RESEARCH STATEMENT

Andrew Koh*

I am an economic theorist and game theorist working across 'pure' and 'applied' aspects of theory. My work can be broadly organized into two interrelated themes:

1. **The economics of technology.** My work here brings economic theory to bear on *the economic impact of new technologies* e.g., on automation [1], mergers and conglomeration [2], market power on digital platforms [3], manipulative and targeted advertising [4], the attention economy [5]. I also work on understanding *technology policy*—how policymakers might optimally steer the trajectory of technological progress through industrial policy [7] and regulation [8, 6].
2. **Learning and dynamic games.** My work here spans social learning [9], dynamic information [5, 10], coordination games and mechanisms [11, 12, 13, 14], and memory in games [15]. In these papers I tackle more conceptual topics, but am always motivated by concrete economic questions. Along the way I also develop methodological tools that might be applied more broadly.

THE ECONOMICS OF TECHNOLOGY

▶ **Market Power, Platforms & Conglomeration.** In '**Market Segmentation through Information**' (Elliott, Galeotti, Koh, and Li, 2025) [resubmitted to *The Review of Economic Studies*] with Matt Elliott, Andrea Galeotti, and Wenhao Li we study how and to what extent providing information about consumer preferences to oligopolists can amplify or soften competition. That is, when can information implement the 'collusive outcome' even without dynamic incentives? We give tight conditions on the distribution over consumer valuations when all surplus—that is, all gains from trade—can be extracted as producer surplus. This condition is surprisingly weak, and is fulfilled, for instance, in 'standard' models e.g., Hotelling. We also solve for the consumer-optimal information design which uses information to amplify competition, as well as characterize all outcomes that can obtain through jointly designing information and matching i.e., which firms consumers can access.

Although digital platforms are typically conceived as intermediaries, they might also shape market power more subtly by exploiting consumer biases. In '**Flexible Demand Manipulation**' (Dai and Koh, 2024) with my classmate Yifan Dai, we analyze the interplay between (persuasive) targeted advertising and market power. While targeted advertising has received considerable attention from policymakers due to privacy concerns, their impact on prices and welfare remain unclear. We develop a simple framework to analyze such questions. A designer (e.g., digital platform) can flexibly manipulate the demand curve by influencing individual valuations at a cost. A monopolist prices against this manipulated demand curve. We completely characterize the form of optimal advertising plans under ex ante and ex post welfare measures. Flexibility *per se* is powerful and can

*MIT Department of Economics; October 2025; prepared for the 25/26 job market

substantially harm or benefit consumers vis-a-vis uniform advertising and our results have sharp implications for regulation, intermediation, and the joint design of manipulation and information.

Digital platforms are not just in the business of selling things, but also extracting users' attention. In '**Attention Capture**' (Koh and Sanguanmoo, 2022) [R&R at the *Journal of Political Economy*] joint with my classmate Sivakorn Sanguanmoo, we ask the basic question of how information can be used to extract attention from a receiver for whom information is either instrumentally or non-instrumentally valuable. We make four contributions in the latest draft. First, we completely characterize the extreme points of all stopping times that can be induced, and the accompanying dynamic information structures that implement them. En route we develop new results on the extreme points of distributions that satisfy conditional moment constraints—we expect this will be more broadly useful in dynamic design problems. Second we show that neither *inter-* nor *intratemporal* commitment is necessary to optimally extract attention. That is, whatever the designer can implement via dynamic information design, she can also implement via dynamic cheap talk. Third, we uncover a tight connection between attention capture under instrumental and non-instrumental value (e.g., 'suspense' or 'surprise') for information. Finally, we show how to make the dynamic information structure *exit proof*, thereby turning the receiver's control problem (continually deciding whether or not to pay attention) into a stopping problem. This is done through a novel encoding of the *meaning* of future messages within past messages such that a receiver who has not consecutively paid attention is unable to interpret new information.

Over the past decade, we have seen the rapid expansion of tech firms into new industries. In '**Capability Accumulation and Conglomeritization in the Information Age**' (Chen, Elliott, and Koh, 2023) in the *Journal of Economic Theory* we develop a parsimonious model to explain why tech firms have grown so quickly via acquisitions, and into so many new industries (mergers of scope rather than of scale). Our framework draws on the language of hypergraphs that associates each firm with bundles of capabilities it possess, and each market with bundles of capabilities it values. For a given market, the firms with capabilities for that market compete, and their relative competitiveness is determined by the number of relevant capabilities they wield in that market. We let firms merge and de-merge, and study stable industry structures. We show that as markets begin to value more of the same capabilities (e.g., the ability to process 'big data'), this leads to a sudden jump in the size of the largest firm across all stable industry structures. This offers a parsimonious explanation for recent merger activity among tech firms and our capability-based framework has attracted substantial attention from European regulators.[1]

▶ Optimal Regulation & Intervention. In my job market paper '**Robust Technology Regulation**' (Koh and Sanguanmoo, 2025), we take three facts about technology seriously: *uncertainty*, *irreversibility*, and *externalities*. We show that an *adaptive sandbox mechanism* that imposes a zero marginal tax up to an evolving quantity limit is (i) robust i.e., it delivers optimal payoff guarantees over all tech firm's learning processes (*how* and *what* they might learn as they do R&D) and preferences (e.g., their risk-aversion); (ii) dominant among robust mechanisms i.e., it performs strictly better away from the worst-case; and (iii) is uniquely time-consistent among

[1] For instance, our paper was cited in 2022 by the Chief Economist at the European Commission's Department for Competition in a speech on platform mergers.

robust mechanisms. We further show that robustness is important—in its absence (e.g., under a linear Pigouvian tax) the worst-case learning processes takes the form of an inhomogeneous bad news process such that in the absence of bad news (e.g., a financial crises, AI disaster, etc.) the tech firm is emboldened to take on excessive risk. This can substantially magnify small wedges in risk-aversion and induce unboundedly poor payoffs. Our results offer a novel perspective on how mechanisms should simultaneously *evolve* with new information (adaptivity) while safeguarding against the worst-case (robustness).[2] They also offer optimality foundations for a number of current policies ranging from FDA drug trials to sandboxes for financial technologies, as well as sharp prescriptions for how new technologies such as AI should be regulated.

My job market paper takes a somewhat 'reduced form' view of externalities that are, in practice, embedded within a complex web of spillovers. For instance, if a firm overdeploys a risky technology, this might disproportionately hurt the firms closest to it in the supply chain. In **'Prices and Symmetries'** (Koh and Martinez-Bruera, 2025) with my classmate Pedro Martinez-Bruera, we take the network structure of externalities seriously and analyze how a planner should intervene to correct externalities when linkages are stochastic. Linear taxes dominate quotas whenever shocks are symmetric because the ensuing equilibrium adjusts endogenously toward the first-best allocation, thereby correcting policy mistakes. By contrast, quotas dominate when shocks are antisymmetric. This speaks directly to the old socialist calculation debate: prices induce responsiveness (an insight by Hayek (1945)) but in the presence of externalities, responsiveness *per se* does not always work well—we make precise that symmetry is what is important for equilibrium allocations to adjust in the right direction *toward* efficiency. We further show—perhaps surprisingly—that flexible interventions (e.g., nonlinear taxes) can implement the first-best allocation *network-by-network* if and only if shocks to the network are symmetric or antisymmetric. Taken together, 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 of links.

In '**Data-Driven Automation**' (Farboodi, Koh, and Xia, 2025) joint with Maryam Farboodi and Anchi Xia, we develop a simple model that takes three facts about data and modern AI systems seriously: data is *heterogeneous*—sectoral data is specialized, *exhibits spillovers*—one task's data can augment the productivity of another, and is *accumulated endogenously* in equilibrium. First, 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. Second, data-driven automation is contagious—any cross-task spillovers via transfer learning is sufficient for full automation irrespective of task substitutability. We characterize the long-run behavior of this economy, and show that the composition of tasks produced and data converge proportionately to the *leading eigenfunction* of the graphon of cross-task spillovers. Third, data-driven automation is generically inefficient as atomistic firms do not internalize the value of data on future productivity. Automation in the decentralized equilibrium can be inefficiently fast or slow depending on the degree of cross-task elasticity. Our results have sharp implications for how policymakers might steer the *direction* of data accumulation.

[2] Our analysis of instruments for regulation is in the spirit of Weitzman (1974) whose work has influenced mine enormously.

LEARNING AND DYNAMIC GAMES

▶ **Learning & Dynamic Information.** In '**Persuasion and Optimal Stopping**' (Koh, Sanguanmoo, and Zhong, 2024) joint with Sivakorn and Weijie Zhong, we build on 'Attention Capture' (referenced above) by developing a general analysis of the *joint* distribution over actions, states, and stopping times that can be implemented with dynamic information. Our contribution is fourfold. First, we develop a powerful and flexible toolkit based on a functional first-order approach and strong duality to analyze such problems. Second, we show that intertemporal commitment is unnecessary in such environments, even when the designer has arbitrary preferences over actions and/or states. This is surprising, and disproves a conjecture by (Ely and Szydlowski, 2020)—we provide an algorithm to restore commitment, and apply to their 'moving the goalposts' strategy. Third, we develop broad principles on how (i) relative time-risk preferences; (ii) persuasion gain; and (iii) the direction of persuasion-delay complementarity or substitutability shapes optimal dynamic information. Fourth, we develop applications. For instance, when the state is continuous a *dynamic tail policy*—a generalization of the upper- and lower-censorship policies typical in static persuasion—is often optimal. This offers a novel perspective on the influential real options literature in macroeconomics and finance (Bernanke, 1983), and rationalizes the common practice of using fan charts (central bank communication) or reaffirming earnings guidance.

In '**Balanced Social Learning**' (Koh and Li, 2025) with Ricky Li in the year below, we completely characterize when and how efficient social learning can be achieved in the presence of multiple actions and heterogeneous preferences. In such settings efficiency requires *balance*: too few good actions might be learnt about leading to forgone surplus from horizontal matches. Under transparency in which all past signals are disclosed, efficient learning is achieved if and only if (i) the underlying experiment approaches noisy bad news where Type I dominate Type II errors; and (i) preference dispersion satisfies a path balance condition. To show this, we show that in the noisy experiment limit, the multidimensional belief process solves the 'martingale problem' posed by the infinitesimal generator of a conjectured Feller process—the latter's behavior can be tractably analyzed to shed light on the probability of balanced learning. Our results uncover a fundamental tradeoff between *speed* and *breadth* of learning in environments with many actions—contra Acemoglu et al. (2022), we show that slowing learning down can achieve efficient long-run outcomes and offer sharp prescriptions for how signals (e.g., reviews) should be coarsened. We partially characterize this speed-balance frontier and show that the success probability of learning goes to 1 exponentially with the noisiness of the bad news signal. Although transparency is widespread in practice, platforms and governments can and do censor past information e.g., reviews are routinely coarsened and/or removed. We completely characterize the form and value of these kinds of adaptive policies, leveraging techniques from ongoing work on *feasible adapted beliefs* when agents are learning from outside information (Koh and Sanguanmoo, 2025).

In '**Memory Correlated Equilibrium**' with Sivakorn, we develop an expressive framework for modelling random and correlated memory in extensive form games. Our starting point is the observation that the 'standard' approach to modeling imperfect recall (insofar as there is one) uses information sets (Piccione and Rubinstein, 1997). But this seemed needlessly restrictive—we develop the notion of *memory correlated equilibrium* (MCE) that comprises (i) a base extensive-form game; and a (ii) memory structure that delivers private self-locating information at each history. We develop progressively stronger notions of MCE and show they generally exist. We then develop a revelation principle, and an efficient procedure to characterize the set of all outcomes in MCE for a fixed extensive form game. Our results pave 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.

▶ **Dynamic Coordination & Mechanisms.** In '**Informational Puts**' (Koh, Sanguanmoo, and Uzui, 2024) [extended abstract in *EC'24*] with Sivakorn and Kei Uzui in the year below, we analyze how dynamic information should be provided to uniquely implement the largest equilibrium in binary-action coordination games in which players face arbitrary switching frictions. We show that the policymaker can offer an *informational put*—she stays silent if players choose her preferred action, but adaptively injects asymmetric and inconclusive public information if they lose faith (i) closes 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.[3] Our results have stark implications for the design of policy in coordination environments. Informational puts leverage players' uncertainty about fundamentals. Can we do the same under *complete* information? Yes—in '**Delaying the Deviation**' (Koh and Merotto, 2025) with Anna Merotto, we analyze a complete information and frictionless coordination game in which a large population of players make irreversible stopping ('attack') decisions. If the designer has access to transfers, carefully randomizing delay incentives across players and across time (i) full implementation: 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 (but without conditioning on) heterogeneous payoff types; and (v) strategically simple: it relies only on players' rationality and first-order certainty of rationality (as opposed to CKR). This has sharp implications for networked pricing, political regimes, and security design e.g., it rationalizes the use of staggered contracts.

In '**Inertial Coordination Games**' (Koh, Li, and Uzui, 2024) [extended abstract in *EC'25*] with Ricky and Kei, we develop a new model of dynamic coordination in which past actions propagate into the future with some (potentially small) lag. We show that the speed of learning determines long-run equilibrium dynamics: the risk-dominant action is played in the limit if and only if 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. Our model and results reconcile the 'history vs expectations' view of business cycles (Krugman, 1991), as well as the 'fundamental' vs 'coordination' view of bank runs.

In '**Speed vs Resilience in Contagion**' (Koh and Morris, 2022) with Stephen Morris, we uncover a fundamental tension between resilience (the requisite measure of 'infected' players choosing an action e.g., adopt a product, protest etc. for it to spread) and speed (conditional on

[3] Our results stand in contrast to the important papers of Frankel and Pauzner (2000) who show the risk-dominant equilibrium is always selected when the learning process is Brownian.

spreading, the extra measure of 'infected' players per period) in networked environments. When the network exhibits geometry such that 'nearby' players (e.g., those similar to each other) are more likely to be connected, we show that more resilient networks *must* propagate the action more quickly conditioned on spreading. In the latest version we use graphexes—sparse continuous graphs recently introduced in probability theory—to formalize these tradeoffs, and show that they well-approximate the large population limit where players are linked via a geometric random graph. Our results on the speed-resilience tradeoff help us understand the dynamics of equilibrium (e.g., norm) transitions, and offer new insights into how fragility should be understood in financial and production networks.

REFERENCES

[1] Maryam Farboodi, Andrew Koh, and Anchi Xia. Data-driven automation. *MIT Working Paper*, 2025.
[2] Jun Chen, Matthew Elliott, and Andrew Koh. Capability accumulation and conglomeratization in the information age. *Journal of Economic Theory*, 210:105647, 2023.
[3] Matthew Elliott, Andrea Galeotti, Andrew Koh, and Wenhao Li. Market segmentation through information. *Available at SSRN 3432315*, 2025.
[4] Yifan Dai and Andrew Koh. Flexible demand manipulation. *arXiv preprint arXiv:2410.24191*, 2024.
[5] Andrew Koh and Sivakorn Sanguanmoo. Attention capture. *arXiv preprint arXiv:2209.05570*, 2022.
[7] Andrew Koh and Pedro Martinez-Bruera. Prices and symmetries. *Available at SSRN 5132200*, 2025.
[8] Andrew Koh and Sivakorn Sanguanmoo. Robust technology regulation. *arXiv preprint arXiv:2408.17398*, 2025.
[6] Gillian K Hadfield and Andrew Koh. An economy of ai agents. *NBER Handbook on The Economics of Transformative AI*, forthcoming.
[9] Andrew Koh and Ricky Li. Balanced social learning. *Available at SSRN 5699342*, 2025.
[10] Andrew Koh, Sivakorn Sanguanmoo, and Weijie Zhong. Persuasion and optimal stopping. *arXiv preprint arXiv:2406.12278*, 2024.
[11] Andrew Koh, Sivakorn Sanguanmoo, and Kei Uzui. Informational puts. *arXiv preprint arXiv:2411.09191*, 2024.
[12] Andrew Koh, Ricky Li, and Kei Uzui. Inertial coordination games. *arXiv preprint arXiv:2409.08145*, 2024.
[13] Andrew Koh and Stephen Morris. Speed vs resilience in contagion. *Available at SSRN*, 2022.
[14] Andrew Koh and Anna Merotto. Delaying the deviation. *MIT Working Paper*, 2025.
[15] Andrew Koh and Sivakorn Sanguanmoo. Memory correlated equilibrium. *MIT Working Paper*, 2025.
[16] Martin L Weitzman. Prices vs. quantities. *Review of Economic Studies*, 41:315–329, 1974.
[17] Friedrich A Hayek. The use of knowledge in society. *American Economic Review*, 35(4):519–528, 1945.
[18] Jeffrey C Ely and Martin Szydlowski. Moving the goalposts. *Journal of Political Economy*, 128(2): 468–506, 2020.
[19] Ben S Bernanke. Irreversibility, uncertainty, and cyclical investment. *The quarterly journal of economics*, 98(1):85–106, 1983.
[20] Daron Acemoglu, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. Learning from reviews: The selection effect and the speed of learning. *Econometrica*, 90(6):2857–2899, 2022.
[21] Andrew Koh and Sivakorn Sanguanmoo. Feasible adapted beliefs. *MIT Working Paper*, 2025.
[22] Michele Piccione and Ariel Rubinstein. On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior*, 20(1):3–24, 1997.
[23] David Frankel and Ady Pauzner. Resolving indeterminacy in dynamic settings: the role of shocks. *The Quarterly Journal of Economics*, 115(1):285–304, 2000.
[24] Paul Krugman. History versus expectations. *The Quarterly Journal of Economics*, 106(2):651–667, 1991.