A Dialogue on a Classic Interview Puzzle

Salviati: I have a puzzle[1] for you!

Puzzle 1. In my hand I have two envelopes, each containing a real number. I will hand you an envelope, you will look inside, and learn the number. You may either keep that envelope, or swap it for the other envelope. If you swap, you may not swap back. Your goal is to end up with the envelope containing the larger number. What is your best strategy?

Simplicio's Answer and Salviati's Response

Simplicio: There is no best strategy. We don't know anything about the envelopes. Picking either at random will be fine.

Salviati: Oh not so, my simple minded friend. There is an ingenious strategy to get a better-than-chance outcome. You see, we may pick some reference number $r \in \mathbb{R}$. After choosing an enumeration of the envelopes, let x_1 and x_2 denote the numbers contained in envelopes one and two respectively. Since these two numbers are chosen randomly, there is some probability $P > 0$ that $x_1 \leq r \leq x_2$.

We may adopt the following strategy. Initially we learn x_1. If $x_1 \leq r$, we switch to the second envelope. Otherwise, we stay with the first envelope. In the event that r lies between x_1 and x_2, we are guaranteed a win by the strategy.

We easily see by an exhaustive enumeration of cases that this strategy wins with a probability of $\frac{1}{2} + P$. Since $P > 0$, this is a winning strategy.

Simplicio: How are x_1 and x_2 being selected? Are you assuming a uniform distribution on the real numbers? Are x_1 and x_2 being drawn uniformly from some bounded subset of the reals? This would make sense, since only so big a number can fit in an envelope.

Salviati: Good heavens no! This is a math problem. We are not constrained by the bounds of notation and A5 paper. We must not assume that x_1 and x_2 are drawn from a bounded subset of the reals. Although they might be, for we do not specify how x_1 and x_2 are generated.

Instead, consider this. Suppose that x_1 and x_2 are generated by some random process. Indeed, we may even suppose that x_1 and x_2 are generated by distinct independent random processes. Thinking of x_1 and x_2 as observations of random variables X_1 and X_2, for our chosen reference number r, we still have a probability $P := \mathbb{P}[x_1 \leq r \leq x_2] \geq 0$. Our reference number strategy leaves us no worse off than chance, but possibly better off if $P > 0$.

I suppose you are correct that it could be the case that $P = 0$, such as when $x_1 = x_2$, but that's just being pedantic. The genius of the strategy remains.

[1] If you haven't seen this puzzle, I encourage you to think about it before reading this dialogue.

Simplicio's First Objection

Simplicio: I plead the fifth on pedantry, but I will raise an objection. I know you are more sophisticated than I am, but I think you have become too comfortable in the Platonic realm of mathematics. You tell me that we may not apply the real world constraints of notation, but you assume that x_1 and x_2 are generated by some probabilistic processes.

Salviati: Well how else would they be chosen?

Simplicio: I don't know. Perhaps a Goddess wrote the numbers on your tongue before you put them in envelopes. Are you sure that Goddesses can be coherently modeled by probability distributions? In short, I think your strategy relies on the unstated assumption that the numbers can be modeled probabilistically.

The concept of **Knightian Uncertainty** may be useful here. In *Risk, Uncertainty, and Profit* [Kni21], Frank Knight sought to disentangle the concepts of "risk" and "uncertainty" in economics, finance, and decision theory. To Knight, **risk** is quantifiable by nature. Like statistical noise, it can be estimated and accounted for. **Uncertainty**, on the other hand is completely unquantifiable. The range of possible outcomes cannot be accounted for, and cannot even be described by a probability distribution. At its core, uncertainty reflects an essential unpredictability that lives beyond the scope of probabilistic tools.

Salviati: Must I accept the existence of Knightian uncertainty? While I concede that there are events and processes that I cannot describe probabilistically, it is merely a reflection of my own lack of knowledge. It is a matter of ignorance, not ontological indeterminacy. It is wholly consistent that I cannot provide a probabilistic model for Goddesses, and that Goddesses are perfectly described by some probability distribution.

Simplicio: Tread lightly. You are very close to arguing against your own free will. In his paper *The Ghost in the Quantum Turing Machine* [Aar16], Scott Aaronson argues that the existence of Knightian Uncertainty is logically entailed by free will. If there were a probability distribution that would perfectly describe all the choices you could ever make from birth until death, would you really have free will? Even if we lack the tools to probe at that distribution, it seems to me that its mere existence would imply a limit to your agency.

Salviati: I believe you have yet again conflated the real world with the domain of mathematics. I do not wish to stake out a strong position on the non-existence of free will, so I will concede that Knightian Uncertainty may-or-may-not exist in the universe we inhabit. However, this concession does not introduce Knightian uncertainty to the world of mathematics my puzzle inhabits.

Simplicio: I disagree. Your puzzle only works because it's a real-world hypothetical. Without the numbers being put in envelopes, the notion that they were "chosen" would be changed. Without this real-world framing, the puzzle reduces to the following:

Puzzle 2. Let x_1 and x_2 be real numbers. Which is bigger?

Simplicio: In this version of the puzzle, it would be outrageous to assume that x_1 and x_2 are observations of random variables. My answer is still that there is no strategy better than choosing randomly. Please notice, by the way, that it does not matter if my choice here can be described by a probability distribution.

Salviati: I am not convinced that removing all references to concrete objects reduces my puzzle to yours. You have, among other attributes, lost the decision-making aspect of the puzzle entirely. Yet in the interest of kindness and avoiding more economic terminology, I shall amend the statement of my puzzle[2] to appease you.

> *Puzzle* 3. In my hand I have two envelopes, each containing a real number, each generated by a separate independent random processes. I will hand you an envelope, you will look inside, and learn the number. You may either keep that envelope, or swap it for the other envelope. If you swap, you may not swap back. Your goal is to end up with the envelope containing the larger number. What is your best strategy?

Simplicio's Second Objection

Simplicio: This change of wording does nothing to address my objection.

Salviati: Oh dear… please no more economics.

Simplicio: Don't worry, I'm done with economic philosophy. Let's briefly review your solution.

x_1 and x_2 are observations of independent random variables X_1 and X_2 respectively. There is no declared relationship between X_1 and X_2 beyond independence. You set $P := \mathbb{P}[X_1 \leq r \leq X_2]$ for some reference number r, and argued that $P \geq 0$. Your strategy is to change envelopes if and only if the observation $x_1 \leq r$, and claimed that this strategy will succeed with probability $\frac{1}{2} + P$.

My objection is simple. I don't believe the probability of success is $\frac{1}{2} + P$. Why do you believe that?

Salviati: Well that should be trivial, even for you. For ease, let us assume that X_1 and X_2 are continuous random variables, but the argument can be adapted to handle the discrete and mixed cases. There are almost surely exactly six possible orders for our observations and the reference number. If $[x_1 < r < x_2]$, $[x_1 < x_2 < r]$, $[x_2 < r < x_1]$, or $[r < x_2 < x_1]$, the strategy results in a win. On the other hand, if $[x_2 < x_1 < r]$ or $[r < x_1 < x_2]$, the strategy results in a loss. Now we set:

$$
\begin{aligned}
P &:= \mathbb{P}[X_1 < r < X_2] = \mathbb{P}[X_2 < r < X_1] \\
Q &:= \mathbb{P}[r < X_1 < X_2] = \mathbb{P}[r < X_2 < X_1] \\
R &:= \mathbb{P}[X_1 < X_2 < r] = \mathbb{P}[X_2 < X_1 < r]
\end{aligned}
$$

[2] Salviati could have changed the puzzle by explicitly declaring the distributions from which x_1 and x_2 are picked. For example, maybe they're both uniformly chosen from the interval $[0, 1]$. This ruins the puzzle though! For the example given, it's immediately clear that the best strategy is to pick the second envelope if the number in the first envelope is $\leq \frac{1}{2}$. The puzzle isn't much of a puzzle anymore! This is why it is in Salviati's best interest to leave as much about the puzzle unspecified as possible.

From this, we see that $2(P + Q + R) = 1$, and we win with probability $2P + Q + R$. Basic algebra gives us that our probability of winning is $\frac{1}{2} + P$.

Simplicio: Why do you say that $\mathbb{P}[X_1 < r < X_2] = \mathbb{P}[X_2 < r < X_1]$ when this is clearly not true? Since X_1 and X_2 can be different distributions, these two probabilities could be completely different.

To be super explicit, your argument is implicitly assuming, among other things, that $\mathbb{P}[X_1 \leq X_2] = .5$, which is equivalent to assuming that X_1 and X_2 have the same median. This was not an assumption of the problem.

Salviati: Ah but you see, we don't know the distributions X_1 and X_2. It certainly could be that the medians are unequal, but *a priori* it is just as likely that $\text{median}(X_1) \geq \text{median}(X_2)$ as it is that $\text{median}(X_2) \geq \text{median}(X_1)$. Algebraically, these two cases cancel out and our win probability is $\frac{1}{2} + P$.

Simplicio: Then you've left my objection unaddressed. How did you pick the random variables X_1 and X_2? Did you sample them from a probability distribution of all possible probability distributions on $\mathbb{R}$? All you've done with this new puzzle is push the problem of Knightian Uncertainty up to the probability distributions themselves.

Salviati: I thought you said no more economics.

Simplicio: I used my free will to lie to you. Anyway, you must see how this game continues. If you assume you picked the distributions X_1 and X_2 from some distributions, I will point out that you need to have picked the distributions you drew *them* from. We'll push it up layer by layer until we're blue in the face. It's turtles all the way up.

Salviati: You think you're so clever, but there's an easy way out. Once we make two draws from the same distribution, we're safe. Perhaps I tried to be too general. Let me present one last version of the puzzle.

Puzzle 4. In my hand I have two envelopes, each containing a real number, each generated by independent draws of the same random process. I will hand you an envelope, you will look inside, and learn the number. You may either keep that envelope, or swap it for the other envelope. If you swap, you may not swap back. Your goal is to end up with the envelope containing the larger number. What is your best strategy?

Salviati: Now that X_1 and X_2 are independent identically distributed random variables, it is clear that $\mathbb{P}[X_1 < r < X_2] = \mathbb{P}[X_2 < r < X_1]$. With that, my solution works.

Simplicio's Final Objection

Simplicio: Usually I'd drop my objections here, but I can't resist. I still have an issue with your strategy. Let me ask you one more question. You ask for the *best* strategy, yet you leave the way we choose r unspecified. How should we pick r? After all, we want the best r for the given distri—

Salviati: Stop right there. I can already see the rest of this dialogue. Do you think I'm slow? You believe I'm going to say that the choice of r doesn't matter because we don't know anything about the distribution governing X_1 and X_2. Then you're going to argue that I'm effectively saying that we're picking the median value for the underlying distribution uniformly over all of $\mathbb{R}$, and I can't do that so I need to have specified the distribution from which the random variables X_1 and X_2 were drawn. Then I'll have to adapt the puzzle again.

But you've overextended yourself. Whether the uncertainty of the way X_1 and X_2 were picked is Knightian or not doesn't matter this time! The simple fact that the distribution *has a median* is enough. Sure we hope our choice of r is close to the actual median, but we have no way of knowing what it is. If you are going to assert that there's Knightian Uncertainty behind how I chose the distribution of X_1 and X_2, I'm just going to lean in further and pick r in a way that can't be probabilistically modeled as well. I'll just use *my* free will. Nothing goes wrong. Asshole.

Simplicio: Oh yeah I know that. That wasn't my objection. We're done with all this Knightian Uncertainty stuff. This time, it doesn't matter how you choose the distribution from which x_1 and x_2 are drawn, be it by free will or by a probabilistic process.

For an arbitrary distribution, why would you ever think that $\mathbb{P}[x_1 \leq r \leq x_2] > 0$? In light of the uncertainty of the distribution, for any choice of r, almost every distribution will put less than ϵ mass on the event that $x_1 \leq r \leq x_2$. There is, famously, no canonical null-set structure on the space of probability measures on $\mathbb{R}$, so we have some wiggle room here. I'm a simple man, but it seems to me that the only sensible assignment for P is 0. Your strategy is just my strategy with more steps.

Salviati: I'm never sharing a puzzle with you again.

References

[Aar16] Scott Aaronson. *The Ghost in the Quantum Turing Machine*, page 193–296. Cambridge University Press, 2016.

[Kni21] Frank H. Knight. *Risk, Uncertainty and Profit.* Houghton Mifflin Co, Boston, MA, 1921.