# CCS Problem Set 3
## Due: 30 May @ 11.55pm

Your writeups should include text and figures. Many problems can be solved either analytically or with code: unless specified otherwise, you may do whichever one you wish (and need not do both!). If you code something up, you should always include the (commented) code in your write-up in addition to reporting the solution. We have written the problem set assuming that you are coding in R; if you would like to use a different language, please contact Wai Keen first. Writeups should be submitted as pdfs; LaTeX is preferable but not obligatory.

1. **Prospect theory** (25 points)

   In the first exercise in this problem set, we look to apply prospect theory to solving a "simple" problem in politics.

   **(1a)** Much of politics – and much of life – involves making trade-offs between competing interests and competing desires. In this question, we want to look at what prospect theory has to say about these trade-offs. However, in order to do so, we need to make some extensions to prospect theory. In particular, we need to consider value functions that are defined over two types of event. That is, every possible outcome is a vector $\mathbf{z} = (z_1, z_2)$ (in the example we're going to consider below, these two types of events will relate to "number of kittens received" and "number of puppies received"). Similarly, our reference point $\mathbf{r} = (r_1, r_2)$ now corresponds to a location in a two-dimensional space. So, if we let $Z$ denote the set of possible outcomes $\mathbf{z}$, the value of some action $a$ to a decision-maker who adopts the reference point $\mathbf{r}$ is now written as follows:

$$v(a, \mathbf{r}) = \sum_{\mathbf{z} \in Z} P(\mathbf{z}|a) v(\mathbf{z}, \mathbf{r})$$

   where $v(\mathbf{z}, \mathbf{r})$ is the value of outcome $\mathbf{z}$ relative to $\mathbf{r}$. If we assume that all combinations of $z_1$ and $z_2$ are possible (formally: that $Z = Z_1 \otimes Z_2$), we can rewrite this as:

$$v(a, \mathbf{r}) = \sum_{z_1 \in Z_1} \sum_{z_2 \in Z_2} P(z_1, z_2|a) v(z_1, z_2, r_1, r_2)$$

   The nested summations in this expression are annoying, so it would be nice to make them go away. In general, we can't do this, but if we make some additional assumptions it is possible. Specifically, let's assume that the outcomes $z_1$ and $z_2$ are conditionally independent of one another given the action $a$ chosen by the decision maker:

$$P(z_1, z_2|a) = P(z_1|a) P(z_2|a).$$

   Moreover, we will assume that the value function is additive, in the following sense:

$$v(z_1, z_2, r_1, r_2) = v_1(z_1, r_1) + v_2(z_2, r_2)$$

   for some individual value functions $v_1(\cdot, \cdot)$ and $v_2(\cdot, \cdot)$. Your first task is to show that, under these assumptions, prospect theory allows us to simplify the calculation of the total value of $a$ to the sum of it values with respect to the two outcomes. That is, show that:

$$v(a, \mathbf{r}) = \sum_{z_1 \in Z_1} P(z_1|a) v_1(z_1, r_1) + \sum_{z_2 \in Z_2} P(z_2|a) v_2(z_2, r_2)$$

Analytic solutions preferred [Hint: This problem is easier than it looks. It's just algebra], but numerical work is acceptable.

**(1b)** Okay, time to start thinking about politics. The town of Cutesville is having an election. The people of Cutesville like both kittens and puppies, but they have a fixed budget of \$100 to pay for both. You are considering running for office in Cutesville, and you need to figure out what your "Kittens and Puppies Budgetary Policy" is going to be. Specifically, you must determine how many dollars $d_k$ to spend on kittens and how many dollars $d_p$ to spend on puppies. Let $\mathbf{d} = (d_k, d_p)$ denote your spending policy, and note that $d_k + d_p = 100$, $d_k \geq 0$, and $d_p \geq 0$. Next, let $\mathbf{n} = (n_k, n_p)$ denote the outcome of your policy, where $n_k$ denotes the number of kittens your policy generates, and $n_p$ denotes the number of puppies it generates. Careful testing has determined that the kitten-generation and puppy-generation mechanisms are independent, and that voters evaluate their kitten-prospects separately from their puppy-prospects. As a consequence, you may assume the result from (1a): the value of your spending policy $\mathbf{d}$ is

$$v(\mathbf{d}, \mathbf{r}) = \sum_{n_k=0}^{N_k} P(n_k|d_k)v(n_k, r_k) + \sum_{n_p=0}^{N_p} P(n_p|d_p)v(n_p, r_p)$$

where $N_k$ is the maximum number of kittens, and $N_p$ is the maximum number of puppies. Scientific testing has shown that for your budget, these maxima are $N_k = 10$ and $N_p = 20$. Moreover, if you spend $d_k$ dollars on kittens, then the probability that your policy generates $n_k$ kittens is binomial:

$$P(n_k|d_k) = \begin{cases} \frac{10!}{n_k!(10-n_k)!}\left(\frac{d_k}{150}\right)^{n_k}\left(1-\frac{d_k}{150}\right)^{10-n_k} & \text{if } 0 \leq n_k \leq 10 \\ 0 & \text{otherwise} \end{cases}$$

If you spend $d_p$ dollars on puppies, the probability of generating $n_p$ puppies is:

$$P(n_p|d_p) = \begin{cases} \frac{20!}{n_p!(20-n_p)!}\left(\frac{d_p}{300}\right)^{n_p}\left(1-\frac{d_p}{300}\right)^{20-n_p} & \text{if } 0 \leq n_p \leq 20 \\ 0 & \text{otherwise} \end{cases}$$

[Note that the `choose()` function in R provides an efficient way to calculate $n!/(k!(n-k)!)$.] Polling has determined that the voters are expecting to receive a total of $r_k = 4$ kittens and $r_p = 3$ puppies (so this is their reference point). As before, you may assume that the value function is

$$v(z, r) = \begin{cases} \ln(1 + 8(z - r)) & \text{if } z \geq r \\ -\log_2(1 + 8(r - z)) & \text{if } z < r \end{cases}$$

for both puppies and kittens. For bureaucratic reasons, your policy must specify an integer number of dollars (e.g., spending \$2 on kittens is allowed, as is \$3, but \$2.33 is not). Given this, what amount of money should you spend on kittens and puppies if you want to maximise the perceived value $v(\mathbf{d}, \mathbf{r})$ of your spending policy? Include any code or mathematical work you used to arrive at your answer as well as the answer itself in the write-up.

**(1c)** What happens to your policy when the voters give up their expectations that you will deliver puppies? (i.e., set $r_p = 0$, $r_k = 4$) What is the optimal spending policy now? On average, how many kittens does it deliver, and how many puppies does it deliver? Why is the optimal policy different to the policy in part b?

[Hint: calculating "the average number of puppies" doesn't require any complicated programming. Because this is a binomial distribution, you can safely assume that the average number of puppies will be $20 * d_p/300$ and the average number of kittens will be $10 * d_k/150$. In other words, just divide the number of dollars spent by 15 and you'll arrive at the answer]

**(1d)** What if the expectations for kittens and puppies both drop to zero? (i.e., set $r_p = 0$, $r_k = 0$) What is the optimal spending policy now? On average, how many kittens does it deliver, and how many puppies does it deliver? Why does this policy differ from the previous two?

**(1e)** What about if the expectation increases to 10 in both cases? (i.e., set $r_p = 10$, $r_k = 10$) What is the optimal spending policy now? On average, how many kittens does it deliver, and how many puppies does it deliver? Why is this different to the best policy in part c?

2. **Sequential Sampling Models** (25 points)

In class we discussed a general family of "sequential sampling models" that human decision-makers might adopt when the utility function takes account of the amount of time it takes to make a particular decision. In these questions, we want to think about a slightly more general class of sequential sampling models than the simple random walk model that was derived in class. Specifically, the algorithm we want to consider is this one:

```
set time t = 0
set initial evidence x_0
do while |x_t| < γ_t
  increment time, t = t + 1
  collect sensory sample; and evaluate the log-odds for that sample, y_t
  increment evidence tally, x_t = x_{t-1} + y_t
if x_t ≥ γ_t, choose option A (at time t)
elseif x_t ≤ -γ_t, choose option B (at time t)
```

The only difference between this decision model and the random walk models discussed in class is that $\gamma$ is not necessarily a constant value: it may be different as a function of time $t$. (This isn't an arbitrary change to the algorithm – although it wasn't discussed in class, there are some computational analyses that predict that $\gamma$ should change over time).

**(2a)** In this question, $\gamma$ is just a constant (i.e., exactly the same algorithm as the one presented in class), and your task is to implement the random walk model using the simple information function,

$$y_t = \begin{cases} +1 & \text{with probability } p \\ -1 & \text{with probability } q \end{cases}$$

where $q = 1 - p$. As it happens, this analytic properties of this model are well studies: Feller (1968) proved that the probability that this algorithm terminates at the upper boundary $+\gamma$ (i.e., the probability that the decision-maker chooses $A$) is given by

$$P(\text{“choose } A\text{”}) = \begin{cases} \dfrac{w^{2\gamma} - w^{\gamma - x_0}}{w^{2\gamma} - 1} & \text{if } p \neq q \\ \\ 1/2 & \text{if } p = q \end{cases}$$

where $w = p/q$. Now that you've implemented the model yourself, your task it to verify that Feller's result is correct in the following way. Firstly, fix $p = .55$ and $x_0 = -1$, and estimate $P(\text{“choose } A\text{”})$ for $\gamma$ values of $2, 3, \ldots, 15$. For each value of $\gamma$, generate 1000 random decisions using your code. Plot the model output against Feller's predictions. Include the code you used to implement the model and Feller's predictions (but there is no need to include the code you used to generate the plots as long as you include the plots themselves).

**(2b)** Repeat the exercise above, this time fixing $x_0 = -1$, $\gamma = 5$, and varying $p$ as follows $.05, .10, .15, \ldots, .95$. Finally, try holding $p = .55$ and $\gamma = 10$, and varying $x_0$ as follows: $-9, -8, \ldots, 9$. As with (2a) you should plot the results against the predictions made by Feller's derivation. Since this is the same code with different parameters, there is no need to include the code in this part.

**(2c)** Alter the information function so that the increments $y_t$ are normally distributed with mean $\xi$ and variance 1,

$$y_t \sim \text{Normal}(\xi, 1).$$

Run the model with the parameters set at $\xi = .2$, $\gamma_0 = 3$, $x_0 = -1$ and include the changes you have made in the code in your write-up. What do the first passage time distributions look like? Plot them separately for the "choose A" (upper boundary) case and the "choose B" case? What is the probability that the decision maker chooses option $A$ (i.e., process terminates at $+\gamma$)? What is the average first passage time for those decisions where the decision-maker chooses $A$? Similarly, what is the probability that the decision-maker chooses option $B$, and what is the average response time for those decisions?

**(2d)** Alter your algorithm so that the decision threshold $\gamma$ now changes as a function of time (and, again, include the code that does this). Specifically, set

$$\gamma_t = \begin{cases} \gamma_0 - t/10 & \text{if } \gamma_0 \geq t \\ 0 & \text{otherwise} \end{cases}$$

Again, set the parameters to be $\xi = .2$, $\gamma_0 = 3$, $x_0 = -1$. Repeat the exercise from (2c) with this new model, answering the same questions.

**(2e)** At the end of your simulations from (2c) and (2d) you should have access to an estimate of the probability that the model makes choice $A$ at time $t$. If we let $P_c(r = A, t)$ denote this probability for the "constant threshold" model in (2c), and let $P_d(r = A, t)$ denote the corresponding probability for the "declining threshold" model in (d), then calculate the *difference* between the two $P_d(r = A, t) - P_c(r = A, t)$ for all $t \leq 50$. Report this difference either in the form of a table or (preferably) in the form of a plot. Do the same thing for choice $B$ (i.e., plot $P_d(r = B, t) - P_c(r = B, t)$. Include the plots/tables in your write-up. Can you explain what it is about these two models that makes these curves look the way they do? Why are the "choose $A$" and "choose $B$" plots different from one another?