Research Statement

Shiyu (Jake) Zhang

California Institute of Technology · October 2026

My research in microeconomic theory examines how informational frictions shape individual decision-making and strategic coordination. Across both single-agent and interactive settings, behavior is governed by more than the quantity of available information. Environments that look equivalent by standard measures of information can produce distinctive choice patterns and very different prospects for coordination. I model the structure of informational frictions, on what people can tell apart and what they can infer about others' information. My work helps economists separate cognitive constraints from preferences in individual choice, and identify when the common beliefs assumed in equilibrium analysis actually emerge.

Costly Perception and the Shape of Mistakes

"A Costly-Information Foundation for Psychometric Curves" (Job Market Paper)

Many consequential decisions require classifying cases whose evidence is costly to acquire or assess. Lenders evaluate loan applications, clinicians diagnose patient symptoms, and judges assess recidivism risk. A closer inspection improves decision quality but requires time and effort, so decision makers trade off deliberation costs against classification accuracy. My job market paper asks how the structure of these costs determines where mistakes occur and how such decisions should be interpreted, incentivized, and organized.

To capture the intrinsic difficulty of fine distinctions, I study binary choice over a continuous state space using the Fisher information cost of Hébert and Woodford (2021). This cost makes it expensive for behavior to differ sharply across nearby states, so how well a case is handled depends not only on what is at stake. This framework yields two primary sets of results:

Explaining diminishing sensitivity to evidence. The S-shaped choice curve is a form of diminishing sensitivity: choices respond most to evidence near the decision boundary and progressively less as the evidence becomes more extreme, a regularity documented from perception to risky choice (Kahneman and Tversky, 1979). In symmetric environments, I establish a necessary and sufficient condition for the optimal choice rule to be S-shaped across all incentive levels. It says the marginal benefit of adjusting behavior must decline as evidence becomes more obvious. Here diminishing sensitivity is derived rather than assumed, and its degree responds to stakes and to the prior beliefs.

The model also predicts where choices are most informative. When cases are uniformly distributed, discriminability, that is, how well nearby cases are told apart, peaks exactly at the payoff cutoff where the better action switches. Outside symmetric environments, this cutoff need not coincide with the point where each action is chosen half the time. This matters for inference. Researchers often infer a decision maker's standard, such as a lender's risk tolerance or a judge's standard

of proof, from the point where she is equally likely to choose either action. In my model, that 50-percent point need not coincide with the payoff cutoff where the better action switches. Two decision makers with identical preferences but facing different information costs can therefore appear to hold different standards. Separating perception from preference requires modeling how attention is allocated, not only fitting choices.

Comparative Statics of Attention. The model generates concrete predictions for how choices shift with the environment. First, in classic discrimination tasks where two actions are treated symmetrically, higher stakes reduce errors across all states, but with sharply diminishing returns. Average error falls with the inverse square root of the reward, so quadrupling the reward only halves it, contrasting with the exponential error reduction predicted by standard model using mutual information cost.

Second, the environment also matters. Concentrating cases symmetrically toward the decision boundary, holding incentives fixed, increases errors at every state, including distant cases. An inspector whose caseload contains more borderline cases makes more mistakes even on the easy ones. This creates an immediate identification problem: even evaluation of performance case by case can conflate a decision maker's competence with the difficulty of her caseload.

Third, increasing the payoff to an action or shifting prior beliefs in its favor raises the choice probability of that action monotonically across the entire state space, suggesting a neighborhood spillover. Hence, incentives also cannot be targeted case by case. Rewarding an action more in some cases makes the decision maker choose it more often in every case. A larger penalty for approving risky loans also lowers approval of loans that look safer.

These predictions can be tested in a standard discrimination experiment that varies rewards and the stimulus distribution independently. With the cost parameter calibrated in a baseline condition, the model predicts the entire response curve under new payoffs and priors, which a fitted logit or probit curve cannot do.

When Does Individual Learning Become Common Learning?

Successful coordination requires mutual confidence. In speculative investments, debt rollovers, or organizational adoption, an agent's incentive to act depends on her beliefs about others' beliefs, iterated to arbitrary orders. While Rubinstein (1989) showed that higher-order coordination can collapse entirely despite near-certain private signals, Cripps et al. (2008) proved that with finite signal spaces, agents who learn the state eventually achieve common learning. My second agenda investigates the boundary between these outcomes along two dimensions: communication dynamics over time and the spatial geometry of social networks.

"Common Learning in Email Games" (work in progress)

I consider a setting where two agents repeatedly learn about a state through a generalized version of Rubinstein's email protocol. In each period, one agent receives news about the state and, if the news is favorable, emails the other; the other confirms, the first acknowledges the confirmation, and so

on until a message is lost. The probability that a message is lost may vary with how many messages have already been exchanged. Once the first message is sent, later acknowledgments carry no extra information about the state, so protocols that differ only in the reliability of later messages generate the same individual beliefs. Their implications for common learning can nonetheless be opposite.

I show that common learning is governed by an agent's ability to localize the point where communication stopped. If the posterior probability that one's last message was lost on her side settles to a positive limit as exchanges grow long, common learning occurs if and only if this limit equals one for at least one agent. Such an agent can infer the other's record of past exchanges almost exactly, so doubts about the other's information stop compounding from one level of confirmation to the next. In contrast, in the counterexample of Cripps et al. (2008), both limits fall short of one and common learning never occurs.

This mechanism produces two surprising conclusions. Making late messages more reliable can undermine common learning while letting late messages fail with increasing probability restores common learning. And because only the tail matters, two protocols that agree up to any fixed finite number of messages can have opposite implications for common learning.

"Local Global Games and Network Common Learning" (with Olga Rospuskova and Omer Tamuz; working paper)

When agents coordinate with their network neighbors, the network decides how information travels. We consider a model where agents relay signals through the network, so that each observes the private signals of everyone within graph distance r. As $r \to \infty$, every agent learns the true state, and neighbors' observations overlap almost completely. We show that whether efficient coordination is attainable nonetheless depends decisively on the network's geometry. On a line, when coordination requires sufficiently high confidence, only the safe action survives in equilibrium, however large r is. On grids of dimension two and higher, efficient coordination becomes possible as r grows.

To analyze equilibria, we define *network q-belief*, a local analogue of Monderer and Samet's (1989) common p-belief that chains beliefs along the edges of the network, and show that it is equivalent to the existence of an efficient equilibrium. Whether it emerges as r grows turns on the *information boundary*: the set of signals one agent observes that her neighbor does not. On a line, this boundary contains a single signal regardless of r, so uncertainty about a neighbor's information stays significant and propagates from neighbor to neighbor. On a grid, the boundary grows with r, and concentration of measure makes a neighbor's unobserved signals predictable. Building on this, we introduce *network common learning*, the network analogue of common learning, and give sufficient conditions for it that cover grids of dimension two and higher and bounded-degree trees whose neighborhoods grow fast enough.

The two papers share an economic lesson: more plentiful or more reliable information need not make coordination easier. Information structures cannot be ranked for coordination by how much they reveal, since statistically equivalent information can support different equilibrium outcomes. In communication, making late confirmations more reliable can backfire, while deadlines that cap

the number of confirmations restore common learning. In networks, who shares information with whom matters more than how much is shared. Along a supply chain, where each firm's investment pays off only if its neighbors in the chain also invest, a sufficiently risky investment cannot be coordinated however much information flows; in grid-like clusters or branching hierarchies, it can.

Ongoing and Future Work

Multiple Alternatives and Ordered Categorization. Many economic classifications involve multiple ordered categories. This raises questions about how local discrimination costs shape choices across multiple alternatives. A natural conjecture is that errors concentrate between adjacent alternatives and that distant categories exert minimal substitution pressure, which can generate violations of the independence of irrelevant alternatives from random utility models. This setup can also explain findings from absolute-identification experiments, where extreme stimuli are identified more quickly and accurately than interior ones (Stewart et al., 2005): boundary categories face confusion from only one adjacent alternative, whereas interior categories face two-sided competition. I am also exploring how to interpret the Fisher metric as deliberation time and link the model to response-time distributions.

Empirical and Experimental Tests. The comparative statics of my job market paper offer testable implications. A laboratory design independently varying reward levels and stimulus distributions can evaluate three specific predictions: error convergence rate, the presence of neighborhood spillovers, and the divergence of maximum discrimination from the 50-percent choice threshold. Methodologically, descriptive choice models such as logit and probit typically treat the noise in choice as fixed across conditions. If attention instead adjusts to stakes and context, that noise may vary with the regressors. I plan to compare the estimation under my model with those from logit and probit fits in settings where stakes vary systematically.

Information Sharing and Network Formation

I also plan to study how organizations share information internally while limiting leakage. Baccara and Bar-Isaac (2008) model a criminal organization as a network of who holds evidence on whom, so that detecting one member exposes everyone whose information he holds. I would instead model the authority as a Bayesian court that convicts only when its posterior clears a standard of proof. What members reveal to one another then serves two inference problems with different thresholds: insiders need enough information to discipline each other, while the state needs enough to convict. Organizational design can then be cast as a multimarginal optimal transport problem. I want to study when threshold secret-sharing structures emerge endogenously, in which no small group of captured members can supply enough evidence to convict, and how the evidentiary standard determines whether organizations are hierarchical or diffuse. The answers has implications on corroboration requirements and leniency programs.

References

Baccara, M. and Bar-Isaac, H. (2008). How to Organize Crime. *Review of Economic Studies*, 75(4):1039–1067.

Cripps, M. W., Ely, J. C., Mailath, G. J., and Samuelson, L. (2008). Common Learning. *Econometrica*, 76(4).

Hébert, B. and Woodford, M. (2021). Neighborhood-Based Information Costs. *American Economic Review*, 111(10).

Kahneman, D. and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2).

Monderer, D. and Samet, D. (1989). Approximating Common Knowledge with Common Beliefs. *Games and Economic Behavior*, 1(2).

Rubinstein, A. (1989). The Electronic Mail Game: Strategic Behavior under "Almost Common Knowledge". *American Economic Review*, 79(3).

Stewart, N., Brown, G. D. A., and Chater, N. (2005). Absolute Identification by Relative Judgment. *Psychological Review*, 112(4).