

OFFICE CONTACT INFORMATION

MIT Department of Economics
77 Massachusetts Avenue, E52-301
Cambridge, MA 02139
ajkoh@mit.edu
<https://www.andrewjkoh.com/>

HOME CONTACT INFORMATION

50 Rogers St
Cambridge 02142
Mobile: 857-928-4892

MIT PLACEMENT OFFICER

Professor David Autor
dautor@mit.edu
617-253-4669

MIT PLACEMENT ADMINISTRATOR

Mrs. Shannon Robinson
shmay@mit.edu
617-324-5857

DOCTORAL STUDIES **Massachusetts Institute of Technology (MIT)**
PhD in Economics, Expected completion June 2026
DISSERTATION: “*Essays on Learning and Games*”

DISSERTATION COMMITTEE AND REFERENCES

Professor Stephen Morris
MIT Department of Economics
77 Massachusetts Avenue, E52-422
Cambridge, MA 02139
+1 617-253-1000
semorris@mit.edu

Professor Drew Fudenberg
MIT Department of Economics
77 Massachusetts Avenue, E52-418
Cambridge, MA 02139
+1 617-253-1000
drewf@mit.edu

Professor Daron Acemoglu
MIT Department of Economics
77 Massachusetts Avenue, E52-446
Cambridge, MA 02139
+1 617-253-1000
daron@mit.edu

Professor Matthew Elliott
Cambridge University Faculty of
Economics, Austin Robinson Building, 44
Cambridge, UK CB3 9DD
+44 1223-335250
matthew.elliott@gmail.com

PRIOR EDUCATION **Oxford University** 2020-21
MSc in Mathematical Sciences, *with Distinction*

Cambridge University 2017-20
BA in Economics, *Triple-First Class Honors*
Ranked 1st/160 and Adam Smith Prize for Best Dissertation

CITIZENSHIP Singapore **GENDER:** Male

FIELDS Primary Fields: Economic Theory
Secondary Fields: Finance, Economics of Technology, Behavioral

MIT Economics

ANDREW KOH

OCTOBER 2025-- PAGE 2

TEACHING EXPERIENCE	<i>Advanced Mathematical Methods for Financial Engineering (MFin)</i>	2022
	<i>Teaching Assistant to Paul Mende (teaching rated 6.60/7)</i>	
	<i>Principles of Microeconomics (Undergraduate)</i>	2023
	<i>Teaching Assistant to Jonathan Gruber (teaching rated 5.80/7)</i>	
	<i>Mathematical Economic Modeling (Undergraduate)</i>	2024
	<i>Teaching Assistant to Muhamet Yildiz (teaching rated 7.00/7)</i>	
FELLOWSHIPS, HONORS, AND AWARDS	<i>2025 Stripe Economics of AI Fellowship</i>	
	<i>2024 Journal of Industrial Economics Fellowship</i>	
	<i>2024 Gordon B. Pye Dissertation Fellowship, MIT</i>	
	<i>2023 Winter Research Fellowship, Centre for the Governance of AI (GovAI)</i>	
	<i>2022 Robert Solow Fellowship, MIT</i>	
	<i>2021-23 Global Priorities Fellowship, Forethought Foundation</i>	
	<i>2021-23 Armen Avanessians Fellowship, MIT (deferred from 2020)</i>	
	<i>2020 Full scholarship for the master's in Maths, Jesus College Oxford</i>	
	<i>2020 Ranked 1st with highest overall mark in Cambridge cohort of 160 candidates</i>	
	<i>2020 Adam Smith Prize for Best Dissertation in Economics, University of Cambridge</i>	
PROFESSIONAL ACTIVITIES	<i>2018-20 E.M. Burnett Prizes, Hughes Hall Cambridge</i>	
	<u><i>Referee</i></u>	
	<i>American Economic Review, Review of Economic Studies, Journal of Political Economy, American Economic Review: Insights, Journal of Economic Theory, American Economic Journal: Microeconomics, RAND Journal of Economics, Games and Economic Behavior, Journal of Mathematical Economics, Economic Theory</i>	
	<u><i>Selected talks</i></u>	
	<i>2025: SITE ('Dynamic Games Contracts and Markets'), AEA Meetings ('Policy Implications of Transformative AI'), Johns Hopkins</i>	
	<i>2024: Harvard Computer Science, Berkeley-Columbia-Duke-MIT-Northwestern IO Conference (Duke), ACM EC'24 (Yale), Nuffield College Oxford, Cambridge University, National University of Singapore, Econometric Society Summer Meeting</i>	
	<i>2023: Singapore Management University</i>	

**PUBLICATIONS/
FORTHCOMING**

1. **“Capability Accumulation and Conglomeratization in the Information Age”** (with Jun Chen and Matthew Elliott) *Journal of Economic Theory* 210 (2023): 105647.
2. **“An Economy of AI Agents”** (with Gillian Hadfield) [\[link\]](#)
Forthcoming, NBER Handbook on The Economics of Transformative AI

In the coming decade, artificially intelligent agents with the ability to plan and execute complex tasks over long time horizons with little direct oversight from humans may be deployed across the economy. This chapter surveys recent developments and highlights open questions for economists around how AI agents might interact with humans and with each other, shape markets and organizations, and what institutions might be required for well-functioning markets.

**JOB MARKET
PAPER**

3. **“Robust Technology Regulation”** (with Sivakorn Sanguanmoo) [\[link\]](#)

We analyze how uncertain technologies should be robustly regulated, and how regulation should evolve with new information. An adaptive sandbox comprising a zero marginal tax on R&D up to an evolving quantity limit is (i) robust: it delivers optimal payoff guarantees when the agent’s learning process and/or preferences are chosen adversarially; (ii) dominant: it outperforms other robust mechanisms across all agent learning processes and preferences; (iii) time-consistent: it is the only robust mechanism that can be implemented without commitment. We argue that robustness is important—absent robust regulation, worst-case payoffs can be arbitrarily poor and are induced by weak but growing optimism that encourages excessive risk-taking. Our results offer optimality foundations for existing policy and speak directly to current debates around managing emerging technologies.

**UNDER
REVISION**

4. **“Attention Capture”** (with Sivakorn Sanguanmoo) [\[link\]](#)
Revise and resubmit, Journal of Political Economy

We develop a unified analysis of how information captures attention. A decision maker (DM) faces a dynamic information structure and decides when to stop paying attention. We characterize the convex-order frontier and extreme points of feasible stopping times, as well as dynamic information structures which implement them. This delivers the form of optimal attention capture as a function of the designer and DM’s relative time preferences. Intertemporal commitment is unnecessary: sequentially optimal information structures always exist by inducing stochastic interim beliefs. We further analyze optimal attention capture under noninstrumental value for information. Our results speak directly to the attention economy.

5. **“Market Segmentation through Information”** (with Matthew Elliott, Andrea Galeotti, Wenhao Li) [\[link\]](#)

Revise and resubmit, The Review of Economic Studies

We explore the power that precise information about consumers’ preferences grants an intermediary in shaping competition. We think of an intermediary as an information designer who chooses what information to reveal to firms, which then compete a la Bertrand in a differentiated product market. We characterize the information designs that maximize consumer and producer surplus, showing how information can be used to segment markets to intensify or soften competition. Our analysis demonstrates the power that users’ data can endow intermediaries with, and speaks directly to current regulatory debates of digital marketplaces.

WORKING PAPERS

6. **“Balanced Social Learning”** (with Ricky Li) [\[link\]](#)

When and how can efficient social learning be achieved in the presence of multiple actions and heterogeneous preferences? We completely characterize the value and form of policies that control the flow of information from past to future. *Full transparency* in which all past information is disclosed guarantees efficiency if and only if (i) the experiment approaches noisy bad news where Type I errors dominate Type II errors; and (ii) preference dispersion satisfy a path-balance condition. *Adaptive evidence design* in which subsets of past signals are dynamically disclosed can guarantee efficiency if and only if priors are initially balanced, a far weaker condition than (i) and (ii). Finally, if efficiency cannot be guaranteed, we fully characterize the maximum probability it can be achieved under *adaptive information design*.

7. **“Data-Driven Automation”** (with Maryam Farboodi, Anchi Xia) [\[link\]](#)

We build a model of data-driven automation in which data (i) is *heterogeneous* – different kinds of data are differentially valuable across different sectors; (ii) exhibits *spillovers* – one sector's data can augment the productivity of another (via transfer learning); and is (iii) *accumulated endogenously* in general equilibrium. We derive tight conditions for the economy to either exhibit a balanced data path to achieve full limit automation, or an imbalanced data path in which automation fizzles out. Data-driven automation can be contagious: cross sector spillovers via transfer learning is sufficient for full automation; but inefficient: small firms do not internalize the value of future data and equilibrium automation can be inefficiently fast or slow.

8. **“Flexible Demand Manipulation”** (with Yifan Dai) [\[link\]](#)

We develop a simple framework to analyze how targeted persuasive advertising shapes market power and welfare. A designer flexibly manipulates the demand curve by influencing individual valuations at a cost. A monopolist prices against this manipulated demand curve. We fully characterize the form of optimal advertising plans under ex-ante and ex-post welfare measures. Flexibility per se is powerful and can substantially harm or benefit consumers vis-a-vis uniform advertising. We discuss implications for regulation, intermediation, and the joint design of manipulation and information.

9. **“Informational Puts”** (with Sivakorn Sanguanmoo, Kei Uzui) [\[link\]](#)
Extended abstract in 25th ACM Conference on Economics and Computation

We analyze how dynamic information should be provided to uniquely implement the largest equilibrium in binary-action coordination games. The designer offers an informational put: she stays silent if players choose her preferred action, but injects asymmetric and inconclusive public information if they lose faith. Informational puts (i) *close the multiplicity gap*: the largest partially implementable equilibrium can be implemented uniquely; is (ii) *sequentially optimal*; and is a (iii) *universal equilibrium selector*: it can be added atop a general exogenous learning process to uniquely implement the largest rationalizable equilibrium. Our results have sharp implications for the design of policy in coordination environments.

10. **“Inertial Coordination Games”** (with Ricky Li, Kei Uzui) [\[link\]](#)
Extended abstract in 26th ACM Conference on Economics and Computation

We analyze inertial coordination games: dynamic coordination games with an endogenously changing state that depends on (i) a persistent fundamental that players privately learn about over time; and (ii) past play. The speed of learning determines long-run equilibrium dynamics: the risk-dominant action is played in the limit if and only if learning is slow such that posterior precisions grow sub-quadratically. This generalizes results from static global games and endows them with a learning foundation. Conversely, when learning is fast such that posterior precisions grow super-quadratically, shocks can propagate and generate self-fulfilling spirals.

11. **“Prices and Symmetries”** (with Pedro Martinez-Bruera) [\[link\]](#)

We analyze optimal interventions in networked economies with stochastic linkages. Linear taxes dominate quotas whenever shocks are symmetric because the ensuing equilibrium adjusts toward the first-best allocation, thereby correcting policy mistakes. Vice versa for antisymmetric shocks. Flexible interventions can implement the first-best allocation for each

realized network whenever shocks are symmetric. Our results offer foundations for price interventions when shocks have a large common component, and for quantity interventions when policymakers are concerned about the correlation structure across links.

12. **“Persuasion and Optimal Stopping”** (with Sivakorn Sanguanmoo, Weijie Zhong) [\[link\]](#)

We analyze the interplay between persuasion, timing, and commitment. A principal conducts a sequence of statistical experiments to persuade an agent to stop at the right time, in the right state, and choose the right action. We develop a revelation principle which delivers a first-order approach for solving the principal’s problem under commitment, and an anti-revelation principle which establishes that commitment is unnecessary and transforms the solution via indirect recommendations to restore dynamic consistency. We further characterize how time and action preferences jointly shape optimal strategies featuring a suspense-generation stage which optimally concentrates the agent’s stopping time, followed by an action-targeting stage which maximally correlates/anticorrelates persuasion and delay.

13. **“Speed vs Resilience in Contagion”** (with Stephen Morris) [\[link\]](#)

We highlight a trade-off between speed (the rate at which behaviors propagate in the population) and resilience (the measure of initial adopters required for spreading) in models of threshold contagion: contagion is faster in networks where it is harder to initiate contagion. We derive various orderings over networks under which this trade-off is stark. While this trade-off holds between pairs of networks for possibly different contagion thresholds, we also outline conditions under which, for a given contagion threshold, one network is both less resilient and propagates behaviors more quickly; this highlights the role of intermediate links as bulwarks against contagion.

14. **“Memory Correlated Equilibrium”** (with Sivakorn Sanguanmoo)
[draft online soon; available upon request]

We develop a framework for modeling random and correlated memory in games: a memory correlated equilibrium (MCE) comprises of (i) a base extensive-form game; and a (ii) memory structure that delivers private self-locating information at each history. We show MCE exist and develop a revelation principle which paves the way for memory design in games. We illustrate applications to cooperation, mechanisms, coordination, and hold-up problems, as well as discuss implications for the design of algorithms and artificial agents.

15. **“Delaying the Deviation”** (with Anna Merotto)
[draft online soon; available upon request]

We show that delaying the deviation—randomizing incentives to delay attacks both across players and across time—uniquely implements not attacking in dynamic coordination games. Our mechanism is (i) *universal*: it prevents attacks even when attacking is both payoff- and risk-dominant; (ii) *first-best*: it incurs zero cost on-path and vanishing cost off-path; (iii) *sequentially optimal*: the designer is incentivized to follow-through at every history; (iv) *heterogeneity-robust*: it works against any distribution of players’ payoff types; and (v) *strategically simple*: it relies only on players’ first-order certainty of rationality. We discuss applications to bank runs, platforms, security design, and political regimes.