Research Statement (Full)

Jonathan Libgober

1 Overview

I am an information economist whose work explains how uncertainty and beliefs shape economic decisions. My research elucidates how the prospect of learning and information arrival shapes economic interactions and guides the design of incentives. My work has studied disparate applications, such as how monopolists set prices, how scientists decide which experiments to conduct, or why incorrect beliefs persist despite corrections. While these applications may seem unrelated, my work illustrates how they can be studied under a unified approach using common theoretical tools. My research uses these applications to motivate further development of foundational principles in decision theory, game theory, and mechanism design. All of my contributions have been either theoretical or theory-driven. Broadly, my agenda follows two main lines of research.

The first line examines learning in the presence of *model misspecification* and *robustness concerns*. This work continues the tradition of emphasizing how uncertainty and asymmetric information shape institutions and markets. These models are often criticized for assuming that individuals or organizations possess unrealistic knowledge of their environment, a feature apparently inconsistent with the reality that simpler heuristics are often used in practice. To accommodate this criticism, my work distinguishes between the *objective characteristics* of an economic environment and the *subjective perception* of that environment held by decision makers. My work on model misspecification essentially posits that individuals are mistaken about objective characteristics and do not know this is the case (He and Libgober 2026; He and Libgober 2025; Gonçalves, Libgober, and Willis 2026). My work on robustness instead allows individuals to maintain uncertainty about the objective characteristics, resolving this uncertainty under minimal confidence (Libgober and Mu 2021; Cho and Libgober 2025; Z. Li, Libgober, and Mu 2025).

The second strand of my research examines credibility in information transmission. My work in this area falls into a larger agenda aiming to bridge the gap between two main benchmark models of asymmetric information: The *hard information* benchmark under which information has concrete verifiable content, and the *soft information* benchmark under which any uncertainty is perfectly falsifiable. While many important contributions fall squarely within one of these benchmarks, most information in reality falls in a middle ground. A major application in some of my work on this topic (Libgober 2022; Abrams, Libgober, and List 2026) focuses on scientific research, where certain

outcomes—such as estimated empirical effects—may be difficult to manipulate undetected, even though the research process itself is impossible to fully document. Other projects examine forecasting (Y. Li and Libgober 2025; Libgober 2025), where forecaster information may be inherently difficult to manipulate but can be evaluated against observable outcomes. Additional work considers cases in which signal accuracy is uncertain (Libgober, Michaeli, and Wiedman 2026) or where transforming soft information into hard, verifiable data involves nontrivial costs (Z. Li and Libgober 2025). Across these settings, my work shows that the hard-soft dichotomy neglects important real-world phenomena.

Several themes have arisen while pursuing both of these lines of work, such as how the prospect of information arrival in the future shapes actions in the present, or how the method by which information is generated can influence its interpretation. In some cases, these themes have suggested new lines of work falling slightly outside of these two areas, but which can nevertheless be analyzed using similar methodology. But for all of these new areas of work, the consistent theme has been to elucidate how information plays a central role in shaping economic outcomes across markets, organizations, and institutions.

2 Learning with Misspecification and/or Robustness

As alluded to above, I divide my work into those that anticipate the lack of knowledge of the true environment (Robustness) and those that do not (Misspecification).

2.1 Robustness

My work on learning originated from my interest in addressing the strong assumptions of economic models about the knowledge agents possess of an underlying informational environment. Several papers in this line of work are inspired by ideas from computer science that consider optimal or near-optimal performance guarantees without relying upon strong assumptions of correctly specified Bayesianism.

This interest started with my job market paper, *Informational Robustness in Intertemporal Pricing* (Libgober and Mu 2021, **Review of Economic Studies**). This paper considers a classic problem: how should a monopolist, able to commit to a selling strategy, post prices over time to maximize profit? In particular, should prices change or not? Stokey 1979 showed that while lowering prices may allow the seller to capture new buyers, doing so may also incentivize other buyers to delay. In the simplest setting where future purchases are discounted independently of willingness-to-pay, the tradeoff resolves in favor of the monopolist using a constant price path. Our starting point was to ask whether the same conclusion holds if buyers might be learning about their willingness-to-pay over time. In principle, this possibility could have dramatic impacts: the possibility of learning could encourage the buyer to "wait and see," but might also lead the seller to expect that they would be able to extract more revenue from buyers who learn that their willingness-to-pay is actually

quite high. We point out that the elegance of the Stokey 1979 result—which does not depend sensitively on hard-to-measure variables such as the discount factor or the exact value distribution—is lost.

We instead approach the problem by considering the robust approach described above: Rather than imagine the monopolist trying to maximize profit against a *particular* way the buyer may learn, we consider the case where they do not know the learning process and try to maximize the profit level that can be guaranteed, no matter how they actually learn. We show that under this objective, we recover the well-known result on the optimality constant price paths. The particular price used will differ from the case where the buyer already knows her value, but we show an analogy between the solution to our model with a single period and the classic monopoly pricing problem under known values. After establishing this main benchmark, we seek to provide qualitative descriptions of when it may fail, in order to guide when one might expect sellers to depart from constant prices. This work was the first paper to show that the worst-case approach could be used to tractably handle complexities associated with dynamic information arrival in a design setting; to my knowledge.

Our application of the robust approach to a dynamic setting was (at the time) relatively novel; on the other hand, the assumption of commitment was necessary not only for the economic substance of our findings, but also for the coherence of our approach. To be more precise, the particular way of learning that yields the profit guarantee against a given price path at one time might imply a relatively high profit level at a later time; conversely, the learning process yielding the profit guarantee starting at some later time might be relatively favorable at an earlier time. Such *dynamic inconsistency* issues associated with such an objective were well-understood in single-agent decision problems, but relatively underexplored in settings where a designer chooses a policy.

I explored these issues in *Sequentially Optimal Pricing under Informational Robustness* (Z. Li, Libgober, and Mu 2025, Working Paper), in order to see if any further economic insights could be delivered by extending the robust approach to this domain. Specifically, we formulate the robust objective without commitment by considering the benchmark where the seller correctly anticipates what the "worst case" will be across different times. We do so by formulating this as a zero-sum game against nature; while such devices are common in the literature, our novelty was to formulate dynamic consistency in terms of the commitment power of this player. We show that the worst-case learning process minimizes the seller's payoff period by period, yielding sharp prescriptions about economic parameters such as buyer and seller surplus. Perhaps more surprisingly, we show that, under a permissive class of willingness-to-pay distributions, this learning process remains worst-case as long as the seller does not anticipate a departure to some other learning process. Taken together, our results provide an optimality justification for why sellers might not need to consider the details of a learning process when setting prices, as we identify a close connection between outcomes in these settings and others that do not feature learning.

While both of these papers considered a seller choosing prices to optimize this particular objective, in practice, sellers often resort to algorithms to do so on their behalf, aiming to choose one that performs "well" according to some notion. It is difficult to reconcile the use of such algorithms, which are often simple relative to the complexity involved with pricing decisions, with the focus on optimality in such models.

I address this issue in *Learning Underspecified Model* (Cho and Libgober 2025, **Journal of Economic Theory**), which considers a seller's choice of a pricing algorithm. A key innovation of this paper is formulating the seller's objective using the same robustness notions in the above work. We suppose the seller wants to learn the true demand curve "quickly," at a rate that is uniform over the set of possible demand curves considered. We also formulate a notion of complexity in terms of the number of parameters the algorithm must keep track of. We show that it is possible to achieve the seller's goal using only very limited data relative to the set of possibilities available in principle: Rather than keep track of the entire history of sales, the seller can learn just as well by keeping track of only two parameters. The main conclusion is that simple algorithms may appear to perform just as well as more complex ones.

2.2 Misspecification

Most of the rest of my work focuses on situations where all parameters of an economic actor's environment are specified, but mistakes persist. Examples of such mistakes in practice are ubiquitous: prominent examples include beliefs about vaccine efficacy and side effects or political misinformation. Such mistakes are often accompanied by an apparent *retraction failure*: some information is disseminated, but later shown to be inaccurate, yet the retraction seems less effective at dispelling the belief compared to the initial information propagating it. From the perspective of economic theory, a natural question is whether such failures as reflecting a departure from objectively rational belief formation—some departures from which have been consistently demonstrated. While many of these illustrations were done in psychology and political science, to my knowledge, no other work had considered directly testing in a setting resembling a simple economic model.

My paper *Retractions: Updating from Complex Information* (Gonçalves, Libgober, and Willis 2026, **Review of Economic Studies**) presented results from an experiment performing this very test. Our goal was to understand whether theoretical models of information processing necessarily should consider the method by which information is generated, using retraction failure as a case study to illustrate why this is indeed the case. In an online experiment, we presented subjects with balls from a box, replacing each ball after showing it to the subjects. The box contained a number of "noise balls," half of which were yellow and half of which were blue, or a "truth ball" which was yellow or blue with equal probability. After showing a number of these draws, we would then tell subjects whether a randomly selected draw was the truth ball or a noise ball. After each

piece of information was given, we would ask subjects for their beliefs about the color of the truth ball. We referred to a "retraction" as the event that one of their previous draws was shown to be noise.

For a rational Bayesian, being told a ball was a noise ball would result in that draw being ignored. However, we found, with striking consistency, that subjects would treat the retraction as less informative compared to a new draw. We show that this finding is not explained by many well-known biases in the literature. At the same time, we describe why it is natural to view retractions as relatively more complex, and provide indirect evidence in our data consistent with the complexity explanation. Our findings suggest that practitioners should not necessarily expect context-specific factors to prevent retraction failures, and instead that this should be the norm, rather than the exception, for most views held by the public.

My other papers on misspecification address the question of whether there are performance-based motivations for the emergence of misspecifications. This approach has a long tradition in the literature on evolutionary game theory, and this particular question has been studied in what is known as the indirect evolutionary approach. Whereas traditional models of evolutionary fitness focus on which strategies might emerge in a population, the indirect evolutionary approach instead posits that evolution acts on preferences. This literature then studies the properties of preferences that might emerge as a function of the relevant environment.

In *Misspecified Learning and Evolutionary Stability* (He and Libgober 2025, **Journal of Economic Theory**), I studied how to extend this approach to cases where evolution selects models rather than preferences, in a sense taking the indirect evolutionary approach one step further. The motivation for doing so was inspired by recent work studying misspecified learning, with the idea being that a decisionmaker who is misspecified about some parameter of the environment may rationalize this by inferring other parameters incorrectly. This rationalization is unique to model inference, and our paper identifies certain evolutionary phenomena that might emerge only when this inference ability is present.

While this paper discussed the implications of our framework on a few applications, a more thorough application is presented in *Higher-order Beliefs and (Mis)learning from Prices* (He and Libgober 2026, **American Economic Journal: Microeconomics**). We apply the framework from the previous paper to the selection of misspecified belief formation, which several papers (using similar paradigms as my retractions paper) have shown is a general and robust finding. In particular, we consider a setting where duopolists with private information about demand may be incorrect about the correlation of these signals with other players. Our motivation for doing so originated in part from scrutinizing the claim in Friedman 1953 that evolutionary pressures should select at least "as-if rational" payoff maximization. In our framework, we find that whether rationality is favored over a bias depends on whether inference is possible. Specifically, exaggerations of signal

correlation are beneficial when elasticity is learned from prices, but harmful when dogmatic beliefs about elasticity are held. This claim follows from studying how the misinference of parameters allows players to make beneficial commitments. We thus show the limitations of the logic from Friedman 1953, and that, in particular, market forces may favor misperceptions of information.

3 Credibility in Information Transmission

My work on credibility can similarly be divided into two different threads: One focusing on contracting problems where information is revealed over time, while another studies market inferences in settings where neither fully hard nor fully soft information is conveyed.

3.1 Dynamic Contracting Without Transfers

Work that has studied the design of allocation mechanisms typically starts with the presupposition that preferences or information (what is referred to as "types") are privately known, with the designer concerned with how to do so optimally. Combining this observation with the previous point that information is often partially but not fully verifiable, the question naturally arises how the movement away from complete falsifiability influences the design of mechanisms. This problem has been recognized as important in the literature on corporate finance and organizational governance. My particular contributions have sought to extend other work studying this question to settings where dynamics play a significant role.

This first paper on this thread concerns the case where the information of interest to the designer can be observed by paying a fixed cost. Taking this cost to be arbitrarily large nests the soft-information benchmark, while taking this cost to be arbitrarily small nests the hard-information benchmark. In *The Dynamics of Verification when Searching for Quality* (Z. Li and Libgober 2025, Second revision requested at **Review of Economic Studies**), we solve for how a designer should perform costly verification when able to do so in conjunction with specifying a final decision rule. The particular problem we study is where the designer seeks to make a selection from some set of possibilities, aiming to choose one of high quality. The agent, by contrast, observes quality but does not benefit from it, hoping only that some selection is made. The agent searches for a possibility over time, and the question of interest is how the mechanism should unfold as a result.

When the verification cost is sufficiently low (but nonzero), the optimal mechanism is characterized by a property which we dub decreasing skepticism: The verification probability decreases over time, until a deadline (chosen by the designer) at which any selection is made with probability 1. The intuition is simple: The agent must be rewarded for truthfully reporting quality, implying that eventually allocation will occur—and since the cost of verification is nonzero, it is not possible to search forever. However, to deter the detection of a lie, the best punishment available is to instead refuse to ever make a selection, inducing the agent to tell the truth. Since verification is costly, the designer seeks to do so with as low a probability as possible. So, the closer to the deadline the

interaction is, the lower this probability needs to be, yielding the finding. While this intuition is simple, we show that it hinges critically on the cost of verification not being too high, in which case the principal would not want to always utilize it, and also on the horizon being long, in which case not enough periods may be available to make a promised allocation worthwhile. We relate these findings to questions in corporate governance, showing that oversight should be expected to evolve dynamically as part of a project-selection process.

Incentivizing Forecasters to Learn: Summarized vs. Unrestricted Advice (Y. Li and Libgober 2025, Working Paper) studies the case in which the designer can instead condition rewards on an ex-post state, and aims to elicit the information about this state. Our leading application of this model is forecasting: In applications where a forecast is sought, the natural goal is to obtain as much information about this future event as possible. To induce the forecaster to actually attempt to learn about this future event, the designer can condition a recommendation or endorsement on whether or not their prediction is correct, or do so with some probability. The question of interest is how to assign these rewards to learn as much as possible. Crucially, these rewards are of fixed value and cannot be adapted flexibly. While the computer science literature has studied how to elicit information using this particular class of contracts, our innovation was to use such a framework to speak to how to provide the best incentives for dynamic information acquisition.

Our main question of interest is whether the designer should try to obtain continuous reports from the agent, or if a simple, one-time report suffices to incentivize the most information acquisition. We show that the answer is subtle: in principle, the best way of providing rewards should depend on the forecaster's belief, but the problem is that this belief may move around over time. Despite this, we identify ways of maximizing incentives for information acquisition using static elicitation when obtaining a piece of information moves beliefs in one direction, or when it is perfectly informative. We also show that otherwise, dynamic elicitation may be necessary. We relate these findings to the question of when richer contracting capabilities are more conducive to more informed forecasts, which is relevant to understanding how consultation and forecaster relationships should be structured.

3.2 Inferences with Partially Verifiable Information

My other papers in this agenda focus on modelling particular applications or scenarios of interest, rather than design problems. Two of these papers are focused on a particularly meaningful application, the incentives of scientific researchers. While both of my papers on this topic study particular institutional design choices, these are not flexibly set as in the prior line of work. Instead, the main observations compare outcomes under different scenarios and speak to policy questions by determining which of these regimes yields preferable outcomes.

The first of these two papers, *False Positives and Transparency* (Libgober 2022, **American**

Economic Journal: Microeconomics), discussed transparency requirements. The main consideration was whether extensive documentation of research outcomes (e.g., the number of tests considered) would yield more socially advantageous experimentation. I show that the answer to this question is perhaps more subtle than many would have assumed. The naive intuition suggests that less transparency yields more opportunities to create bias. While this indeed is an important consideration, my paper points out that researchers also have a natural incentive to counteract the perceptions of bias. In some cases, this might lead to more informative experiments. Formally, I developed a particular model of strategic experimentation choice, where a "sender" chooses an experiment parameterized by multiple dimensions. I provided conditions such that keeping some of these dimensions hidden would induce the sender to choose more informative experiments, ultimately helping the receiver. I connected this observation to some policy-oriented discussions regarding transparency requirements, thus bridging the gap between the theoretical modelling and applied messages.

The second of these two papers focused on a related question, the design of research registries for experimental research. *Research Registries and the Credibility Crisis: An Empirical and Theoretical Investigation* (Abrams, Libgober, and List 2026, **Economic Journal**) presented a novel theoretical model of research registration, together with an empirical documentation effort aimed at assessing the current effectiveness of registration in economics. Empirically, we focused on evaluating the AEA RCT Registry, finding that despite requirements, preregistration is not as dominant as one might have expected (although we also documented changes in this over time), and that many registrations leave significant latitude. We also compared our findings to ClinicalTrials.gov and documented similar patterns.

Theoretically, we focused on the question of whether allowing for post-registration could mitigate the incentives for researchers to register. It is immediate to think why it would not: If registration decisions were made without consideration of incentives, then one would expect that more opportunities to register would increase registrations, even if after the experiments were conducted. While it is less clear whether these registrations are as socially valuable, in fact, we show a more dire possibility when researchers are strategic: under some natural restrictions on what researchers know about experiment outcomes before conducting them, allowing for late registration can crowd out preregistration. The idea is simple: A researcher who does not know the outcome of an experiment might prefer to keep their options open by delaying registration, even despite the market potentially making a negative inference as a result. So, allowing for late registration can crowd out preregistration, and in particular, leave open the possibility that researchers never register if negative results are obtained. While this possibility requires particular parameters, perhaps surprisingly, a numerical calibration exercise finds uniformly that a late registration ban would increase registrations. Consistent with this theory, the *Journal of Political Economy: Microeconomics* decided to break from the AEA journal, which allowed for postregistration of research studies, instead mandating preregistration—but, using a much more flexible list of possible venues, in line with our analysis suggesting this would be important.

My other paper on this topic considered a similar theoretical structure, but with a different baseline model (disclosure of verifiable information, rather than signalling) and a different leading application (firm valuation, as is typical for this application). Specifically, *With a Grain of Salt: Investor Reactions to Uncertain News and (Non)Disclosure* (Libgober, Michaeli, and Wiedman 2026, **Journal of Accounting and Economics**) develops a model where a manager can disclose firm value to the market, but may also remain silent. The market, by contrast, does not know if the manager is silent because there is no opportunity to disclose, or because the information would be unfavorable. Such models have a long tradition in both economics and accounting. Our twist was to consider the case where, after the disclosure decision is made, the market would observe a signal which could be either accurate or inaccurate. The key innovation here was to focus on the implications of this uncertainty about the accuracy of the information. We showed that this innovation might result in prices being nonmonotonic in news quality: Even as signals became more favorable, the market would interpret the firm to be of lower value.

The driving force behind this paradoxical finding is the manager's strategic incentives. Markets understand that managers with the highest firm values would prefer to disclose, although in some cases, disclosure might not occur for exogenous reasons. Nevertheless, if the news were favorable, the market would understand it is less likely to be noise. This is asymmetric with unfavorable news, which the manager would have preferred to keep hidden. Thus, favorable news is discounted in a way that unfavorable news is not. We show that the possibility of favorable news incentivizes managers to "piggyback" on its realization, implying that when the market places more weight on the manager being uninformed, rather than the signal being accurate, the market price declines. We describe how this asymmetric reaction to news may explain potentially surprising empirical findings, thus suggesting that these forces may underlie important factors determining asset market prices.

4 Other Work

In addition to these projects, I have actively attempted to push this agenda into new directions, yielding some contributions that do not fit neatly into the above categories but which I nevertheless hope may spur interesting questions. As the above conveyed, my work has touched on themes regarding the theoretical models of dynamic information arrival, "information about information," and mechanism design without transfers. These interests have led me to pursue projects exploring each independently.

My paper *Identifying Wisdom (of the Crowd): A Regression Approach* (Libgober 2025, **Journal of Political Economy: Microeconomics**), considered the problem of how an analyst might be able to infer the way a population of individuals forms beliefs about uncertain events. This problem is related to the literature on crowd wisdom, inspired by (Prelec, Seung, and McCoy 2017), which aims to develop procedures for uncovering a "ground truth" based on the population of beliefs.

My paper, instead, focused on the problem of inferring the informational environment, essentially dictating how beliefs would be formed. This object is a fundamental one in several lines of work above, and so this project was focused on whether it would be possible to identify it from data. My work joined a small number of other papers in considering what could be done when the analyst also had access to the population's belief distribution over what others believe (Chen, Mueller-Frank, and Pai 2021; McCoy and Prelec 2025). I describe a striking procedure that would work to do so when there are more possible beliefs of experts than possible states, which stems from the observation that the information structure has a natural interpretation as the "best fitting line" of the beliefs about others as a function of the beliefs about the state. As described in McCoy and Prelec 2025, this procedure may facilitate the aggregation of expert predictions to arrive at more informed forecasts.

I also study the implications of dynamic learning in my paper *Electric Vehicle Sharing: Crowding Adoption Out or In?* (Libgober and Song 2025, Under revision at **Journal of Public Economics**). This paper considered the impact of a car-sharing program, BlueLA, on the adoption of electric vehicles. This program was unique compared to other carsharing programs in that it was a public-private partnership, allowing it to target lower-income households, and also in that all cars were electric vehicles. The hypothesis is that increased familiarity with these vehicles would spur an increase in its adoption; while a natural hypothesis, empirical methodologies which focus on dynamic learning can create challenges in identification, stemming from the need to use differences in the timing of the informational intervention (in our case, BlueLA entry) to separate out the impact of the innovation as opposed to other changes in the EV market which occurred over the same study period. Despite this, we were able to document a causal impact of BlueLA on adoption, and also document a number of other patterns consistent with the hypothesized informational channel.

Both of these projects naturally grew out of my two main lines of research—the former started as an exploration of how the informational environment could be observed from data, and the latter related to studying the observable implications of dynamic learning. Going forward, I am extending this agenda to organizational design, competition, and mechanism design under informational complexity. One line of work I plan to pursue is to extend models of individual decision making, which underlie much work on misspecification and complexity, to organizational settings. These questions are related both to the design of organizations and to how contracting between firms may work in light of these departures from correctly specified Bayesianism. I am also hoping to bring the models of intermediate verifiability of information to settings in more traditional mechanism design environments. These directions suggest a set of open questions about how learning, complexity, and information design interact in markets and institutions, which I plan to pursue in future work.

References

Abrams, Eliot, Jonathan Libgober, and John A. List (2026). "Research Registries and the Credibility Crisis: An Empirical and Theoretical Investigation". In: *Economic Journal*. Forthcoming.

Chen, Yi-Chun, Manuel Mueller-Frank, and Mallesh M. Pai (2021). "The Wisdom of the Crowd and Higher-Order Beliefs". Working paper.

Cho, In-Koo and Jonathan Libgober (2025). "Learning Underspecified Models". In: *Journal of Economic Theory* 226, p. 106015.

Friedman, Milton (1953). *Essays in Positive Economics*. University of Chicago Press.

Gonçalves, Duarte, Jonathan Libgober, and Jack Willis (2026). "Retractions: Updating from Complex Information". In: *Review of Economic Studies* 93.1, pp. 476–516.

He, Kevin and Jonathan Libgober (2025). "Misspecified Learning and Evolutionary Stability". In: *Journal of Economic Theory* 230, p. 106082.

— (2026). "Higher-Order Beliefs and (Mis)learning from Prices". In: *American Economic Journal: Microeconomics*. Forthcoming.

Li, Yingkai and Jonathan Libgober (2025). "Incentivizing Forecasters to Learn: Summarized vs. Unrestricted Advice". Working paper.

Li, Zihao and Jonathan Libgober (2025). "The Dynamics of Verification when Searching for Quality". In: *Review of Economic Studies*. Second-round R&R.

Li, Zihao, Jonathan Libgober, and Xiaosheng Mu (2025). "Sequentially Optimal Pricing under Informational Robustness". Working paper.

Libgober, Jonathan (2022). "False Positives and Transparency". In: *American Economic Journal: Microeconomics* 14.2, pp. 478–505.

— (2025). "Identifying Wisdom (of the Crowd): A Regression Approach". In: *Journal of Political Economy: Microeconomics* 3.4, pp. 798–826.

Libgober, Jonathan, Beatrice Michaeli, and Elyashiv Wiedman (2026). "With a Grain of Salt: Investor Reactions to Uncertain News and (Non)Disclosure". In: *Journal of Accounting and Economics*. Forthcoming.

Libgober, Jonathan and Xiaosheng Mu (2021). "Informational Robustness in Intertemporal Pricing". In: *The Review of Economic Studies* 88.3, pp. 1224–1252. DOI: 10.1093/restud/rdaa046.

Libgober, Jonathan and Ruozi Song (2025). "Electric Vehicle Sharing: Crowding Adoption Out or In?" R&R, Journal of Public Economics.

McCoy, John and Drazen Prelec (2025). "Belief Decomposition: A Mechanism for Collective Inference". Working paper.

Prelec, Drazen, H. Sebastian Seung, and John McCoy (2017). "A Solution to the Single-Question Crowd Wisdom Problem". In: *Nature* 541, pp. 532–535.

Stokey, Nancy L. (1979). "Intertemporal Price Discrimination". In: *The Quarterly Journal of Economics* 93.3, pp. 355–371.