Research Statement (Abbreviated)

Jonathan Libgober

I am an information economist whose work explains how uncertainty and beliefs shape economic decisions. My research elucidates how the prospect of learning—particularly in dynamic environments—affects economic interactions and informs incentive and policy design.

1 Simplicity and Complexity Beyond Bayesianism

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, while my work on robustness allows individuals to instead maintain uncertainty about the objective characteristics, instead resolving this uncertainty under minimal confidence.

This interest started with my job market paper, *Informational Robustness in Intertemporal Pricing* (Libgober and Mu 2021, **Review of Economic Studies**), which considered the problem of how a monopolist seller should set prices over time when unsure of how a representative buyer might be learning about their willingness-to-pay. We show that without any restrictions on the learning process, constant price paths are optimal, but also that pricing dynamics could emerge if the seller ruled out some learning processes ex-ante. For instance, increasing prices can help the seller in the presence of known informational externalities. A key innovation of this paper was accommodating designer uncertainty over learning—formally, a "worst-case objective"—in a dynamic setting. Dynamic information arrival can be notoriously difficult to study, since information places difficult constraints on how perceived willingness-to-pay can evolve over time. While prior work had suggested that worst-case objectives often allow analysts to restrict attention to instead focus on simple individual possibilities rather than complicated beliefs over them, our innovation was in showing that the same idea could be used to make dynamic analyses tractable.

At the same time, the assumption of commitment was necessary not only for the economic substance of our findings, but also for the coherence of our approach. It was previously known

that worst-case objectives with *dynamic decisions* can lead to *dynamic inconsistency*, in which the worst-case possibility at one point in time may be inconsistent with that at other points in time. While such issues do not arise if the seller chooses prices only once (as with commitment), they do arise if the seller chooses a new price at each point in time. In *Sequentially Optimal Pricing under Informational Robustness* (Z. Li, Libgober, and Mu 2025, Working Paper), my coauthors and I argued that these issues may be less severe than previously thought in the context of dynamic pricing. Specifically, we formulate the robust objective without commitment by considering a benchmark in which the seller correctly anticipates the "worst case" at different times. We do so by framing this as a "game against nature;" while such devices are common in the literature, our novelty is to frame dynamic consistency in terms of this player's commitment power. 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.

My other work featuring departures from rational Bayesianism instead considers the case where individuals are mistaken about objective characteristics and do not know this is the case, rather than being aware of their uncertainty and seeking a favorable guarantee. One challenge in this agenda is identifying why a departure from seemingly optimal behavior necessarily reflects a failure of rational Bayesianism, rather than some other context-specific factor. One particular ubiquitous case where such apparent departures arise relates to *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. While some work in psychology and political science had shown that such behavior can arise in particular contexts (e.g., fake news in politics or vaccine misinformation), to my knowledge, no other work has 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. Using a simple design where subjects were asked to assess the composition of an urn after sampling with replacement, we implemented retractions by presenting information such that subjects should "disregard" some draws. We found that subjects treat the retraction as less informative compared to a new draw. We argued why this finding was indeed novel relative to other well-known biases in the literature, and interpreted the implications of our findings for policy.

Another contribution to the literature on misspecification addresses 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. *Higher-order Beliefs and (Mis)learning from Prices* (He and Libgober 2026, **American Economic Journal: Microeconomics**) applies this framework to the selection of misspecified belief formation, which several papers 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. However, we find that whether rationality is favored over a bias depends on whether inference is possible. Exaggerations of signal correlation is beneficial when elasticity is learned from prices, but harmful when dogmatic beliefs about elasticity are held. Thus, we identified—within a canonical framework—how the identified misinference channel can interact with the ability to make the beneficial commitments described in past work on the indirect evolutionary approach.

2 Credibility in Information Transmission

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. 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. The dynamic aspect of the problem was a key innovation in our model—in particular, the fact that the principal

could *either* make a selection tomorrow or today, and focusing on how the principal should optimally use costly verification as opposed to modifying the selection rule. 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. We also show that this result 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) further considers dynamic contracting in cases where some aspect of private information can nevertheless be verifiably contracted on. In this paper, we study a designer who can condition rewards on an ex-post state, but aims to elicit the information about this state beforehand. Our leading application of this model is to forecasting: In applications where a forecast is sought, the natural goal is to obtain as much information about this future event as possible. 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.

One particular application of interest in my work, where the credibility of information is a significant factor in institution design, is the analysis of the incentives faced by scientific researchers. While both of my papers on this topic study particular institutional design choices, these are not flexibly set as in either Z. Li and Libgober 2025 or Y. Li and Libgober 2025. 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, one of the first to concretely speak to these issues using a communication model. The main consideration was whether extensive documentation of research outcomes (e.g.,

the number of tests considered) would yield more socially advantageous experimentation. While naive intuition suggests that less transparency yields more opportunities to create bias, 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. By making this point formally, I articulate precisely why 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, bridging the gap between the theoretical modelling and applied messages.

The second of these two papers, *Research Registries and the Credibility Crisis: An Empirical and Theoretical Investigation* (Abrams, Libgober, and List 2026, **Economic Journal**), focused on the design of research registries. 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 at all. We show that when researchers are strategic, then under some natural restrictions on what researchers know about experiment outcomes before conducting them, allowing for late registration can crowd out preregistration. Our formal model allowed for precise articulations of the tradeoffs previously described informally in discussions about requirements for research conduct. We articulated why allowing for late registration could actually leave open the possibility that researchers never register in cases where negative results are obtained. Consistent with this theory, the *Journal of Political Economy: Microeconomics* decided to break from the AEA journals, 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.

3 Future Directions

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.

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.

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.