

Why the grass is greener on the other side: Behavioral evidence for an ambiguity bonus in human exploratory decision-making

Robert C. Wilson, Andra Geana, John Myles White, Elliot A. Ludvig & Jonathan D. Cohen

When you go to a restaurant, do you always order the same thing or do you try something new? Going with an old favorite guarantees a happy meal, but you might miss out on something better unless you're willing to explore. The decision between choosing what you know and trying something different is called the Exploration-Exploitation problem.

Theoretical accounts suggest that explore-exploit decisions should be driven by both the expected value and the ambiguity associated with each option. When tractable, optimal strategies assign value to both of these factors, such that an ideal agent can choose an option with lower expected value when its ambiguity is high. Although the effects of expected value on human decision making are well understood, the effects of ambiguity have received relatively little attention and the optimal model's proscribed behavior seems counter to decades of research showing that humans are *averse* to ambiguity. We believe that this may reflect the fact that tests of ambiguity aversion usually involve a single choice, without the opportunity to reduce ambiguity by sampling. When given this opportunity, we predicted people would actually become ambiguity seeking.

We designed a task to test this prediction and tease apart the contributions of expected value and ambiguity in sequential decision making. In our "forced-play bandits task", subjects played a series of games lasting ten trials each. On each trial subjects selected one of two options which paid out points probabilistically. Subjects were instructed to maximize the number of points earned over the experiment and thus faced the explore-exploit dilemma afresh in each new game.

In the first four trials of each game participants were forced to select one option. By varying the number of times a given option was chosen during these forced plays, we were able to manipulate the level of ambiguity subjects faced when making their first free decision. Specifically, we used the forced play trials to set up two different ambiguity conditions: an "equal" condition, in which both options were played twice; and a "different" condition, in which one option was played once and the other three times.

We analyzed the data on the first free play trial by computing psychometric choice curves for each ambiguity condition as a function of the difference in means for the two options. In most subjects we found a significant bias in the "different" condition that was consistent with the ambiguity *seeking* behavior of the optimal model.

These results confirm our initial hypothesis and suggest that, rather than having a blanket aversion to ambiguity, humans modulate their attitudes to ambiguity according to the demands of the task.