<u>Seminar 5 – Learning</u>

The question here generally relates to how we can use theory to design experiments.

Consider the following simple experimental test of rational informational herding theory:

"*A white urn contains x white ball(s) and y red ball(s), while a red urn contains x red ball(s) and y white ball(s)*, and subjects know the values of *x* and *y*. One of these urns is chosen at random with probability 0.5 and emptied into a green urn without the subjects being able to see which one was used. The subjects are asked to approach the green urn in sequence and draw a single ball from the urn, look at its colour and then replace it into the urn, without looking into the urn itself. This is done privately so only the subject and experimenter can see the colour of the ball drawn from the green urn. The subject is then asked to write down his best guess of which urn (red or white) was emptied into the green urn on a public whiteboard. The next subject then approaches the green urn and follows the same procedure but is also allowed to look at the whiteboard which contains the first subject's best guess. This process continues with the third subject now having access to both his private draw and a list of two best guesses on the whiteboard in order (so he can see which was made by which subject). This continues further with subject *t* in general having both his private draw and access to the whiteboard containing *t-1* best guesses from the previous subjects, all in order. Subjects are paid £10 if they guess correctly and £0 otherwise and know this payment scheme in advance."

(a) Consider the case where *x + y = 4.*

(i) Go through each possible combination of values of *x* and y and discuss the properties of the information gained when drawing a ball form the green urn under the different combinations of values of *x* and *y*.

(ii) What values of *x* and *y* would make most sense if we want to stick within the standard assumptions of rational herding theory? Why?

Next insert your chosen values of *x* and *y* from (a)(ii).

(b) Consider subject 6. Imagine first that he drew a red ball from the urn. Next consider the following possible sequences of publicly observable actions (written on the whiteboard). What is the optimal choice of subject 6 under the following observed sequences?

(i) red, white, white, white, white

(ii) red, white, red, white, red

(iii) white, red, white, red, white

(iv) red, red, red, red, white

(c) Imagine next that the following sequence has occurred: *red, red, red, red, red, red*

Now subject 7 has been asked to make a choice but has *not* been allowed to see the whiteboard and so has not seen this sequence. He draws a ball and then publicly chooses *white*. Subject 8 is shown the sequence of choices of the first 6 and subject 7, and he is told the circumstances surrounding the special situation faced by subject 7. He draws a white ball and then chooses red. Is this rational? If not, why not? If so, why so? If this choice is observed across numerous experiments with the same design does this reveal anything interesting?

Sketch of Answers

(a) (i)

If *x = 0* or *y = 0* then we have a perfectly revealing draw from the urn and this is at odds with the theory. (Notice that if *x = 0* then drawing a white ball perfectly reveals the urn to be red!).

If *x = y = 2* then the draw from the green urn contains no information at all and so the draws are worthless to the subject when trying to guess correctly. (The two urns are exactly the same, so draws are useless in providing information)

If *x = 1* and *y = 3* then information is informative and not perfectly revealing, but has to be inverted (so a white ball suggests the correct choice is red and vice versa).

White urn contains less white balls and vice versa. So, if the first subject draws a white ball from the green urn, more likely to be drawing from the red urn.

This can be seen via Bayes rule.

$$P(white\ urn|draw\ red) = \frac{P(Draw\ red\ and\ urn\ is\ white)}{P(draw\ red)}$$

$$= \frac{P(draw\ red|white\ urn)P(white\ urn)}{(P(draw\ red|white\ urn)P(white\ urn) + P(draw\ red|red\ urn)P(white\ urn)}$$

$$= \frac{\frac{3}{4}\times\frac{1}{2}}{\frac{3}{4}\times\frac{1}{2}+\frac{1}{4}\times\frac{1}{2}} = \frac{3}{4} > 1/2$$

If *x = 3* and *y = 1* then the information is informative and not perfectly revealing and a single draw from the green urn has the nice property that a red ball suggests the correct choice is red, and a white ball suggests the correct choice is white.

Same calculation as above.

(ii) Either $x = 3, y = 1$ or $y = 1, x = 3$.

Basically, we want to look at how people learn from information provided by earlier observations from others in a scenario where no one is never perfectly certain[1] sure about which answer is correct. Thus, we need a choice where information from others guesses actually has some (incomplete) information content at least in the beginning. Thus, either x=3 or x=1 will suffice. Let us choose x=3 for simplicity of exposition (this would also probably be easier for participants to understand).

Theoretically, $x = 1, y = 3$ and $x = 3, y = 1$ has no difference for cascade predictions. But given that people might be using heuristics, there might be differences in the results obtained: e.g. more likely to make the "correct" choice in the latter case due to correspondence between the correct colour choice and ball drawn. This might matter for the experimental design: perhaps $x = 3, y = 1$ will be less confusing to participants.

[1] Because of this, information on other's choices thus becomes important in influencing own choices.

(b) Given $x = 3, y = 1$, we have a result that information that sequences like RW and WR negative each other (before a cascade starts) and an information cascade occurs once more than two observations of the same colour are chosen in a row (e.g. __RR or __WW) Why?

Here, we focus on the case of red cascades, the reasoning of white cascades should be the same. We want to show that on observing RR, the 3rd subject will choose R no matter what he draws.

If the 2nd player observes that the 1st player guessed red (he knows rationally that 1st must have drawn red). From before, this means that the posterior probability (his new prior) of the urn being red has been updated to ¾.

More generally, we can perform calculations recursively by substituting the posterior probability of Subject 1 as the new prior of Subject 2.

$$P(red|draw\ X, observe\ H) = \frac{P(draw\ X|\ red)\pi(red)}{P(draw\ X|red)\pi(red) + P(draw\ X|white)\pi(White)}$$

Where $\pi(Z) = P(Z|observe\ H)$ and H is any set of histories which is displayed on the board.

Proof:

$$P(red\ |draw\ X, observe\ H) = \frac{P(red, draw\ X, observe\ H)}{P(observe\ H, draw\ X)}$$

Expanding the numerator:

$$P(red, draw\ X, observe\ H) = P(Draw\ X|observe\ H, red)^{2}P(observe\ H, red)$$

$$= P(draw\ X|red)P(red|observe\ H)P(observe\ H)$$

Expanding the denominator:

$$P(observe\ H, draw\ X) = P(red, observe\ H, draw\ X) + P(white, observe\ H, draw\ X)$$

$$= P(draw\ X|red)P(observe\ H, red) + P(draw\ X|white)P(observe\ H, white)$$

Combining the two, and using $\frac{P(observe\ H)}{P(observe\ H, red\ (or\ white))} = \frac{1}{P(red\ (or\ white)|observe\ H)}$, we get the above. ##

[2] Note that here, that conditioning on the urn is red, any observations will not change the probability of drawing X.

For the 2nd player who observes R from the 1st period:

If he draws a white ball,

$$P(red\ urn|draw\ white, observe\ R) = \frac{P(draw\ W|\ red)\pi(red)}{P(draw\ W|red)\pi(red) + P(draw\ W|white)\pi(White)}$$

$$\frac{\frac{1}{4}\times\frac{3}{4}}{\frac{1}{4}\times\frac{3}{4}+\frac{3}{4}\times\frac{1}{4}} = 1/2 \rightarrow indifferent$$

Let us assume that in this case, since he is indifferent, he chooses white with less than full probability, say ½.

If he draws a red ball,

$$P(red\ urn|draw\ red, observe\ R) = \frac{P(draw\ R|\ red)\pi(red)}{P(draw\ R|red)\pi(red) + P(draw\ R|white)\pi(White)}$$

$$\frac{\frac{3}{4}\times\frac{3}{4}}{\frac{3}{4}\times\frac{3}{4}+\frac{1}{4}\times\frac{1}{4}} = \frac{9}{10} \rightarrow choose\ red$$

Assuming players guesses at random when indifferent, then for the 3rd player, if he observes (R W), then it must be that actual sequence of draws. If instead he observes (R R), then there is some chance that the actual sequence of draws was actually (R W) due to indifference by the second player and randomisation into choosing R. This leads to some complication for the (RR) observation.

Following the notation from before, we want to calculate the posterior (new prior) of the 3rd subject given that he/she observes RR:

$$\pi(red) = P(red\ urn|observe\ RR) = \frac{P(red\ urn, observe\ RR)}{P(observe\ RR)}$$

$$= \frac{P(red\ urn, observe\ RR)}{P(red\ urn, observe\ RR) + P(white\ urn, observe\ RR)}$$

For the numerator we have:

$$P(red\ urn, observe\ RR) = P(red\ urn, observe\ RR\ actual\ RW) + P(red\ urn, observe\ RR, actual\ RR)$$

$$= P(observe\ RR|actual\ RW, red\ urn)^3 P(red\ urn, actual\ RW) + P(observe\ RR|actual\ RR, red\ urn)P(red\ urn, actual\ RR)$$

$$= P(observe\ RR|actual\ RW)P(red\ urn\ , actual\ RW) + P(observe\ RR|actual\ RR)P(red\ urn, actual\ RR)$$

$$= 0.5\ P(red\ urn, actual\ RW) + P(red\ urn, actual\ RR)$$

$$= 0.5\ P(actual\ RW|red)P(red) + P(actual\ RR|red)P(red)$$

$$= 0.5 \times (\frac{1}{4} \times \frac{3}{4}) \times \frac{1}{2} + (\frac{3}{4} \times \frac{3}{4}) \times \frac{1}{2} = \frac{3}{64} + \frac{9}{32} = \frac{21}{64}$$

Likewise, for the 2nd term in the denominator:

$$P(white\ urn, observe\ RR) = P(observe\ RR|actual\ RW)P(white\ urn, actual\ RW) + P(observe\ RR|actual\ RR)P(white\ urn, actual\ RR)$$

$$= 0.5\ P(white\ urn, actual\ RW) + P(white\ urn, actual\ RR)$$

$$= 0.5 \times P(actual\ RW|white) \times P(white) + P(actual\ RR|white) \times P(white)$$

$$= 0.5 \times \frac{1}{4} \times \frac{3}{4} \times \frac{1}{2} + \frac{1}{4} \times \frac{1}{4} \times \frac{1}{2}$$

$$= \frac{3}{64} + \frac{1}{32} = \frac{5}{64}$$

So,

$$\pi(red) = \frac{\frac{21}{64}}{\frac{21}{64} + \frac{5}{64}} = \frac{21}{26} > \frac{3}{4}$$

We have thus derived player 3's posterior when observing RR (our new prior), let us look at the case where the 3rd person observes RR and draws W (in the R case, it is obvious that he will choose R)

$$P(red|\ observe\ RR, draw\ W) = \frac{P(draw\ W|\ red)\pi(red)}{P(draw\ W|red)\pi(red) + P(draw\ W|W)\pi(White)}$$

[3] Here, our observations only depend on the actual draws because subject 2 does not have knowledge of the urn and his choice (which determines what we observe) only depends on the actual draws.

Notice that as long as the prior probability of red is strictly greater 3/4, then will choose red[4]. This is true here. When observing RR, the second R is like a weak observation (due to uncertainty), so while the calculated posterior (current prior) should be greater (>3/4), it does not increase to as much as when one actually observes the draws RR (9/10).

Given that Subject 3's choice is independent of his own draw, it does not contribute to an update to posteriors for Subject 4. This means that Subject 4 will always choose red as well.

(i) white. Cascade starts from 5th person, so should choose white.

(ii) red. No cascade yet. First 4 are redundant. So, will choose red.

(iii) indifferent – could choose either/ randomize.

(iv) It looks like subject 5 has made an error. Subject 6 should stick with red in any case even if subject 5 did draw a white ball: All information is useless except first two and maybe white. But given that he draws red, still rational for him to choose red.

These first two parts are about how rational individuals can exhibit herding behaviour via standard Bayesian updating. They serve as a benchmark for what we predict we will observe.

From the Anderson and Holt paper, results seem somewhat *consistent* (with an addition of perceptions of errors in others' decisions) with rational herding theory.

As an aside, imagine we observe exactly what we predicted, i.e. that after 2 consecutive similar observations, we observe cascades. This does not necessarily mean that the mechanism: Bayesian updating is definitely correct. This behaviour could also be consistent with social pressure/ desires to conform, simple heuristics (majority rule?) etc. The above experimental design without other modifications cannot distinguish between these theories.

For example, to rule out a pure majority rule-based decision, could possibly have treatments with different theoretical thresholds for indifference: 2R 1W, any difference; currently is 1R 1W which resembles majority rule. Another example is given in part c) below.

---

[4] Remember from the earlier calculation (2nd player) that when his prior was ¾ and he drew white, he was indifferent. Hence, when the prior is >3/4, he should strictly prefer red.

(c) If rational, subject 7, having not seen the public information on the whiteboard would have been best off to just follow his draw, and so subject 8 can infer from subject 7's choice that he saw a white ball. Subjects 3, 4, 5 and 6 were in a herd and so nothing can be inferred from their choices. Hence subject 8 has the following information which can be inferred from the sequence and his own private draw from the green urn:

• 2 white draws (subjects 7 and 8).

• 1 red draw (subject 1).

• Subject 2 may have seen white (and so chose red on indifference) or red (and so chose red).

If we know exactly 2R 2W were drawn, should be indifferent. Since here, we assume that Subject 2 may have randomised, we have less that 2R. So, should strictly choose white.

His choice of red was therefore *not* rational. If this sort of non-rational choice was repeated across many iterations of this experiment it would suggest that subjects do not realise how fragile a herd is: they might be incorrectly incorporating the decisions of subjects 3, 4, 5 and 6 into their choice even though these decisions are not based on any private information.

The key aim of the design here is to examine how well subjects recognise the non-informativeness of choices in earlier rounds. Perfect rationality predicts that in such a case, the cascade should be broken. The reason for this as explained in one of the readings is that information cascades mean that the set of information accumulates at the border, therefore a relatively small amount of publicly new information can tilt the situation. Thus, these cascades could be fragile; they use this theory to explain fads.

Note that fragility here is implemented through Subject 7's draw being perfectly informative, while the first 6 are only partly informative (With only first 2 draws being almost fully informative). In particular, we have that majority of information about choices are R, but a perfect revelation of Subject 7's subsequent draw leading to a totally different prediction.

This would provide a test of whether people are indeed perfectly rationally "imitating" others, or irrationally imitating others (in the sense of having some kinds of heurestics / other preferences).

There are some experiments[5] on this and seem to indicate less fragility than predicted by this sort of rational updating procedure: in particular, more guesses of R observed leads to a higher likelihood of choosing R.

The gist of the story here is that we need to be aware of how other theories may be consistent with the predictions of our own theory. In order to distinguish better between different theories, variations of an experiment which control for different avenues can be conducted. Of course, these theories might not always be thought of concurrently: future research thus helps to clear any possible misconceptions from earlier incomplete/incorrect theories.

It also gives us a good example of the links between theory and experimental methodology.

[5] "Fragility of information cascades: an experimental study using elicited beliefs" has a summary of research in this.