Research and Teaching Statement

Kai Hao Yang

08/31/2025

I am an economic theorist working on mechanism and information design. My research has two main streams: (i) exploring the interplay between information, privacy, and markets; and (ii) developing integrated theoretical tools for various areas of economics.

Advances in information technology have reshaped economic activities. Firms now use granular consumer data to personalize pricing. AI-driven decisions raise concerns around fairness and privacy. And digital platforms restructure how markets function. My research develops models to understand these transformations and to inform policy. This includes formalizing the concept of privacy, analyzing the sale and use of consumer data, and examining the role of information intermediaries in digital markets. These models and methods address new problems but also yield breakthroughs on classical ones, including public goods provision, redistricting, and security design—placing them in a unified framework.

Information, Privacy and Markets

Privacy and Discrimination

In many settings, certain information, such as race, gender, or health is protected by law or norms. In **"Privacy-Preserving Signals" (Strack and Yang 2024, *Econometrica*)**, we study how to convey information without revealing protected characteristics. Consider two worker groups, A and B where group A is more productive on average. What information about productivity can be shared without revealing anything about group identity? Full disclosure would not be allowed, since a worker with high productivity is more likely to be in group A. However, revealing one's *quantile* within their own group—without revealing the group—preserves privacy. For example, knowing someone is at their group median reveals nothing about group identity, yet, when combined with group identity, it fully reveals productivity.

The paper characterizes all privacy-preserving signals—those leaving beliefs about protected characteristics unchanged—and shows that the most informative such signals are (reordered) *quantiles*: revealing the (possibly reordered) quantile of the state conditional on the protected characteristic. The notion of privacy here ties to statistical discrimination: in many settings, such as AI-assisted and data-rich decisions like loans and bail, concerns of discrimination stem from correlations between protected characteristics and other predictors. Our result prescribes a simple optimal step to eliminate statistical discrimination: transform raw predictions into their quantiles conditional on protected traits. For example, to make loan decisions, rather than raw default probabilities, using an applicant's default-probability percentile within their protected group is discrimination-free. Moreover, it is optimal: no other discrimination-free signal si more informative on the probability of default.

Non-Discriminatory Personalized Pricing

In addition to criminal justice and loan decisions, pricing also raises concerns about unintended discrimination: granular personal data let firms tailor prices, and consumers with different protected traits may face different prices, even when pricing rules "appear neutral on their faces" (12 CFR, §B). Such disparities are often prohibited in credit, insurance, and housing, making it crucial to identify personalized pricing rules that guarantee non-discriminatory outcomes. In **"Non-Discriminatory Personalized Pricing" (Strack and Yang 2025)**, we characterize the optimal pricing rule under a *non-discrimination* constraint, where price distributions faced by different protected groups must be identical.

The optimal rule charges low-value consumers in each protected group high prices, which allows the firm to tailor prices to high-value consumers and fully extracts their surplus, while preserving the overall price distribution faced by each group. Therefore, unlike standard screening, neither high- nor low-value consumers retain rents, while intermediate-value consumers do. Either the advantaged or disadvantaged group can benefit from anti-discrimination regulations, depending on the parameters. Furthermore, tightening the requirement from parity in offered prices to parity in transaction prices and outcomes reverses some effects, shifting surplus to advantaged consumers.

Information Intermediaries

Advances in information technology have created new business models. Online platforms and social media let firms advertise and deliver personalized information, while consumer data collection enables more effective pricing and design.

As data-driven pricing becomes easier, consumer data have become valuable. *Data brokers* collect and sell data to retailers to refine pricing. In **"Selling Consumer Data for Profit" (Yang 2022, *AER*)**, I model consumer data brokership, where a broker sells data to a producer with privately known production costs, who uses the data to tailor prices.

Because production costs affect optimal prices, producers value the same dataset differently. This makes it profitable for the broker to offer a menu of datasets, allowing producers to "self-select." Unlike classical versioning problems, where products vary along one dimension, datasets are high-dimensional: different combinations of variables yield distinct datasets. I characterize the broker's revenue-maximizing menu, which offers datasets enabling producers to target consumers above certain value thresholds, withholding information about those below. Higher thresholds reveal less, attracting higher-cost producers; richer datasets with lower thresholds attract lower-cost producers. This self-selection yields *quasi-perfect price discrimination*: all purchasing consumers pay their values, though some with values above cost do not buy. Under this optimal menu, consumers retain no surplus, and giving them property rights over their data is Pareto-improving.

While data brokers inform sellers, some intermediaries—such as platforms or influencers—provide product information to consumers. As their business models often depend on consumer engagement, they may also value consumer surplus. In **"Consumer-Minded Information Intermediaries" (Xu and Yang forthcoming, *RAND*)**, we show that greater emphasis on consumer surplus can reduce welfare for *all* parties from a hold-up effect: the seller tends to inefficiently raise price to counter the intermediary's incentive to benefit consumers. Meanwhile, some large platforms provide information to sellers *and* consumers. **"Equivalent Mechanisms for Information Intermediation" (Xu and Yang 2025)** shows that such platforms can achieve the same outcomes as direct price control by giving information only to consumers while revealing nothing about them to sellers.

Information and Welfare

Information shapes markets by altering demand and market structure. In **"Efficient Demand in a Multi-product Monopoly" (Yang 2021, *JET*)**, I characterize demand curves yielding Pareto-undominated welfare in non-linear pricing settings—those with *affine-unit elasticity*. In **"Regulating Oligopolist Competition" (Yang and Zentefis 2023, *JET*)**, we show that the welfare-maximizing market structure is equivalent to *yardstick price caps*, where each firm's cap depends on rivals prices. In **"Explaining Models" (Yang, Yoder, and Zentefis 2025)**, we analyze whether simplified representations of complex models

can improve welfare when full understanding is impossible, matching full understanding for utilitarian decision-makers, who cares only about the average, but offering no benefit for Rawlsian ones, who cares about the worst-case.

New Methods in Economic Theory

Motivated by modern challenges in privacy and information, I develop methodological tools that address new problems and unify classical results across domains. Rooted in convex analysis and stochastic ordering, these methods emerge from data markets and information design, yet yield breakthroughs into long-standing problems like redistricting and public goods, revealing shared structures and contributing to an integrated foundation for economics.

Monotone Function Intervals

The paper **"Monotone Function Intervals" (Yang and Zentefis 2024, *AER*)** is motivated by a canonical result in information theory: Blackwell's theorem (Blackwell, 1953; Strassen, 1965). For a one-dimensional state with a given prior, Blackwell's theorem characterizes the distribution of posterior means induced by *any* signal: it must be a *mean-preserving contraction* of the prior, and conversely, any such contraction can arise from some signal. This result has broad applications, but in some settings other statistics matter more. In voting, the median voter theorem highlights the median voter's position; in ranking or grading, quantiles often matter.

This paper develops an analog of Blackwell's theorem for quantiles. It characterizes the distributions of posterior *quantiles* induced by a signal: such a distribution arises iff it is bounded by two truncated versions of the prior in first-order stochastic dominance. This complements Blackwell's theorem and applies broadly, e.g., to legislative compositions through redistricting or identifying ranking data consistent with rational Bayesian updating.

The proof builds on a general characterization of extreme points of monotone functions bounded above and below by two functions. This result also yields insights in other fields: securities with limited liability correspond to such functions, so the characterization generalizes classical optimal security design results. The paper offers a unifying method connecting problems across information design, political economy, financial economics.

Multidimensional Monotonicity

Likewise, multi-dimensional monotone functions are also at the heart of many economic models. In multi-agent mechanism design problems, allocation rules are multi-dimensional, and incentive compatibility often implies monotonicity. Unlike the one-dimensional case, however, multi-dimensional monotone functions are far more complex, making mechanism design in multi-agent settings challenging, except for a few well-known settings.

The paper **"Multidimensional Monotonicity and Economic Applications" (Yang and Yang 2025)** provides a general characterization of the extreme points of multi-dimensional monotone functions, along with their one-dimensional projections. This characterization uncovers common structures underlying optimal mechanisms in diverse settings such as bilateral trade and auctions with endogenous values, and yields explicit solutions in problems where standard methods fail, including public good provision with ex-post incentive constraints.

These extreme points also relate to *sets of uniqueness* from mathematical tomography: sets uniquely identified by their projections. This connection yields an "anti-equivalence" result: two cornerstone equivalence theorems, Bayesian-dominance and stochastic-deterministic equivalence, are completely incompatible in certain settings.

Teaching

I teach components of two MBA core courses: *Basics of Economics* and *Competitor*. I emphasize rigor and practical applications. My lectures use interactive games along with examples drawn from my research in data and privacy. At the PhD level, I taught *Advances in Microeconomic Theory* in the Economics department, covering foundations and frontiers of mechanism and information design.

References

Strack, P. and K. H. Yang (2024): "Privacy-Preserving Signals," *Econometrica*, 92, 1907–1938.

——— (2025): "Non-Discriminatory Personalized Pricing," Working Paper.

Xu, W. and K. H. Yang (2025): "Equivalent Mechanisms for Information Intermediation," Working Paper.

——— (forthcoming): "Consumer-Minded Information Intermediation and Welfare Losses," *RAND Journal of Economics*.

Yang, F. and K. H. Yang (2025): "Multidimensional Monotonicity and Economic Applications," Working Paper.

Yang, K. H. (2021): "Efficient Market Demands in a Multi-Product Monopoly," *Journal of Economic Theory*, 197.

——— (2022): "Selling Consumer Data for Profit: Optimal Market-Segmentation Design and its Consequences," *American Economic Review*, 112, 1364–1393.

Yang, K. H., N. Yoder, and A. Zentefis (2025): "Explaining Models," Working Paper.

Yang, K. H. and A. Zentefis (2023): "Regulating Oligopolist Competition," *Journal of Economic Theory*, 212.

Yang, K. H. and A. K. Zentefis (2024): "Monotone Function Intervals: Theory and Applications," *American Economic Review*, 114, 2239–2270.