Inquiry under bounds

Humans are bounded agents. We have limited cognitive abilities and we must ration scarce cognitive and physical resources to achieve the best outcomes possible given our abilities. The study of *bounded rationality* asks what rationality comes to for bounded agents.

Rational choice theory is typically *outcome focused*, asking how features of beliefs and actions count in favor of adopting them. However, cognitive resources and abilities constrain not only the beliefs and actions we adopt, but also the process of inquiry by which we adopt them. For this reason, bounded rationality is *process focused*, focusing on the higher-order question of how rational agents should inquire.

My project is to develop an account of rational inquiry for bounded agents. I carry out this project in four parts. Part 1 introduces the project and raises a skeptical challenge: empirical work in judgment and decisionmaking seems to suggest that humans often inquire irrationally. I suggest that existing epistemological views are unlikely to resolve the skeptical challenge, but argue that we can make good progress on the skeptical challenge by recasting empirical observations as the results of rational inquiry by bounded inquirers.

Part 2 develops a consequentialist account of rational inquiry for bounded agents. On this account, inquiries are rational if they are expectedly best given available information. Part 3 constructs three arguments for this account.

The first argument considers duties to gather evidence as a case study. I develop a consequentialist account which relates duties to gather evidence to the expected value of information gathered. I argue that this view outperforms competing accounts of duties to gather evidence developed by ethicists and epistemologists.

The second argument draws on Niko Kolodny's value question: why be rational? I argue that this question makes trouble for a number of traditional epistemological views. These views open a gap between the good and the rational, raising the challenge of why we should do what is rational instead of what is best. Consequentialism provides the natural answer: we should do what is rational because it is best.

The third argument returns to the skeptical challenge raised in Part 1. I argue that a consequentialist view is best suited for vindicatory epistemology, the project of recasting seemingly irrational behaviors as instances of rational inquiry or metacognition. I discuss three examples: conditional reasoning, anchoring and adjustment, and error monitoring.

The consequentialist account of rational inquiry raises a puzzle. How can bounded agents comply with consequentialist norms without engaging in expensive cognitive calculations? This is a question about rational metacognition: monitoring and control of inquiry. Part 4 compares a Bayesian account of rational metacognition with the two-process account favored by many metacognition researchers. I provide descriptive evidence in favor of the two-process account and use this account to explain how bounded agents manage to inquire well without undue cost.