

Information Design in Misspecified Social Learning Models

Aislinn Bohren and Daniel N. Hauser

Extended Abstract

January 2018

Information plays a key role in shaping economic behavior and market outcomes. Faced with uncertainty about the quality of a product, the technology of a competing firm or the likely outcome of a policy, individuals often observe information before making a decision. A rich literature explores how long-run learning outcomes depend on the type of information that individuals observe and the way that individuals process this information. This literature establishes that when individuals use misspecified models to interpret information, long-run inefficiencies may arise and persist, even when sufficient information arrives for a correctly specified agent to identify the state.¹ In such settings, an information designer, such as a public health agency or government, may have scope to *strategically* release information to influence long-run behavior.

To fix ideas, suppose that individuals are learning about whether a new vaccine is safe. Each individual observes a private signal about the safety of the vaccine, as well as the vaccination choices of a sequence of predecessors. They correctly interpret their private signals, but have a misspecified model of how other individuals' actions reflect their private information. In particular, they are naive and believe that all prior actions solely reflect private information. They do not account for the fact that prior individuals also observed the actions choices of their predecessors, and therefore, these actions reflect some repeated information. We know from [Eyster and Rabin \(2010\)](#) and [Bohren \(2016\)](#) that incorrect learning can arise in such settings – in particular, when the vaccine is safe, individuals may come to believe that the vaccine is almost surely unsafe if the first few individuals have a negative signal about the vaccine's safety. Now suppose that an information designer, such as a public health agency, seeks to ensure that individuals accurately learn. Each period, the public health agency can release a costly public signal about the vaccine's safety, such as a public service announcement or online promotion. What is the optimal way to release information? In particular, how does the *timing*, *precision* and *frequency* of the information campaign depend on how severely individuals overweight the informativeness of prior actions?

In this paper, we study the optimal way for an information designer to release public

¹For example, inefficient learning can arise when individuals have confirmation bias ([Rabin and Schrag 1999](#)), overweight information ([Epstein, Noor, and Sandroni 2010](#)), misunderstand the correlation between other individuals' actions ([Bohren 2016](#); [Eyster and Rabin 2010](#)), use a level-k model of inference ([Bohren and Hauser 2017](#)) or systematically slant information ([Bohren and Hauser 2017](#)).

signals in a sequential social learning setting with model misspecification. A sequence of agents observe a private signal about a hidden state and the actions of their predecessors, then make a one-shot decision. In each period, the designer can release a signal whose precision she chooses, where the cost of this signal is increasing in the precision. In order to interpret signals, agents need a model of the signal generating processes, and in order to interpret prior actions, agents need a model of how such actions reflect other agents' information. We allow for agents who have a misspecified model of signals and other agents' decision rules (Bohren and Hauser 2017). This framework can capture many common information processing biases, including confirmation bias (Lord, Ross, and Lepper 1979), overconfidence (Moore and Healy 2008) and motivated reasoning (Kunda 1990), as well as misspecified models of strategic interaction such as cognitive hierarchy (Camerer, Ho, and Chong 2004) and the false consensus effect i.e. misspecified beliefs about others' payoffs (Ross, Greene, and House 1977).

When agents have misspecified models, in the absence of an informational intervention, they may learn the incorrect state, have beliefs that perpetually oscillate, or asymptotically disagree with each other, despite observing the same information (Bohren and Hauser 2017). We establish that in such settings, an information designer may be able to influence long-run learning by releasing public signals – either as a social planner, who seeks to facilitate efficient actions choices, or a self-interested agent, who seeks to sway agents towards a particular action. We restrict attention to settings in which sufficient information arrives for asymptotic learning to obtain in the correctly specified model – in other words, correctly specified agents will eventually choose the efficient action almost surely, and there is no scope for an information designer to influence long-run behavior. This allows us to isolate the impact of misspecification on information design.

The designer faces a trade-off between the precision of information to release, the timing of the intervention and the frequency of the intervention. How she makes this trade-off depends crucially on the form of misspecification. For instance, when facing naive agents who believe prior actions solely reflect private information, as in the vaccine example above, releasing imprecise signals is very costly. Agents' naive model will cause any potentially misleading signal to be vastly overweighted by future agents. But if the designer fails to intervene early, inefficient action choices can accumulate and cause beliefs to become entrenched on the incorrect state, making it harder and harder to overcome with future information releases. Therefore, it is crucially important that any intervention is immediate and precise to minimize the likelihood that agents become entrenched on an inefficient action choice. However, once agents' behavior converges

to the efficient action, no further intervention is needed – an infrequent intervention is sufficient to achieve efficiency. But this intervention would be ineffective for agents with other forms of misspecification. For instance, suppose agents have a misspecified model in which their preferences are biased towards one action, but they do not take this into account when interpreting others’ actions. In this environment, misspecified inferences drawn from the actions of others will eventually overcome any single precise signal, and it is much more important for the designer to intervene frequently, but less important for the intervention to be precise or immediate.

These results provide a better understanding of how to design informational interventions in complex learning environments. They yield insight into how to design policies that improve the accuracy of learning and combat misinformation.

References

- BOHREN, A. (2016): “Informational Herding with Model Misspecification,” *Journal of Economic Theory*, 222–247.
- BOHREN, A. AND D. HAUSER (2017): “Bounded Rationality And Learning: A Framework and A Robustness Result,” *PIER Working Paper*.
- CAMERER, C. F., T.-H. HO, AND J.-K. CHONG (2004): “A Cognitive Hierarchy Model of Games,” *The Quarterly Journal of Economics*, 119, 861–898.
- EPSTEIN, L. G., J. NOOR, AND A. SANDRONI (2010): “Non-Bayesian Learning,” *The B.E. Journal of Theoretical Economics*, 10.
- EYSTER, E. AND M. RABIN (2010): “Naïve Herding in Rich-Information Settings,” *American Economic Journal: Microeconomics*, 2, 221–243.
- KUNDA, Z. (1990): “The case for motivated reasoning.” *Psychological bulletin*, 108, 480.
- LORD, C. G., L. ROSS, AND M. R. LEPPER (1979): “Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence.” *Journal of personality and social psychology*, 37, 2098.
- MOORE, D. A. AND P. J. HEALY (2008): “The trouble with overconfidence.” *Psychological review*, 115, 502.

RABIN, M. AND J. L. SCHRAG (1999): "First Impressions Matter: A Model of Confirmatory Bias," *The Quarterly Journal of Economics*, 114, 37–82.

ROSS, L., D. GREENE, AND P. HOUSE (1977): "The "false consensus effect": An ego-centric bias in social perception and attribution processes," *Journal of experimental social psychology*, 13, 279–301.