ST222 2017 GAMES, DECISIONS AND BEHAVIOUR
EXERCISE SHEET 5 – SOLUTIONS

(1) **Runs in random sequences.** In the lecture, we calculated the expected number of runs in sequences of fair coin tosses. This exercises asks the same question, but without assuming the coin is fair.

Model a sequence of n fair coin tosses by a sequence of independent random variables X_i ($i = 1, \ldots, n$) with values in $\{0, 1\}$ and with $P(X_i = 1) = p$ and $P(X_i = 0) = 1 - p = q$ for all $i = 1, \ldots, n$.

Let Z_r be the number of runs of heads or of tails of length r in n tosses of the coin. Calculate the expectation of Z_r.

First, calculate the expected number of runs of $1's$ using independence of the coin tosses:

$$\begin{aligned} E\big[Z_r^{(1)}\big] &= E\Big[1_{\{X_1=X_2=\ldots=X_r=1, X_{r+1}=0\}} \\ &\quad + 1_{\{X_{n-r}=0, X_{n-r+1}=X_{n-r+2}=\ldots=X_n=1\}} \\ &\quad + \sum_{i=2}^{n-r} 1_{\{X_i=X_{i+1}=\ldots=X_{i+r-1}=1, X_{i-1}=X_{i+r}=0\}}\Big] \\ &= P(X_1 = X_2 = \ldots = X_r = 1, X_{r+1} = 0) \\ &\quad + P(X_{n-r} = 0, X_{n-r+1} = X_{n-r+2} = \ldots = X_n = 1) \\ &\quad + \sum_{i=2}^{n-r} P(X_i = X_{i+1} = \ldots = X_{i+r-1} = 1, X_{i-1} = X_{i+r} = 0) \\ &= p^r q + p^r q + (n - r - 2 + 1)p^r q^2 \\ &= 2p^r q + (n - r - 1)p^r q^2 \end{aligned}$$

To obtain the expected number of runs of $0's$ we simple need to swap the role of p and q in the above calculation:

$$E\big[Z_r^{(0)}\big] = 2pq^r + (n - r - 1)p^2 q^r$$

$$\begin{aligned} E\big[Z_r\big] &= E\big[Z_r^{(1)} + Z_r^{(0)}\big] = E\big[Z_r^{(1)}\big] + E\big[Z_r^{(0)}\big] \\ &= 2pq^r + (n - r - 1)p^2 q^r + 2p^r q + (n - r - 1)p^r q^2 \\ &= 2(pq^r + p^r q) + (n - r - 1)(p^2 q^r + p^r q^2) \\ &= 2pq(q^{r-1} + p^{r-1}) + (n - r - 1)p^2 q^2 (q^{r-2} + p^{r-2}) \end{aligned}$$

(2) **Terrorist identification** (somewhat based on Wikipedia entry about *base rate fallacy.*

In a city with N inhabitants let there be k terrorists and $N-k$ non-terrorists. For simplification it is assumed that all people present in the city are inhabitants. Thus, the base rate probability of a randomly selected inhabitant of the city being a terrorist is $\pi = k/N$.

In an attempt to catch the terrorists, the city installs an alarm system with a surveillance camera and automatic facial recognition software.

Let p be the probability that the camera sends an alarm when it scans a terrorist (*sensitivity*). Let q be the probability that the camera does not send an alarm when it scans a non-terrorist (*specificity*).

The software can fail in two ways:

- *False negative:* the system does not send an alarm despite the individual being scanned is a terrorist
- *False positive:* the system sends an alarm, despite the individual being scanned is not a terrorist

The false negative rate α and the false positive rate β are the probabilities associated with these events.

Answer the following questions in the context of this alarm system in this city.

Let $+$ be the event that the individual's scan is positive.
Let $-$ be the event that it is negative.
Let T be the event that the scanned individual is a terrorist.
With this notation,

$$p = P(+\,|\,T) \qquad \text{and} \qquad q = P(-\,|\,T^c)$$
$$\alpha = P(T\,|\,-) \qquad \text{and} \qquad \beta = P(T^c\,|\,+)$$

We will also use

$$P(+) = P(+\,|\,T)P(T) + P(+\,|\,T^c)P(T^c) = p\pi + (1-q)(1-\pi)$$
$$P(-) = P(-\,|\,T)P(T) + P(-\,|\,T^c)P(T^c) = (1-p)\pi + q(1-\pi)$$

(a) The company that sells the software claims says that $p = 0.95$ and $q = 0.99$ and hence the false negative rate is 5% and the false positive rate is 1%. Is this true? Can you say which are the correct percentages? Discuss how the company may have derived these percentages.

No, it is wrong. It is impossible to calculate concrete percentages without knowing the base rate π. In order to be able to come up with some percentages anyway, the company may have confused α with

$$P(-\,|\,T) = 1 - p = 1 - 0.95 = 5\%.$$

And they may have confused β with

$$P(+\,|\,T^c) = 1 - p = 1 - 0.99 = 1\%.$$

(b) Derive formulas for α and for β.

Using the formulas above,

$$\alpha = P(T\,|\,-) = \frac{P(T)}{P(-)}P(-\,|\,T) = \frac{(1-p)\pi}{(1-p)\pi + q(1-\pi)}$$
$$\beta = P(T^c\,|\,+) = \frac{P(T^c)}{P(+)}P(+\,|\,T^c) = \frac{(1-q)(1-\pi)}{p\pi + (1-q)(1-\pi)}$$

(c) What are α and β for $k = 0$? First explain what the answer should be intuitively and then calculate it.

α is the probability the individual is a terrorist conditioned on another event (negative scan). But since there are no terrorists, it must be 0. β is the probability the individual is not a terrorist conditioned on another event (positive scan). Since nobody is a terrorist, this must be 1.

Using the formulas above confirms this.

(d) What are α and β for $k = N$? First explain what the answer should be intuitively and then calculate it.

α is the probability the individual is a terrorist conditioned on another event (negative scan). But since everybody is a terrorists, it must be 1. β is the probability the individual is not a terrorist conditioned on another event (positive scan). Since everybody is a terrorist, this must be 0.

Using the formulas above confirms this.

(e) Assume the sensitivity p is 95% and the specificity q is 0.99 and that the city in question has 100 terrorists among 1 million inhabitants. Calculate α and β. Comment on the numbers you obtained.

$\pi = 100/1,000,000 = 0.0001.$

$$\begin{aligned}
\alpha &= \frac{(1-p)\pi}{(1-p)\pi + q(1-\pi)} \\
&= \frac{0.05 \cdot 0.0001}{0.05 \cdot 0.0001 + 0.99 \cdot 0.9999} = \frac{500}{500 + 9900 \cdot 9999} \\
&\approx 0.000005050985 \\
\beta &= \frac{(1-q)(1-\pi)}{p\pi + (1-q)(1-\pi)} \\
&= \frac{0.01 \cdot 0.9999}{0.95 \cdot 0.0001 + 0.01 \cdot 0.9999} = \frac{100 \cdot 9999}{9500 + 100 \cdot 9999} \\
&\approx 0.9905885
\end{aligned}$$

That means, the probability someone is a terrorist despite having a negative scan is extremely small (lower than 0.00051%). On the hand, the probability someone with a positive scan was actually misclassified (as terrorist) is extremely high (larger than 99%).

(f) (*Optional*) If you have a computer with R (or another suitable other programme), write some code to do the previous exercise for a range of values of k.

$p = 0.95, q = 0.99, N = 1000000$

alpha(p,q,1/N)= 0.00000001052633

alpha(p,q,10/N)= 0.0000001052642

alpha(p,q,100/N)= 0.000001052736

alpha(p,q,1000/N)= 0.00001053674

alpha(p,q,10000/N)= 0.0001063151

beta(p,q,1/N)= 0.9999802

beta(p,q,10/N)= 0.999802

beta(p,q,100/N)= 0.9980237

beta(p,q,1000/N)= 0.9805654

beta(p,q,10000/N)= 0.8333333

Even though the sensitivity is not perfect (95%), the false negative rate for the range of prevalences considered (1 in 1,000,000 to 10,000 in 1,000,000) is very low. However, even with the high specificity assumed here (99%) the false positive rate is very high across the whole range of prevalences considered, including quite high ones.

(g) Give an lower bound for q as a function of p and π such that β is at most 1%. Calculate the threshold explicitly for $p = 0.95, N = 1,000,000$ and $k = 100$.

$$\frac{(1-q)(1-\pi)}{p\pi+(1-q)(1-\pi)} \leq 0.01$$
$$\Leftrightarrow (1-q)(1-\pi) \leq 0.01p\pi - 0.01(1-q)(1-\pi)$$
$$\Leftrightarrow 1.01(1-q)(1-\pi) \leq 0.01p\pi$$
$$\Leftrightarrow -1.01q(1-\pi) \leq 0.01p\pi - 1.01(1-\pi)$$
$$\Leftrightarrow -q \leq \frac{0.01p\pi}{1.01(1-\pi)} - 1$$
$$\Leftrightarrow q \geq 1 - \frac{p\pi}{101(1-\pi)}$$

Explicitly for $p = 0.99$, $N = 1,000,000$ and $k = 100$, $\pi = k/N = 0.0001$ and hence

$$q = 1 - \frac{0.99 \cdot 0.0001}{101 \cdot (1-0.0001)} = 0.999999$$

(h) (*Optional*) If you have a computer with R (or another suitable other programme), write some code to draw the false negative rate β as functions of the specificity q with fixed sensitivity $p = 0.95$ and fixed prevalence $\pi = 0.0001$. Show a few such curves for different prevalences between $\pi = 0.01$ and $\pi = 0.000001$. Interpret the results.

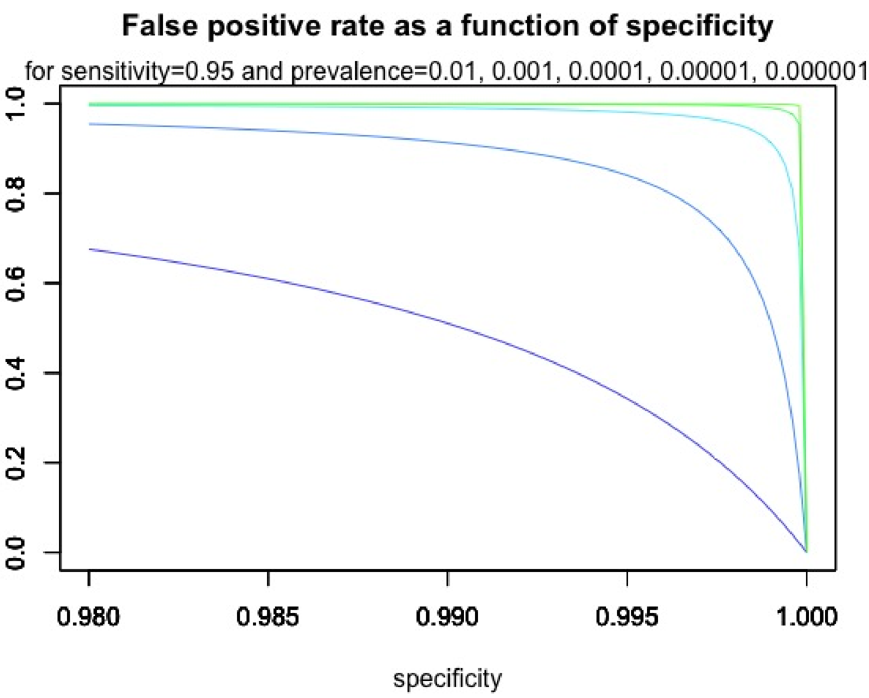

The plot shows that there is a non-linear dependency between the false positive rate β and the specificity q. For low prevalences - and this would be the case in this example - the specificity needs to be extremely high to achieve lower false positive rates.

(i) (*Optional*) If you have a computer with R (or another suitable other programme), write some code to draw the false positive rate β as functions of the prevalence for fixed sensitivity $p = 0.95$ fan fixed specificity $q = 0.999$. Show a few such curves for different specificities between $q = 0.9$ and $q = 0.99999$. Interpret the results.

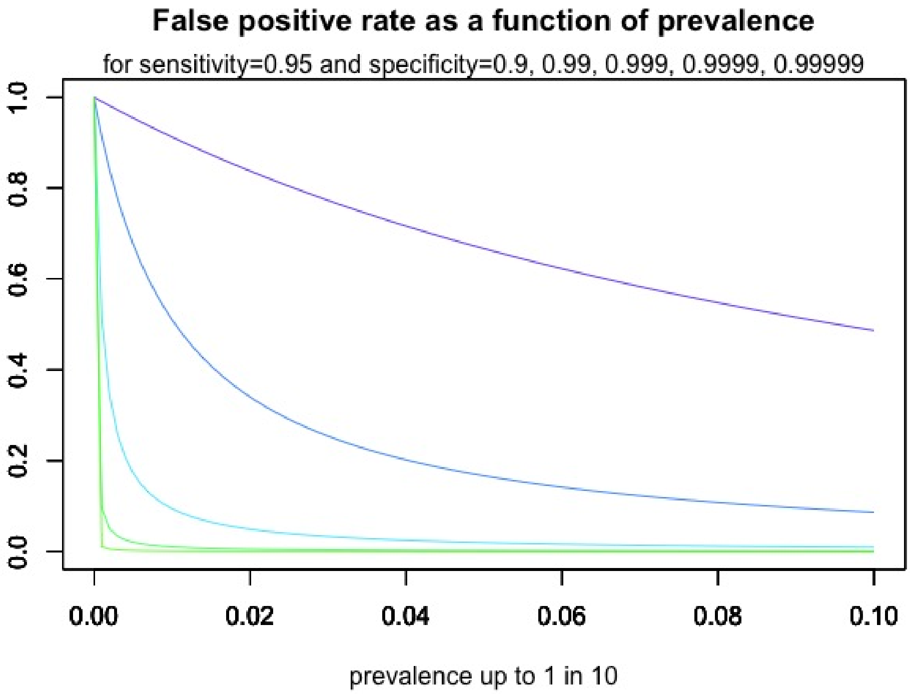

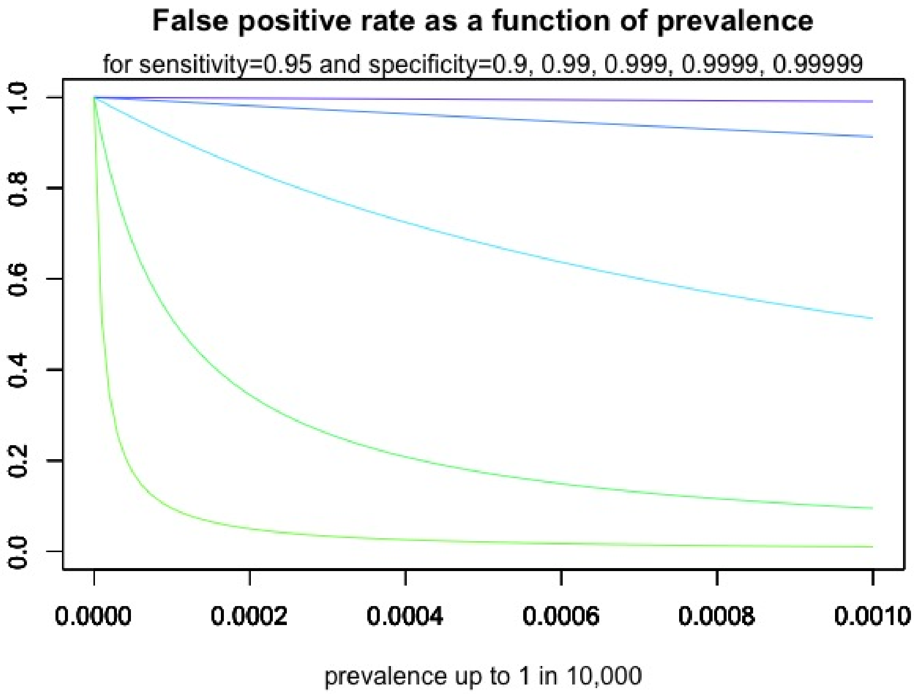

The plots show that there is a non-linear dependency between the false positive rate β and the prevalence q. For low prevalences - and this would be the case in this example - only the curve belonging to the specificity of 0.99999 can achieve false positive rates under about 10%, and once there are fewer than 100 terrorists per 1,000,000 inhabitants, it quickly increases as wel.

(3) **Empirical study on affect heuristic.** You want to empirically study the *affect heuristic* in the context of how people perceive and evaluate risk. Here is some scientific background:

> "Humans perceive and act on risk in two fundamental ways. Risk as feelings refers to individuals' instinctive and intuitive reactions to danger. Risk as analysis brings logic, reason, and scientific deliberation to bear on risk management. Reliance on risk as feelings is described as the affect heuristic."

Source: Paul Slovic and Ellen Peters, Risk Perception and Affect, Current Directions in Psychological Science, Vol. 15, No. 6, pp. 322-325, 2006.

You have the opportunity to collect data in 8 classes of young people between 13 and 15 at your old secondary school. You have 30 minutes of class time, a set-up for showing videos and tablet computers for each students are available.

Design an experiment to compare how students perception and evaluation of risk depends on their mood. Your experiment should start with an exposure to risk that is supposed to impact the subjects feelings one way or another, followed by assessing their perception and evaluation of risk through questions or tasks.

Explain the set-up of the experiment and give reasons for your choices. Acknowledge potential difficulties and limitations of the design (e.g. logistical challenges, biases, variation between subjects, questions regarding conclusions for a general population etc).

There are many possible answer to this question. The following gives just one example.

The 8 classes are being divided into two groups. First, they are being shown a 15 minute video. Group A is shown a video that is expected to provoke anxiety in the students (e.g. devastating epidemic, violence in their town). Group B is shown a video about a more neutral topic (e.g. documentary about an ant colony or the manufacturing of furniture). After that, students are being given a questionnaire that assesses their answers to a variety of tasks measuring risk perception, risk attitude and decision making under uncertainty. This includes lottery experiments that quantify how much someone is willing to pay for reducing uncertainty of different levels. For example, assessing preference between

A: Get £45 with certainty.

B: Toss a coin. If it lands heads get £100, otherwise get nothing.

The experiment divides the students into a group that will experience an intervention (anxiety provoking video) and a group where the intervention is replaced by a neutral activity in order to compare their responses and attribute observed differences to the intervention. For (logistical) simplicity, students will undergo the experiment in their form (class) with some forms being part of the intervention and others part to the neutral activity.

Discussion: Logistic challenges include timing, sufficiently fast distribution of questionnaires, getting students' attention. Another source of bias is any subculture related to their form. Generally, the social situation may interfere with their answers, as at that age students tend to pay a lot of attention to how their peers think about them (eager to please peers, fear of embarrassment). Another question is whether or not videos are suitable to trigger real emotions. From the scientific point of view it would be better to put them in a truly frightening situation (e.g. through wrong information about their grades), but this not pass any ethics panel. Caution has to be applied with generalising these findings to the whole population, because the study sample is a very specific in age and they all go to the same school which may have a subculture.

(4) **Properties of the probability weighting function.** In response to empirical findings dismissing *homo economics* as a model organism for human behaviour, Kahneman and Tversky have proposed a revised version of (subjective) expected utility theory. Prospect theory works with similar ingredients, but allows more flexibility: probabilities are weighted by a function,

utilities can be measured with respect to a reference point, the utility function is asymmetric with respect to the reference point.

(a) K & T defined the probability weighting function

$$w(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{1/\gamma}} \qquad (p \in [0,1])$$

for parameters $\gamma > 0$.

(i) Show that w has the *subcertainty property,* that is,

$$w(p) + w(1-p) < 1 \text{ for all } p \in (0,1).$$

Let $g(p) = w(p) + w(1-p)$. Then the conditions is equivalent to

$$\log(g(p)) < 0 \text{ for all } p \in (0,1).$$

Plugging in, because of symmetry,

$$g(p) = \frac{p^\gamma + (1-p)^\gamma}{(p^\gamma + (1-p)^\gamma)^{1/\gamma}} = (p^\gamma + (1-p)^\gamma)^{\gamma - 1/\gamma},$$

hence

$$\log(g(p)) = (\gamma - 1/\gamma)\log(p^\gamma + (1-p)^\gamma).$$

For this to be smaller than 0 consider all cases. For $\gamma < 1$ the first factor is negative and the second is positive. For $\gamma > 1$ it is the other way round. Either way, their product is negative. (For $\gamma = 1$ both factors are 0, and so is their product, so we can not show the required strict inequality but just equality. But this case is not interesting, since $w(p)$ is just the identity, and it should have been excluded in the formulation of the exercise.)

(ii) For commonly used values of γ, w is monotone. However, this is not the case for small γ. Illustrate this with plots. Using analysis, determine approximately, for which γ the function w is monotone.

Let $F(p) = \log w(p)$. Since $\log$ is strictly increasing, w is monotone if and only if F is.

$$F'(p) = \frac{\gamma p^\gamma + \gamma(1-p)^\gamma - p^\gamma + p(1-p)^{\gamma-1}}{p(p^\gamma + (1-p)^\gamma}$$

Since the denominator is positive, for F' is negative if and only if the numerator is negative. Dividing it by $(1-p)^\gamma$ and defining $x = p/(1-p)$ and $f(x) = (\gamma-1)x^\gamma + x + \gamma$, we see that F' is negative if and only if f is negative. $f(0) = \gamma > 0$, $f(x)$ tends to $+\infty$ for x going to ∞. Calculating the second derivate shows that f is strictly convex:

$$f'(x) = \gamma(\gamma-1)x^{\gamma-1} + 1, f'' = \gamma(\gamma-1)^2 x^{\gamma-2} > 0$$

Hence, f has a unique minimum in $x_0 = (\gamma(1-\gamma))^{1/(1-\gamma)}$ and negative values can only be attained in a neighbourhood of x_0.

$$f(x_0) = \gamma(1 - (1-\gamma)^{(2-\gamma)/(1-\gamma)}\gamma^{(2\gamma-1)/(1-\gamma)}$$

We now need to look at this as a function of γ. There is a critical value γ_0 such that for all $\gamma < \gamma_0$, $f(x_0)$ is negative and, hence, w is not monotone. Finding the critical value means to find γ such that $f(x_0) = 0$, in other words, $\gamma^{(2-\gamma)/(1-\gamma)}\gamma^{(2\gamma-1)/(1-\gamma)} = 1$. By inspection of the graph, or numerical analysis, we find that γ_0 is approximately 0.279. For plots of w for a range of γ see lecture notes.

Note: This exercise is harder than what is expected in the exam.

(b) **Medical treatment choice with additional option.** In a survey (Redelmeier DA, Shafir E. Medical Decision Making in Situations That Offer Multiple Alternatives. JAMA 1995) of a random sample of 373 family physicians in Ontario, two almost identical patient descriptions were used in a decision task. The descriptions are printed below with differences marked in bold. The physicians were randomly assigned to one of the two groups.

> **Group A:**
> The patient is a 67-year-old farmer with chronic right hip pain. The diagnosis is osteoarthritis. You have tried several nonsteroidal anti-inflammatory agents (eg, aspirin, naproxen, and ketoprofen) and have stopped them because of either adverse effects or lack of efficacy. You decide to refer him to an orthopedic consultant for consideration for hip replacement surgery. The patient agrees to this plan.
> Before sending him away, however, you check the drug formulary and find that there is one nonsteroidal medication that this patient has not tried (**ibuprofen**). What do you do?
>
> **Group B:**
> The patient is a 67-year-old farmer with chronic right hip pain. The diagnosis is osteoarthritis. You have tried several nonsteroidal anti-inflammatory agents (eg, aspirin, naproxen, and ketoprofen) and have stopped them because of either adverse effects or lack of efficacy. You decide to refer him to an orthopedic consultant for consideration for hip replacement surgery. The patient agrees to this plan.
> Before sending him away, however, you check the drug formulary and find that there are two nonsteroidal medications that this patient has not tried (**ibuprofen and piroxicam**).
> What do you do?

In Group A, 53% chose the default option of proceeding to surgery, and in Group B it were 72%. (Running standard statistical tests shows a statistically significant difference. The response rate in this study was 77%, so the results are meaningful.)
Give a possible explanation for the difference. Use the terminology introduced in lecture.
Hint: Check in the lecture slides for experiments creating a similar decision situation and discuss the physicians in this study may behave along lines.

The results are surprising in one might have thought that more choices of alternative drugs should have resulted in a higher percentage of physicians opting for a drug, but instead it was substantially lower. This can be explained by avoiding decisional conflict. If there is an additional decision to take, it ends up being avoided all together. This may seem illogical when described as a isolated behaviour, but becomes less irrational when put into the context of a practical situation where factors like time pressure or incomplete information can drive the behaviour.

(c) **Choices involving different types of risk.** Consider the following lab experiment. Fifty-six undergraduates were given the following question:
You have two lotteries to win \$250. One offers a 5% chance to win the prize and the other offers a 30% chance to win the prize.
A: You can improve the chances of winning the first lottery from 5 to 10%.
B: You can improve the chances of winning the second lottery from 30 to 35%.
Which of these two improvements, or increases, seems like a more significant change?

The majority of respondents (75%) viewed option A as the more significant improvement. The same respondents were also given a different question.

You have two lotteries to win $250. One offers a 65% chance to win the prize and the other offers a 90% chance to win the prize.

C: You can improve the chances of winning the first lottery from 65 to 70%.

D: You can improve the chances of winning the second lottery from 90 to 95%.

Which of these two improvements, or increases, seems like a more significant change?

In the second question, only 37% of the participants viewed option C as a more significant improvement.

(i) Give reasons why respondents chose option A over B.

An increase from 5 to 10 appears bigger than an increase from 30 to 35. This can be explained by a ratio based approach. From 5 to 10 means to double the number, whereas the increase from 30 to 35 only amounts to an increase by a sixth.

(ii) Give reasons why respondents chose option D over C.

People like certainty, and 95% gets them very close to that.

(iii) Explain how this apparent incoherence can be explained.

By the rationale leading to people choosing A over B they would have to choose C over D. However, their behaviour simply follows another, more dominant, heuristics that applies in this case, as detailed above. These priorities are capture in the S-shape of typical probability weighting functions.