

Abstract: I study the strategic interaction between a benevolent sender (who provides data) and a self-interested narrator (who interprets data), who compete to persuade a boundedly rational receiver (who takes action). The receiver, who does not know the data-generating model, chooses between the models provided by the sender and the narrator, selecting the one that best fits the data given her prior belief. The sender faces a trade-off between providing precise information and minimizing misinterpretation. Surprisingly, full disclosure can be suboptimal and even backfire. I identify a finite set of models that contain the optimal data-generating model, which maximizes the receiver's expected utility. The sender can guarantee non-negative value of information, preventing harm from misinterpretation. I apply this framework to information campaigns and employee feedback.