24.223 Paper 1 Prompts

Kevin Dorst

Instructions: Please write a roughly 1500-word paper (5–6 pages double-spaced) that responds to one of the following prompts.

ChatGPT and AI: You can use whatever methods you like to brainstorm and outline, but once you start writing your draft, all your writing must be your own. **You must write your draft, from start to finish, in Google Docs.** When you submit your paper, you will **share editing privileges with me for your paper**, so that—if needed—I can use Draftback to verify that you wrote it yourself, rather than using an AI to write it.

The aim is to make one small contribution to the conversation started by the paper and in class. You should begin with a thesis paragraph, explaining what you will argue and how you will proceed. You should then reconstruct the argument you will be engaging with in a way that is understandable to someone who has not read the paper or come to class. (You may presuppose the basics of Bayesian epistemology and decision theory, though it is good to remind your readers what any terms and equations mean.) After doing so, you should make your contribution, which should come in the form of raising or responding to an objection to the argument you have reconstructed. I will grade your paper by evaluating it for (1) sentence-level and structural *clarity*, (2) accuracy and adequacy of its *exposition* of the target argument, and (3) the novelty and substance of its *engagement* with that argument.

Papers should be submitted over email to kmdorst@mit.edu by sharing editing privileges with your Google Doc. The due date is **11:59pm on November 6**. Late papers will be marked down 1/3 of a letter grade for each late day.

Prompts:

1. Hedden argues that hindsight bias is not necessarily a bias. What does this mean, and how does he argue for it? Is he right?

2. Past-Kevin argued that the gambler's fallacy is not a fallacy by showing that given a certain type of information, it's rational to be more than 50% confident that a koin will "switch". Reconstruct how this argument works, and then assess whether it captures what is meant, intuitively, by the "gambler's fallacy" in a range of real-world cases.

3. Using the "value of evidence" result from the Kinney and Bright paper, Kevin argued (in class) that motivated ignorance is *either* irrational *or* driven by non-morally-kosher values. Explain how this argument works, and whether you think it's right.

4. Kelly argues that—due to the possibility of "redemptive preferences"—either we have little evidence that people commit the sunk cost fallacy, or there's little reason to believe it's irrational. Is he right? Why or why not?