S may not internalize the expectation to take on LPTs, while M may. Would this difference in prior experience shape their future behavior in mixed-gender environments? And ultimately, could it result in S earning more than M? To investigate these questions, we conduct a controlled laboratory experiment that sequentially combines two experimental conditions from Babcock et al. (2017). If the experience of not doing LPTs persists when women move into mixed-gender environments, this would represent an unexplored benefit of same-gender education for women. If, however, their behavior changes once they enter mixed-gender settings, then the policy of maintaining same-gender environments may still serve as a tool to reduce the gender pay gap. (*The experiment will be conducted in October and November. IRB approval in progress.)

Political Polarization, Wage Inequality and Preferences for Redistribution

Lionel Page, The University of Queensland

We investigate how beliefs about wage inequality impact preferences for redistribution with 10,000 people across six high-income countries. Using nationally representative, randomized survey experiments and a novel distribution builder tool, we elicit detailed beliefs about wage inequality and examine the impact of providing accurate information on support for redistribution and beliefs. We find that, on average, respondents underestimate wage inequality, and that information treatments have minimal effects, with the exception of those on the far right, who exhibit large increases in support for redistribution. These results suggest appeals emphasising wage inequality to far-right voters are likely to be effective.

A Bliss Point Model of Distributional Preferences

Moira Gidseg, University of Arizona

I propose and empirically evaluate a new "bliss point" model of social preferences, in which one's utility decreases in the distance between others' payoffs and a bliss point defined as a linear function of one's own payoff. Using new experimental data, I estimate individual-level parameters for this model and three alternatives (CES, inequality aversion, maximin), and then compare predictive

accuracy. Preliminary results suggest that the bliss point model is the best predictor for 15-32% of subjects, depending on the task.

Zero-Sum Views Reduce Support for Redistribution Across Borders

Diego Marino Fages, Durham University

This paper provides causal evidence on how zero-sum views shape support for cross-border redistribution and openness to trade and migration. We report results from a pre-registered two-by-two experiment with a broadly representative sample of 2,116 UK adults. In the first treatment, participants were primed to adopt stronger or weaker zero-sum mindsets in general social and economic interactions, without reference to redistribution. Inducing a stronger general zero-sum mindset reduces donations to international anti-poverty organizations and marginally lowers stated support outcomes. In the second treatment, participants were informed of their position in the global income distribution. Learning one's relative global advantage fully offsets the effects of zero-sum priming. The findings show that information on people's global rank can counteract exclusionary zero-sum beliefs amid rising nationalist and protectionist rhetoric.

Preferences over the Timing of Redistribution Policies

Mallory Avery, Monash University

Across the world, redistribution policies vary wildly both in terms of their popularity and their implementation. Motivated by real-world redistribution policies and the literature on loss aversion for others, we study how timing - both of information about the redistribution and of the redistribution itself - relative to the work event affects preferences for redistribution. Using a third-party redistribution design, we study how spectators' willingness to pay to redistribute from a status quo winner-take-all payment scheme to a piece-rate structure is impacted by whether the workers will be informed of the redistribution before or after the work event and whether the workers will or will not learn about their pre-redistribution earnings. Furthermore, we consider the impact on workers' performance as well as workers' emotional responses and spectators' beliefs about workers' emotional responses to the redistribution. The results will be informative for how redistribution policy structure affects public perception of such policies, possibly suggesting implementation strategies that will be more positively received. We are currently in the process of data collection and analysis.

Style over Substance in Market Design

Lucas Coffman, Boston College

The importance of substantive properties of matching algorithms has been well-documented. Less is known about the importance of how the algorithms are presented and communicated to market participants. We source different styles of presentation from the academic literature - e.g. describing the algorithm, explaining stability, providing advice, so on - as well as tactics from over 100 school system websites around the world - e.g. video explainers, example stories, so on. Across hundreds of experimental treatments with thousands of participants, we assess the importance of "style" using the model completeness measure in Fudenberg et al 2022. We benchmark the results comparing how much behavior changes due to "substance", i.e. when the matching mechanism is deferred acceptance versus a Boston mechanism. We discuss the importance of studying how mechanisms are communicated and understood and how those learnings can be incorporated into mechanism design.

Complexity Beyond Incentives: The Critical Role of Reporting Language

Manshu Khanna, Peking University

Many assignment systems require applicants to rank multi-attribute bundles (e.g., institution-major-tuition). We study whether this reporting task is inherently difficult and how reporting interfaces shape accuracy and welfare. In laboratory experiments, participants hold utility over attributes that generate lexicographic, separable, or complementary preferences over programs. We compare a full programs ranking, two simplified interfaces that elicit attribute rankings (lexicographic nesting and weighted attributes), a sequential choice version of serial dictatorship that is obviously strategy-proof, and a baseline that uses full programs ranking, but rewards pure accuracy and not the allocation by a mechanism. Four results emerge. First, substantial misreporting occurs even without strategic incentives in baseline and rises with preference complexity. Second, serial dictatorship induces additional mistakes consistent with misperceived incentives. Third, simplified interfaces do not raise-and sometimes reduce-accuracy, even when they match the structure of preferences. Fourth, sequential choice delivers the highest accuracy, improves efficiency, and enhances fairness (lower justified envy). The findings caution against restricted reporting languages and favor sequential choice when ranking burdens are salient.

Consolidating land through rental swaps: Experimental evidence from Uganda

Tom Wilkening, University of Melbourne

Can market design be used to improve agriculture rental markets and facilitate welfare improving trades? To explore this question, we map ownership of agriculture land in Uganda and identify bilateral land swaps that can consolidate agriculture land. We elicit the prices at which farmers are willing to participate in these rental swaps using an incentive compatible trading mechanism that, in some cases, induces real trades. We find that, on average, farmers are willing to pay just under 10% of the rental price to farm consolidated plots. These consolidation gains exceed estimated trade frictions in many transactions resulting in a positive expected surplus in 40% of the land-consolidating swaps that we offer. Our results suggest that small-cycle swaps can improve welfare outcomes in locations where land is fragmented, and we use our data to evaluate potential larger market redesigns. Simple matching algorithms that use fixed prices based on differences in land sizes can capture just over 50% of the potential surplus identified. Both subsidies and more complicated mechanisms that incorporate farmer-reported prices improve outcomes relative to this simpler design.

Marital Arrangement and Spousal Cooperation

Uzma Afzal, RMIT

We present three lab-in-the-field studies investigating systematic heterogeneity in cooperative decision-making across spouses in arranged and love-matched marriages in Pakistan, where the former is the tradition and the latter is associated with modernization. In Study 1, we engaged married couples in a one-shot, two-person, sequential public goods game, in which we applied the strategy method to the second mover. Using hierarchical clustering to analyze the strategy data, we categorized spouses into cooperative types and found that spouses in love-matched marriages are significantly more likely to be unconditionally cooperative. Spouses in love-matched marriages are also significantly more cooperative overall. In Study 2, we replicated our findings from Study 1 in a new sample of villages similarly close to a city but found that, as distance from the city increased, the love-matched effect declined. We interpreted this as suggestive evidence that there is less tolerance and support for love matches in more remote areas. In Study 2, by also engaging the spouses in games with neighbors, we established that the observed differences in cooperation between spouses

in love-matched versus arranged marriages could not be explained by the selection of unconditionally cooperative people into love-matched marriages. Finally, in Study 3, we confirmed that there is indeed a social norm prescribing arranged marriage and that this norm is stronger in more remote villages.

When Democratic Choice Meets Identity: Cooperation in a Public Goods Game with Third-Party Punishment

Hyoji Kwon, University of Hyogo

This research presents an experimental study that investigates the interplay between majority-rule decision-making and in-group bias within sanctioning systems. Using a public goods game framework, the study examines how cooperation is affected when punishers are determined through majority-rule voting rather than random assignment. Four experimental conditions are implemented, varying along two dimensions: (1) whether punishers are selected by majority-rule election or random assignment, and (2) whether a minimal group identity was induced, creating a salient in-group/out-group context for participants. By comparing these treatments, the study aims to identify the conditions under which cooperative behavior is most effectively promoted, and to assess the independent and interactive effects of institutional legitimacy and in-group favoritism. Our results reveal a powerful synergistic effect: cooperation was highest in the condition combining a majority-rule election with a shared in-group identity. This finding suggests that neither institutional legitimacy (from the election) nor group identity (from shared membership) alone was sufficient to maximize cooperation. The analysis indicates that the majority-rule-only condition failed to produce robust cooperation, likely because it lacked the relational ties and shared community bond forged through the identity formation process. Conversely, the identity-only condition-which featured a shared group identity but lacked a legitimate majority-rule process for selecting the authority-also failed to fully leverage the potential of the community bond. It appears the simple co-presence of an in-group authority, without a legitimizing procedure, was not enough to motivate sustained cooperation. Most strikingly, the combined condition revealed that out-group members (those with a different identity from the monitor elected by majority rule) significantly increased their cooperation, with contribution levels rivaling or even surpassing those of in-group members. This finding provides strong evidence that a majority-rule institution becomes most effective when it is built upon a pre-existing sense of community. Furthermore, a legitimate democratic process, when integrated with shared identity, appears capable of overcoming in-group bias, transforming a potentially parochial in-group bond into a mechanism for inclusive, universal cooperation.

Government Subsidies under Climate Risk: Provision of Rural Public Goods in the Field

Shuwen Li, WU Vienna University of Economics and Business

The increasing frequency of extreme weather events has imposed new challenges to small farmers, especially in developing countries. In this paper, we aim to find efficient ways for governments to partner with the local farmers so that modern irrigation channels can be constructed to avoid losses from drought. Using a framed threshold public goods game of loss avoidance, we test whether two different types of government subsidies increase the rate of successful irrigation system provision, disentangling the price-reduction effect and the government-leading effect. We collect data in the field using farmers in rural Henan, China. Results show that with no subsidy, the rate of provision is below the social optimum. Subsidies substantially increase the likelihood that the irrigation project is undertaken, though groups remain under-insured in relation to the social optimum. A one-to-one matching subsidy appears to be at least as effective as lump-sum grants under same budget.

Solving normative conflicts in collective action by promoting redistribution

Lata Gangadharan, Monash University

Heterogeneous returns from public good contributions within groups create a normative conflict between equality and efficiency. We propose a menu of contribution principles to group members, including one featuring a de-centralized redistribution mechanism ensuring earnings equality in exchange for fully efficient contributions. Although a majority of individuals, when in the position of an impartial observer, consider this principle to be the most appropriate and expect others to agree, they fail to act on it. Designating a leader who endorses this principle and can make non-binding recommendations enables a majority of groups to adopt it successfully. This results in full contributions and earnings equalization through redistribution from advantaged to disadvantaged members, effectively resolving the conflict.

Informational and Financial Nudges for Sugar Reduction in Non-Standardized SSBs: Field Experimental Evidence from Taiwan

Chiu-Lin Huang, National Chi Nan University

Sugar-sweetened beverages (SSBs) are a major source of added sugars globally and contribute significantly to obesity, diabetes, and metabolic diseases. Although more than 45 countries have implemented SSB taxes, these policies primarily target standardized packaged drinks. In contrast, non-standardized and customizable beverages, such as handcrafted coffees and teas, are inherently difficult to regulate because sugar content varies by preparation and individual customization. This lack of standardization makes these beverages challenging to tax and difficult to study rigorously, creating a notable gap in both policy design and academic research. In Taiwan, the bubble tea industry has expanded rapidly, with annual revenue rising from US$1.41 billion in 2009 to US$3.18 billion in 2020. National survey evidence indicates that 42% of adults consume at least one sugary drink daily, and Taiwanese teenagers consume nearly twice as much as their peers elsewhere, contributing to an additional US$2.83 billion in SSB-related healthcare costs annually. This study examines whether informational and financial nudges can effectively reduce added sugar in non-standardized beverages. We conduct a field experiment in a high-volume bubble tea shop in Taipei, implementing three point-of-purchase informational displays (Figures 1-3): Figure1: Numeric-only information (sugar grams + % daily recommended intake) Figure2: Traffic-light imagery + numeric information Figure3: Sugar-cube imagery + numeric information We additionally test a no-added-sugar subsidy (US$0.2 discount), aligned with Taiwan's legal reusable-cup incentive, as a feasible alternative to taxing non-standardized SSBs. The experiment consists of a one-month baseline period, three one-week informational interventions, and a two-week subsidy phase with and without the most effective information. Across all stages, we expect to observe over 15,000 beverage purchases and survey approximately 700 customers. Preliminary evidence from 10,582 walk-in purchases shows that before the interventions, 45.34% of customers choose full sugar, 9.11% half sugar, 23.29% less sugar, and 22.25% no sugar. Controlling for rainfall, order size, reusable-cup use, and weekend effects, all three informational nudges significantly shift consumers toward lower-sugar options. The numeric-only display reduces full-sugar selections by 4.1 percentage points and increases no-sugar choices by 3.0 points. Traffic-light imagery elicits the strongest response, lowering full-sugar choices by 5.1 points and increasing no-sugar choices by 4.0 points. Sugar-cube imagery also produces substantial improvements, reducing full-sugar selections by 4.9 points and increasing no-sugar choices by 3.5 points. Customers who bring reusable cups are more likely to choose lower-sugar beverages. Rainfall is negatively associated with sugar levels, suggesting that consumers opt for lighter beverages on rainy days. Results for the subsidy intervention and survey will be available by December 2025. Overall, this field experiment provides early evidence that simple, low-cost

informational nudges can significantly reduce added sugar in non-standardized SSBs. These findings offer practical insights for policymakers seeking scalable interventions that complement or substitute traditional SSB taxes.

Does the informational promotion nudge credit card repayment?

Simin Tao, University of Technology Sydney

Credit card debt is a prevalent financial issue, and the behavioural factors influencing repayment decisions are critical for developing effective financial policies and consumer protection strategies. This study investigates the effects of informational prompts and repayment options on credit card repayment through a controlled experiment with a 3x2 factorial design. The study recruits 600 Australian participants via Prolific who meet the following eligibility criteria: (1) they are at least 18 years old, (2) they own at least one credit card, and (3) they have made at least one minimum payment in the past 24 months and failed to pay their credit card balance in full at least three times within the same period. Participants are randomly assigned to one of six experimental conditions to assess how variations in repayment presentation and informational framing affect their chosen repayment amounts. Two key independent variables: Informational Prompts, No Prompt (Control): Participants receive no additional informational intervention and make repayment decisions based solely on their existing preferences and financial constraints. With Prompt (Treatment): Participants see the following message before making their repayment choice: "The more you repay now, the faster you'll clear your debt." This statement is designed to encourage higher repayment amounts by making the long-term benefits of reducing debt more salient. Repayment Slider Options (3 Levels) Zero Repayment (0): The repayment slider is pre-set at zero. Midpoint Repayment (Middle of balance): The slider starts at a pre-set middle value between the minimum and the full balance. This condition tests whether providing a midpoint reference influences participants' repayment decisions. Maximum Repayment (Full balance): The repayment slider is pre-set at the maximum outstanding balance, which may encourage participants to consider paying off their full debt obligation rather than making partial payments. Each participant is exposed to only one combination of these factors, creating six unique experimental conditions. By manipulating both the framing of repayment decisions (via informational prompts) and the default option structure (via slider settings), this study provides insights into how consumers navigate credit card repayment choices when presented with different informational and intervention. Two key research questions have been addressed: Does the presence of an informational prompt encourage higher repayment amounts? Prior research suggests that simple reminders and nudges can significantly influence financial behaviour. We find that a simple information nudge can significantly reduce credit card debts, by discouraging individuals to pay the minimum repayment amount. How do default repayment options impact decision-making? The way repayment options are structured may anchor participants to particular repayment amounts. We find that changing the anchor by setting different positions of the initial repayment slider can significantly influence individuals' credit card debt repayments. The broader policy implications of this research are substantial. Credit card debt is a major concern for financial regulators, consumer protection agencies, and policymakers. If certain default settings and informational nudges are found to be effective in promoting higher repayments, financial institutions and regulatory bodies could consider implementing such interventions to improve consumer financial well-being.

Rights and Wrongs: The Effect of Labor Rights Information on Workplace Decisions among Temporary Migrants

Monica Beeder, University of Southampton

Temporary migrant workers power key industries, including hospitality, tourism, and agriculture, yet many are vulnerable to harsh working conditions due to factors including limited knowledge of their working rights, limited language skills, and financial insecurity. For instance, in Australia, young temporary migrants contribute over $3.2 billion annually to the economy and are the primary source of labor in industries such as horticulture, where it is common to work long hours for less than minimum wage. One approach to tackling this issue are information campaigns, aimed to increase awareness about rights to empower temporary migrants to select into better working conditions, if they are able to make these choices. Whilst these campaigns are frequently used in practice, causal evidence regarding their effectiveness in different settings is limited. In this paper, we present results from a randomized controlled trial testing the effect of a working rights information campaign on working conditions in a longitudinal study in Australia. We confirm findings from previous research, showing that those who are unaware of their rights, have limited English skills, and lack financial security face worse employment conditions. Our study boosted knowledge of working rights. This increased awareness is associated with improvements in working conditions over time, but other factors, such as systematic dependency on employers, raise caution regarding the effectiveness of information campaigns. We discuss how these results can inform use of information campaigns in practice, and how working conditions may be improved further through campaigns targeted not only workers but also other parties involved, including employers and consumers.

Intrinsic preferences in prediction markets: An experiment

Senran Lin, Southwestern University of Finance and Economics

We investigate how intrinsic preferences—non-instrumental motivations tied to the outcomes of forecasted events—shape the performance of prediction markets. Prediction markets are widely regarded as efficient mechanisms for aggregating dispersed information, and under standard rational-expectations models, prices should reflect collective beliefs about future events. This efficiency, however, depends on traders responding solely to informational and monetary incentives. In practice, traders often bring behavioural biases that can systematically distort market prices.
A particularly relevant bias is motivated reasoning, whereby individuals interpret or act on information in ways consistent with their preferred outcome. When participants care intrinsically about how the forecast resolves—because of political identity, moral convictions, or personal stakes unrelated to payoffs—this non-monetary utility can lead them to place trades that depart from profit-maximising strategies. Such behaviour not only reflects expressive preferences but also undermines the market's ability to generate accurate forecasts.

We test this mechanism using a controlled laboratory experiment in which participants traded in prediction markets under conditions that varied the presence of intrinsic preferences. When participants held such preferences, mispricing increased significantly relative to baseline markets without expressive stakes. These distortions persisted even though traders could have increased expected earnings by setting their preferences aside and trading solely on information.

These findings highlight a key limitation of prediction markets: while they are powerful tools for information aggregation, their accuracy can be compromised when intrinsic preferences compete with financial incentives.his paper presents an axiomatic characterization of the Grether (1980) model and Bayes' rule that is centered around the preservation of the Monotone Likelihood Ratio Property (MLRP). I show that Bayesian updating can be characterized by the preservation of MLRP from signals and the martingale property, while the Grether (1980) model is characterized by the preservation of MLRP from signals and priors. This framework allows us to identify a class of non-

Bayesian updating rules where we can obtain useful comparative statics across different signal realizations and prior beliefs in canonical belief updating problems. I also conducted an experiment to test the axioms that can be used to characterize Bayes' rule to identify why people are non-Bayesian. I find that subjects broadly comply with the preservation of MLRP, providing validation for the widely used Grether (1980) model.

An Axiomatic Model and Test of Grether (1980) and Bayes' Rule

Kenneth Chan, National University of Singapore

This paper presents an axiomatic characterization of the Grether (1980) model and Bayes' rule that is centered around the preservation of the Monotone Likelihood Ratio Property (MLRP). I show that Bayesian updating can be characterized by the preservation of MLRP from signals and the martingale property, while the Grether (1980) model is characterized by the preservation of MLRP from signals and priors. This framework allows us to identify a class of non-Bayesian updating rules where we can obtain useful comparative statics across different signal realizations and prior beliefs in canonical belief updating problems. I also conducted an experiment to test the axioms that can be used to characterize Bayes' rule to identify why people are non-Bayesian. I find that subjects broadly comply with the preservation of MLRP, providing validation for the widely used Grether (1980) model.

Information representation and meta-cognition of informativeness

ChienHsun Lin, National Taipei University

We examine whether people can correctly evaluate the value of information. In our experiment, we employ the canonical "balls and boxes" task, where subjects update their beliefs about the true state (box chosen) based on balls drawn from the box. Subjects receive different representations of the information, featuring varying degrees of simplification relative to the raw data. Before subjects undertake the tasks, we elicit their willingness to pay (WTP) for each representation. We first find that subjects' WTP does not strictly align with theoretical informativeness; specifically, subjects prefer to see a "proportion" (e.g., 20% orange balls), even though it is not the most informative representation. Furthermore, we find that subjects do not necessarily perform better with the representation they prefer, suggesting a lack of meta-cognitive awareness regarding how effectively they can utilize the information.

The effect of price signals on effort and productivity in a hard task

Hassan Andrabi, Centre for Brain, Mind and Markets, The University of Melbourne

Economists have long theorised that markets can coordinate individual actions toward efficient outcomes. In this paper, we test this hypothesis by asking whether market prices can reveal optimal solutions to complex problems, and whether traders learn from these prices to improve their own problem-solving performance. In a laboratory experiment, we asked participants to solve instances of the Knapsack Problem (KP), a canonical example of a complex task, while trading an asset whose payoff was tied to the best solution identified by any participant. In this setting, trading interactions generated price signals that encoded private knowledge about potential solutions to the KP, which individuals could exploit to potentially improve their own performance. We compare the performance of those who traded in this market to two benchmarks in which trading was not permitted: one in which participants received an explicit signal disclosing the optimal solution to the KP, and a control condition in which no information was provided. We find that traders in the market learned from prices to improve their own problem-solving performance, achieving similar levels of productivity to those who

received the optimal solution signal and outperforming those who had no signal. However, the informativeness of prices diminished with the complexity of the underlying problem. Our findings illustrate the role of market prices to provide information signals that guide individual effort and productivity, even when solutions are computationally difficult to identify.

Augmenting Collective Intelligence: The Effects of AI Cognitive and Coordination Support on Team Performance

Zebang Deng, Nanyang Technological University

Artificial intelligence (AI) technologies are increasingly being used to support human teams, yet we have limited causal evidence on how AI assistance affects team dynamics and performance. This paper presents an experiment that studies the impact of AI support on group collective intelligence in a laboratory setting. We form teams to work on a diverse battery of tasks and introduce AI-based suggestions as a new input into the group problem-solving process. We distinguish between two forms of AI assistance: cognitive information (task-relevant hints or knowledge) and coordination information (process-oriented advice to facilitate teamwork). To elicit teams' valuation of AI support and achieve exogenous variation in its use, we implement a group-level Becker-DeGroot-Marschak (BDM) mechanism: teams bid for AI suggestions and a random threshold price determines whether the AI advice is delivered. This design provides a continuous measure of willingness-to-pay (WTP) for AI help and leverages random assignment from the BDM to identify the causal effect of AI assistance on team outcomes. We hypothesize that AI cognitive support will improve teams' task performance by providing relevant information, while AI coordination support will improve team processes and efficiency. Our analysis examines overall treatment effects and heterogeneous impacts moderated by baseline team characteristics (e.g. prior collective intelligence or trust in AI). The findings will speak to a growing literature on human-AI collaboration and offer insights for organizations on how to best integrate AI to augment team decision-making.

Testing Digital Health Assistants: An Online Experiment

Gyula Seres, National University of Singapore

Recent advances in generative AI models enabled the creation of digital health assistants for patients. However, it remains unclear how patients - especially in the presence of behavioral biases - would utilize them. In the paper, we test a theoretical model in an online experiment. We analyze how bounded-rational patients use AI health assistants to make healthcare decisions. Our decision theory model predicts that cognitive biases lead patients to underutilize these assistants, limiting their potential to prompt high-risk patients to seek necessary care and to reduce unnecessary clinical visits among low-risk patients. Moreover, we found that bounded rational patients become less sensitive to differences in risk, and their decision to seek clinical care is determined primarily by the cost of access to healthcare rather than by the underlying health risk. We test these findings in an online experiment in which healthy participants make incentivized decisions in hypothetical health scenarios. Each scenario is a randomly generated list of symptoms, and the participant makes a binary decision about whether to seek costly medical help. They may or may not require help - their true health status is randomly generated, and more serious symptoms entail a higher probability of being sick. They suffer a loss if they are sick but decide not to seek medical help. There are two treatments. In one of them, the participant receives a noisy computer-generated signal about their true health status. The study shows that the computer-generated signal is underutilized, and decisions are mainly driven by the cost of medical help instead of health risk.

Understanding Effects of Decision Recommendation: Evidence from Large Language Models

Fengfei Sun, NUS

Algorithmic recommendations are ubiquitous in modern decision-making, yet identifying their causal effects remains challenging because recommendations typically correlate with decision-relevant information. We address this using Large Language Models (LLMs) as an experimental platform, exploiting their stochastic variation to decompose recommendation effects into two channels: signaling (conveying information) and perception (altering choice evaluation). We develop a novel experimental paradigm embedding decisions within context-processing tasks. LLMs exhibit dramatic performance heterogeneity---100% accuracy on text analysis versus 58% on visual-spatial tasks---while reporting uniformly high confidence, generating exogenous variation in recommendation quality. In a within-subject experiment with 100+ participants making 40 incentivized binary choices, we implement information-orthogonal treatments where participants know true probabilities, isolating pure perception effects. Three key findings emerge. First, adoption exhibits striking asymmetry: participants readily follow recommendations reducing uncertainty but resist trading certainty for expected value gains, even with perfect LLM accuracy. Second, recommendations substantially alter choices even without informational advantage, increasing safe-option selection by 1.6-1.8 standard deviations. Third, the dominant mechanism is caution ---recommendations systematically shift preferences toward higher minimum payoffs, intensifying in loss frames. Our results demonstrate that algorithmic guidance operates through distinct informational and psychological channels. The cautionary perception effect helps reconcile algorithm aversion (reluctance to abandon safe options) with algorithm appreciation (adoption of risk-reducing recommendations). As AI-assisted decision-making proliferates, understanding these dual channels becomes essential for designing effective human-algorithm collaboration and evaluating welfare implications of automated decision support.

Automating Job Interviews and Interview Evaluations: A Field Experiment

Edwin Ip, University of Exeter

Recent technological advancements have fundamentally changed the way job interviews are conducted and are reshaping the pathways to employment. Most major employers have now automated the interview processes by having applicants carry out asynchronous interviews, where job applicants submit recorded responses via an online platform without interacting with an interviewer. At the same time, many of such employers now automate the evaluation process by using AI algorithms to assess the recorded job interviews. In this paper, we use a field experimental approach to comprehensively study how these new technologies affect job entry. Over 3,000 applicants for tech jobs are randomized into asynchronous audio or video interviews, live online interviews and a control group. We find that asynchronous interviews cause a large decline in application continuation and that this decline is significantly larger for women than men. A vignette study provides evidence that these differences are driven by economic factors such as the job-seekers' perceptions about the competitiveness based on the recruitment process used. The job interviews are then assessed by both professional recruiters and a leading AI-based recruitment platform's proprietary algorithm. We find that the AI assessment tool scores female and underrepresented racial minorities higher than white men, whereas the opposite holds for human evaluators. We track our applicants' labor market outcomes using LinkedIn one year after our initial experiment and find that the AI assessment tool, based on the interview transcripts alone, predicts subsequent employment success significantly better than human recruiters, suggesting that AI may capture soft skills and potential that humans overlook. In addition, we provide evidence that, unlike AI, human recruiters' assessments suffer from fatigue, imprecision, and anchoring. Our findings provide the first evidence that the recent technological advancements are transforming the "entry

gates" to employment: it significant deters applicants, especially female applicants; but women and racial minorities are more likely to be shortlisted, and shortlisted applicants are more likely to have better labor market outcomes.

Expected Effort Allocation Under Task Difficulty and Reward Uncertainty

Michelle Lee, The University of Melbourne

Individuals often avoid complex problems even at the cost of foregoing large rewards, yet the mechanisms underlying this avoidance remain unclear. This raises two important questions: (1) How do individuals form expectations about the amount of effort required when facing a complex task when rewards are uncertain? (2) How does expected effort allocation map onto willingness to exert effort? To address these questions, we conduct a real-effort experiment using a cognitively demanding task: solving a 6x6 Sudoku puzzle under varying task difficulty levels and reward types. Task difficulty is manipulated by varying the number of clues in a puzzle, controlling for the Sudoku solving technique, where more clues indicate easier puzzles. Reward type is manipulated by varying whether the reward for solving a puzzle correctly is fixed or uncertain ex-ante. Each participant completes two blocks of ten puzzle: one for each reward type and each block contains five easy and five hard puzzles. Block order and puzzle order are independently randomized. The experiment consists of two tasks: (i) the main Puzzle Task, followed by (ii) a questionnaire capturing individual characteristics. In the Puzzle Task, there are two stages. In the first stage, participants preview each puzzle before reporting their expected effort and willingness to exert effort. Expected effort is measured as the amount of effort an individual plans to exert to solve the puzzle within a fixed time limit, and willingness to exert effort is measured as the minimum reward they would accept to avoid solving it, elicited using the Becker-DeGroot-Marschak mechanism. For the variable reward block, participants also report their expected reward for solving the puzzle correctly. In the second stage, two puzzles are randomly selected from each block to determine participants' earnings in the Puzzle Task. Participants either solve or avoid each puzzle based on their reported willingness. Finally, participants complete a questionnaire measuring demographics, puzzle familiarity, need for cognition, attitudes towards risk and ambiguity, and logical reasoning ability. Preliminary results are fourfold. First, task difficulty increases expected effort for most participants, suggesting that expected rewards always exceed cognitive costs of solving puzzles for these participants. However, a small fraction choose to exert less effort in hard puzzles, indicating that cognitive costs exceed expected rewards, leading to lower expected effort and potential avoidance. Second, reward uncertainty affects expected effort through beliefs about expected rewards: expected effort increases with higher expected rewards across all participants. Third, the effect of expected reward on intended effort is smaller for hard puzzles than for easy puzzles, consistent with costly effort. Fourth, the relationship between expected effort and willingness to exert effort is negative for most participants, reflecting costly effort, but follows an inverted-U shape for participants whose cognitive costs exceed expected rewards in hard puzzles. By linking expected effort allocation to willingness to exert effort, this study sheds light on how beliefs about effort provision shapes aversion to complex problem-solving.

Complexity and the Use of 'Supposedly Irrelevant' Factors in Decision-Making

Ao Wang, National University of Singapore

Behavioral economics has offered extensive evidence that individuals' decisions are often shaped by factors that are supposedly irrelevant in neoclassical models-such as the status quo, dominated alternatives, or path-dependent factors. We propose an experimental research agenda to study how complexity mediates the decision relevance of these "supposedly irrelevant factors" (SIFs), guided by two broad hypotheses. First, making the economically relevant trade-offs more complex increases

reliance on SIFs, as individuals may lean on these cues when the true payoff-relevant features are harder to integrate. Second, making the SIFs themselves more complex to incorporate should attenuate their influence. We test these hypotheses in three economic environments: Experiment 1 studies binary choices between lotteries (including degenerate lotteries), where the candidate SIF is status-quo earnings serving as a reference point; Experiment 2 uses lottery choices to study decoy effects, where the candidate SIF is the presence of a dominated decoy option; and Experiment 3 examines hiring decisions that may be shaped by path-dependent and expectations-based reference points, both of which are SIFs under neoclassical assumptions.

Minimal Intelligence in the Double Auction: Logit-Choice and Reservation Utility

Brett Williams, University of New South Wales

I propose a new boundedly-rational trader-behavior model for the continuous double auction, fit in the general equilibrium (GE) framework. Traders' order placement and acceptance strategies are driven by GE-adaptions of two classic mechanisms: within-period reservation prices (here utilities), and price-acceptability beliefs. Reservation utility associated with the side of entry is used to gauge immediate acceptability of contra-side orders while contra-side reservation utility informs the trader of admissible orders for placement. Traders use probabilistic choice over their set of admissible orders, based on beliefs formed about the acceptability of each order. The logit choice parameter allows the precision of trader order placement to range from zero intelligence to perfectly maximizing given beliefs. Simulations report market performance for this model, finding promising measures of efficiency and evidence of convergence to nearly Pareto optimal allocations. The model is then applied to experimental data from Williams (2025a) and Friedman et al. (2025), capturing an encouraging amount of the utility-reducing behavior and order placement in general.

Framing Carbon Offsets: Moral Spillovers and Decision Context in Sustainable Supply Chains

Danielle Kent, The University of Sydney

Abstract: As firms increasingly commit to decarbonizing their supply chains, understanding the psychological and contextual drivers behind carbon offset investment decisions has become critical. This study employs a controlled experimental design to examine how message framing and decision context influence individuals' willingness to allocate resources toward carbon offsets, an emerging ethical responsibility within global supply chains. Participants were randomly assigned to either an altruistic allocation task or a return-oriented investment task, with carbon offsetting framed in either an optimistic or a neutral tone. The results indicate that optimistic framing significantly increased offset allocations, particularly when decisions were made in the altruistic context, where decisions are not tied to financial returns. Importantly, task sequencing matters: participants who completed the altruistic task first contributed more to carbon offsets in the subsequent investment round, suggesting path dependency and moral spillover across decision contexts. These findings contribute to behavioural ethics and sustainable supply chain research by revealing how framing and task sequencing jointly shape moral motivation and perceived legitimacy in sustainability decisions. The study concludes by outlining implications for ESG strategy design, supplier engagement, and sustainability communication in ethically complex procurement environments.

Optional Enforcement Backfires: Experimental Evidence on the Regulation of Negative Externalities

Erte Xiao, Monash University

While regulations are a common tool for addressing negative externalities, their implementation often depends on individuals' willingness to bear the costs of enforcement. We conduct a controlled laboratory experiment to examine whether optional enforcement, defined as a setting in which affected stakeholders can choose whether to implement regulation, creates moral wiggle room for producers to justify increases in socially harmful production. Both incentive-based and command-and-control approaches are tested in the experiment. We find that introducing a regulatory option, even when it is not enforced, leads to significantly higher harmful production compared to the no-regulation condition. Importantly, this moral wiggle room is not solely driven by self-serving reasoning, as optional enforcement shifts even third parties' normative beliefs, making higher output appear more socially acceptable. Overall, our results highlight the risk of introducing regulation while allowing enforcement to remain optional.

Community-Based Solutions to Overcome Collective Action Problems: An Experiment in Rural Vietnam

Joseph Vecci, University of Gothenburg

Solving environmental issues like water pollution and climate change requires collective action, yet short-term individual interests often conflict with longer-term group benefits. This project examines key strategies theorized by Ostrom to improve collective action-enhanced community cooperation and environmental monitoring-which have been rarely tested outside the lab. Using a randomized controlled trial in Vietnam's Mekong Delta with over 1000 households, the study evaluates the impact of these strategies on reducing water pollution and long term outcomes. Endline data is currently been collected. This research aims to provide causal evidence to inform effective policies for tackling real-world environmental problems in developing countries.

Fairness views about the International Distribution of Climate Change Costs

Davide Pace, LMU Munich

During climate negotiations, disagreements on how to share the costs of climate change often impede reaching otherwise agreed-upon targets, including stopping global warming well below 2°C. One reason why agreements on cost sharing are hard to reach is that the negotiators from different countries argue in favour of different fairness criteria. Fairness views about cost sharing are also important to determine voluntary transfers from richer to poorer countries to help the green transition. If a rich country would like to make a transfer, which recipients are considered worthy of receiving it by its citizens? We run an experiment that has three objectives related to people's fairness views about how to allocates the costs of climate change. First the experiment estimates citizens' fairness views depending on the type of climate change cost that needs to be shared. Climate change imposes three different cost: mitigation cost, adaptation cost, and the cost of loss and damage. The experiment measures how citizens trade-off different fairness criteria and whether these trade-offs differ depending on the type of climate change costs. Second, the experiment measures citizens' misperception about the characterisitcs of their country which are relevant to determine its fair share of contribution towards climate change costs (i.e., share of global current and historical CO2 emissions, share of the global population living in that country, and share of the world's GDP produced in that country). Third, the experiment tests whether citizens' fairness norms causally impact preferences for cost sharing. To do so, the experiment implements an information treatment that corrects one of participants misperceptions. The idea is that correcting a participant's beliefs about their country characteristics will change what they think their country should do given her fairness views. We measure preferences for cost sharing in three ways. First we ask the participants which percentage of the climate change cost they would ask the UK to pay were they tasked with allocating

costs across countries. Second we implement a dictator game in which the participants have to allocate money either to the UK government or to an international organization fighting climate change. Third we run an obfuscated follow-up in which we ask for the participants' support for concrete climate policies.

Paying to Avoid the Spotlight

Te Bao, Nanyang Technological University

In the digital age, privacy in economic activities is increasingly threatened. In considering policies to address this threat, it is useful to consider what value, if any, people attach to privacy in their economic activities. This valuation may be influenced by a mixture of concerns including the desire for personal autonomy, concerns about the exposure of confidential information, and the risk of reputational damage due to dishonest or stigmatized behavior. Our focus is primarily on reputational concerns as we assess individuals' willingness to pay (WTP) to avoid scrutiny of their potentially dishonest behavior in a simple coin flipping task. We gather and analyze data from Japan, China, and the U.S.A. to determine if there are notable differences across these nations in WTP. Our findings reveal that people's WTP to "avoid the spotlight" is positive and economically sizable across all three countries and is the largest in Japan.

The King's Banquet: Altruistic Wilful Ignorance in the Laboratory

Xiaomeng Chen, Monash University

While previous studies highlight how individuals avoid costless information to create moral wiggle room or deflect blame, we explore a more prosocial rationale: the use of willful ignorance to shield others. Drawing on a novel game inspired by A Banquet without Hat Tassels, we show that a considerable share of principals (Kings) choose not to uncover agents' (Generals') misconduct, even absent personal gain for doing so. Most principals forego information due to generosity instead of trusting that agents will not misbehave. Agents reward the principal more when spared both monetary sanctions and social exposure.

Using fNIRS hyperscanning to test neurocognitive mechanisms of financial advisor misconduct in face-to-face interactions

Xiaofei Niu, Shandong University

We conduct a neuroeconomics experiment in which a dyad of participants (acting as a financial advisor and an investor, respectively) play a face-to-face financial reporting game while we measure their synchronous brain activity using functional near-infrared spectroscopy (fNIRS) hyperscanning. We find that financial advisors' deniable dishonesty (i.e., a severe and common form of misconduct characterized by plausible deniability) in face-to-face interactions between advisor-investor dyads is related to interpersonal neural synchronization associated with theory of mind. Our findings provide the first evidence for the neurocognitive mechanisms of financial advisors' deniable dishonesty in face-to-face interactions. More generally, our study demonstrates that neural data obtained from fNIRS hyperscanning can be used to investigate the underlying mechanisms of interactive economic decision-making among multiple individuals during face-to-face communication.

Playing Dumb to Look Green?

Alice Solda, Emlyon business school

This paper investigates whether individuals use information complexity as an excuse to remain ignorant to behave selfishly. We study this question in a context where individuals face a trade-off between their monetary payoff and their pro-environmental preferences. We propose that individuals use information complexity as an excuse to make self-serving mistakes, which allow them to behave more selfishly without compromising their pro-environmental image. To test this idea, we conducted an online experiment in which we varied (i) the complexity of the information regarding the environmental impact of a donation and (ii) whether there is a trade-off between participants' selfish motives and pro-environmental preferences. In line with our hypothesis, we found that participants make more mistakes when information is complex, but even more so when there is a trade-off between their monetary payoff and their pro-environmental preferences. Our findings suggest that pro-environmental individuals do 'play dumb' when doing so gives them an excuse to behave more selfishly without compromising their image.

Comparative learnings

Mia Tam, University of Melbourne

We characterize the sets of observable actions consistent with (Bayes) Nash equilibrium in otherwise standard symmetric duopoly models with linear demand and constant marginal costs in which each player puts some weight between zero and one on the other player's profit. The comparisons of these sets when agents are privately informed about their conduct parameter and when the conduct parameters are commonly known is descriptive of behavior before and after players have learned each other's type. With the exception of Cournot when both players' conduct parameters are one, there is a one-to-one mapping between BNE outcomes and NE. The transition from BNE outcomes to NE can be stark and surprising. For example, in Cournot, the largest BNE quantities correspond to even larger NE quantities, and if each player produces half of the monopoly quantity in a BNE, both will produce more in the corresponding NE. Statistically, actions are independent while agents learn, that is, in the BNE, and correlated after they have learned, that is, in the NE. We discuss implications of these finding for the interpretation of empirical data.

Uncertain Others: How Ambiguity Attitudes Shape Cooperation Across Social Dilemmas

Tse-Min Wang, National Taipei University

Cooperation in social dilemmas requires individuals to form expectations about others' behavior under strategic uncertainty. While prior research documents how people respond to uncertainty in isolated game environments, we still lack a systematic understanding of how ambiguity attitudes translate into beliefs and cooperative behavior across social dilemmas that differ in complementarity, substitutability, and free-riding incentives. We study how individuals perceive and respond to strategic uncertainty by comparing belief formation and cooperative tendencies across three social dilemmas -Prisoner's Dilemma, Minimum Effort, and Volunteer's Dilemma-. By measuring ambiguity preferences and eliciting expectations about others' actions, we identify how ambiguity aversion and pessimistic beliefs shape cooperation, coordination, and volunteering across structurally distinct interactions Our findings reveal systematic differences in perceived ambiguity, belief bias, and cooperative dynamics across the three dilemmas. By integrating ambiguity attitudes, belief patterns, personality traits, and repeated behavioral adjustment, the study offers new insight into the psychological foundations of collective action and clarifies when ambiguity suppresses cooperation-and when it can unexpectedly promote prosocial behavior.

Modeling Strategies in Indefinitely Repeated Prisoner's Dilemma Games

David Cooper, University of Iowa

We study freely interacting two-person teams playing infinitely repeated prisoners' dilemma (IRPD) games with perfect and imperfect monitoring. Define mutual defection as a case where a team and its opponent both defect in a stage game. A team unilaterally cooperates if it chooses cooperate after observing mutual defection in the previous stage game. Unilateral defection is defined in an analogous fashion and we refer to unilateral cooperation and unilateral defection collectively as unilateral deviation . It is easily confirmed that no standard strategy (e.g. Always Defect, any variation of Grim including semi-grim, any variation of TFT) is consistent with unilateral deviations. Econometric methods that fit a distribution of strategies from observed data typically account for unilateral deviations through trembles- i.i.d. errors that cause agents to choose an action other than their intended action. Existing models fail to address when unilateral deviations occur or what they imply for the underlying distribution of strategies. Likewise, even though it is obvious from casual inspection of experimental data that agents change their initial actions in supergames as they gain experience, little work (exceptions include Dal Bo and Frechette, 2011; Fudenberg and Karreskog Rehbinder, 2024) has been done trying to understand how agents change their initial actions or what this implies for the underlying distribution of strategies and the development of cooperation with experience. To address these issues, we modify the structural model most commonly used to identify the distribution of strategies, the strategy frequency estimatation method (SFEM; Dal Bo and Frechette, 2011), to allow for the possibility that agents switch strategies both within and between supergames. This captures the possibility that unilateral deviations and changes in initial action represent deliberate choices rather than random trembles. Allowing for strategy switches substantially improves the ability of the structural model to fit the data and substantially changes the estimated distribution of strategies. We find that the frequency of substantive chat between teammates is a far better predictor of switches between regimes than features of the observable history.

Strategic sophistication and mixed strategies

Evan Calford, Australian National University

We use three different mechanisms to elicit mixed strategies from subjects in a game theoretic experiment. We find substantial differences in the rate of pure strategy play across the different mechanisms, ranging from 48% to 70%. The distribution of mixed strategies, conditional on playing a mixed strategy, is not affected by the elicitation mechanism and very little of the observed mixing can be explained by subjects being indifferent between pure strategies. Subjects are more likely to play mixed strategies against strategic (i.e. human) opponents than against non-strategic (i.e. coin-flip) opponents. In an out-of-sample test we find that subjects who exhibit higher levels of strategic sophistication are less likely to use mixed strategies. Subjects who exhibit low levels of strategic sophistication can be split into two types - those who exhibit strong predilection towards mixed strategies and those who do not. We discuss the implications of our results for theories of strategic sophistication.

Psychological Pressure: The Case of Penalty Kicks

Romain Gauriot, Deakin University

Emotions and psychological pressure are often thought to distort decision-making, even among highly trained professionals. A prominent claim is that in soccer penalty shoot-outs, players "choke" under pressure. We re-examine this evidence using a large dataset of over 19,000 penalty kicks. First, we define several measures of pressure, explaining how it depends on the state of the game. Second, using the precise locations of shots, we construct different measures of performance. Third, we employ a new identification strategy: shots striking the post create exogenous variation in the

scoreline, and thus in the pressure faced by subsequent players. Fourth, we examine differences across expertise. We find no evidence that pressure affects performance in penalty shoot-outs.

Natural Experiment in Fair Division: Armageddon Chess

John Wooders, NYU Abu Dhabi

We use data on bids and game outcomes from Armageddon chess to test the Minimax Hypothesis. Armageddon games break ties in chess tournaments, with Black winning an Armageddon game either on a draw or by winning outright. The players bid for the right to play Black. The player who bids the smaller number of minutes plays Black, with an amount of time on their clock equal to their bid; the other player plays White with a full clock. We develop a simple model of bidding in Armageddon games. We find that bids and game outcomes are consistent with the Minimax Hypothesis. In particular, each player is equally likely to play Black or White, and a player's probability of winning the game is the same whether playing Black or White.