MIT 14.13 – Problem Set 5

Please make sure to explain your answers carefully and concisely, i.e. do not simply write a numerical answer without an explanation of how you arrived at this answer. Answers without adequate explanation will not receive full credit.

Part 1: Climate change denial (40 points)

Climate change arrives tomorrow. It will either be very bad (event A) or very, very, very, very, very, very bad (not A). Today, it is common knowledge that climate change will only be "very bad" with probability $p \in [0, 1]$.

Charles is the President of France. Elections are the day after tomorrow. Charles discounts with a constant daily factor $\delta \in (0, 1)$ and his utility depends only on his re-election. If he is re-elected, Charles obtains utility of -4 (it is hard work to be president!). If he loses the election, Charles obtains -8.

His Facebook advertising team has used sophisticated machine learning techniques to determine that he will be re-elected if climate change is very bad, but not if climate change is very, very, very, very, very, very bad (his key constituents are located in coastal regions).

1. (5 points) What is Charles's expected utility today?

2. (12 points) Charles always tells the truth. He knows he will get questions today in the presidential debate about his view on the probability p of very bad climate change. These questions affect his utility today, which is now given in expectation by
$$f(p) + \delta^2 \mathbb{E}[u(\text{election result})].$$
In particular, he knows that if climate change seems likely to be very, very, very, very, very, very bad, Charles will get a lot of touch questions about his current government's response. So $f(p)$ is increasing in p.

   MIT climate researchers in Building 54 have discovered whether climate change will be very bad or very, very, very, very, very, very bad. Their *Nature* paper is embargoed until tomorrow, but Charles can email them this morning to find out the results of their discovery.

   Suppose $f(p) = \ln(1 + p)$. Will Charles email the MIT researchers? Does this depend on δ? Does this depend on Charles's payoffs from re-election?

3. (5 points) Now suppose instead $f(p) = (1 + p)^2$. Will Charles email the MIT researchers? Does this depend on δ? Does this depend on Charles's payoffs from re-election?

4. (5 points) Now suppose Charles can lie. He may announce whichever p he likes at the presidential debate to alter his payoff $f(p)$. Because the news cycle is only twenty-four hours, voters will not remember which p he announced when they go to the polls (i.e., they cannot punish him for being wrong or reward him for being right). Which p does Charles choose?

5. (8 points) As in question 4, Charles is no longer necessarily truthful and can announce any p in the debate. Now suppose that his announcement of p affects Congress' decision today about whether or not to set up sea walls to protect the coast. The sea walls are free, will guarantee Charles's re-election, but will only be installed if climate change is likely to be very, very, very, very, very, very bad—specifically, if and only if Charles's announcement satisfies $p \leq \theta < 1$.

   Taking p as given, for which values of θ will Charles announce $p = \theta$?

6. (5 points) Fix θ and the (true) p. Show that if Charles's discount factor is low enough (low δ), he will always lie and report $p = 1$ no matter what.

Part 2: In Search of Lost Time (40 points)

Note: Attribution bias will be covered in recitation on April 23 and 24.

Marcel travels between Paris and Balbec. Denote his location at each moment t by $s_t \in \{0,1\}$, where 0 is Paris and 1 is Balbec. At each moment t, he always does one of two things: spend time with his friend, Albertine ($a_t = 0$), or attempt to write his novel ($a_t = 1$).

Marcel's instantaneous utility $u(a_t, s_t)$ depends on

$$u(a_t, s_t) = \begin{cases} -a_t - 2(1-a_t) & \text{if } s_t = 0 \\ 5(1-a_t) & \text{if } s_t = 1. \end{cases} \tag{1}$$

Due to his unusual childhood, Marcel suffers from attribution bias based on his location. In moment t, he imagines that his utility for action a_τ at all moments $\tau \geq t$ will equal

$$\hat{u}_t(a_\tau, s_\tau) = (1-\gamma)u(a_\tau, s_\tau) + \gamma u(a_\tau, s_{t-1}(a_\tau)), \tag{2}$$

for some $\gamma \in [0,1]$, where $s_{t-1}(a) \equiv \{s_{\tau^*} : \tau^* = \max\{t' \leq (t-1) : a_{t'} = a\}\}$ is the last place that he took action a prior to t.

1. (5 points) Interpret Marcel's instantaneous utility function. Which action a_t does Marcel prefer when he is in Paris? Which does he prefer in Balbec? Where would he rather be?
2. (5 points) Interpret Marcel's imagination of his utility in future periods. What does γ measure? What special cases do $\gamma = 0$ and $\gamma = 1$ represent?
3. (12 points) Marcel now faces a conundrum at t. He is in Paris, and has just received a letter from Albertine. She has written to tell him that she will only continue to spend time with him if they marry. In this case, Marcel must then spend all of his time with her.

   In addition, sea level rise has made future trips to Balbec impossible, so Marcel is stuck in Paris forever ($s_\tau = 0$ for all $\tau \geq t$).

   Marcel is utility-maximizing; specifically, he will choose the course of action that maximizes the utility $\hat{u}_t(a_\tau, s_\tau)$, which is the utility that he imagines he will have in each moment $\tau > t$ for the rest of his life.

   1. If Marcel most recently spent time with Albertine in Paris and attempted to write his novel in Paris, will he propose the marriage? Does your answer depend on γ? Why (not)?
   2. If Marcel most recently spent time with Albertine in Paris and attempted to write his novel in Balbec, will he propose? Does your answer depend on γ? Why (not)?
   3. If Marcel most recently spent time with Albertine in Balbec and attempted to write his novel in Paris, will he propose? Does your answer depend on γ? Why (not)?
   4. If Marcel most recently spent time with Albertine in Balbec and attempted to write his novel in Balbec, will he propose? Does your answer depend on γ? Why (not)?
4. (8 points) In which of the above cases (if any) would you say Marcel makes a mistake (defined as: being better off had he made a different choice)? And if he does make a mistake, do you think he will ever realize this mistake? Can you think of potential policies that Marcel's mother could impose to help her son avoid such mistakes?

Part 2, continued. Time Regained.

Before Marcel can reply to her letter, he receives news that Albertine has died in an unfortunate horseback-riding accident.

Now that he has to write his novel, Marcel faces a new conundrum in every moment t: should he work hard to write his novel ($w_t = 1$), or daydream about being a famous writer, but in fact do nothing ($w_t = 0$)?

With Albertine out of the picture, Marcel's preferences have changed.

His short but vivid memory means that he values what he does at t based on what he remembers doing at $t-1$. That is, Marcel's instantaneous utility at moment t depends on his choice w_t and his state s_t, where the state is now his previous action ($s_t = w_{t-1}$). His utility is given by

$$u(w_t, s_t) = \begin{cases} 5(1 - w_t) & \text{if } s_t = 1 \\ -2(1 - w_t) - 12w_t & \text{if } s_t = 0. \end{cases} \tag{3}$$

For example, if Marcel daydreamed at $t - 1$, so that $s_t = 0$, then he has become accustomed to daydreaming, so writing incurs a huge cost (-12), whereas being lazy gives him -2.

Despite his remarkable literary talent, Marcel suffers from projection bias. At time t, he predicts his future utility of his choice w_τ at time $\tau > t$ to be

$$\hat{u}_t(w_\tau, s_\tau) = (1 - \alpha) \cdot u(w_\tau, s_\tau) + \alpha \cdot u(w_\tau, s_t), \tag{4}$$

where $\alpha \in [0, 1]$.

5. (5 points) Provide a brief interpretation of Marcel' situation.
   1. Does Marcel enjoy writing?
   2. How does writing today affect the level of his utility tomorrow?
   3. How does writing today affect his marginal utility of writing tomorrow?
   4. What kind of good would you say writing is for Marcel?

6. (5 points) What does α measure? Which does it mean for α to equal 0? Which does it mean for α to equal 1?

Part 3: Learning from antibody tests (up to 20 extra credit points)

Note: This question is optional extra credit.

Suppose that the fraction of people with SARS-CoV-2 antibodies is 2.8%. Suppose also that there exists a lateral flow immunoassay test, which

(i) correctly detects the antibody if it is present with probability $\alpha \in [0, 1]$; and

(ii) correctly rejects the presence of the antibody if it is not present with probability $\beta \in [0, 1]$.

1. (4 points) Suppose that $\alpha = 0.938$, and $\beta = 0.956$, as with the first FDA-approved antibody test (from Cellex). Researchers are sampling the population at random for testing. Pete is chosen and tests positive for the SARS-CoV-2 antibody. What is the probability that Pete has the antibody?

2. (4 points) Let the fraction of people with SARS-CoV-2 antibodies be $\pi \in [0, 1]$, instead of just the previous 2.8%. Still taking $\alpha = 0.938$ and $\beta = 0.956$, and still assuming Pete has tested positive, plot the probability that Pete has the antibody as a function of π over the range $[0, \frac{1}{2}]$.

3. (4 points) How large does the prevalence π need to be for a positive test to mean that Pete does, in fact, have antibodies with probability 93.8%?

   (An approximate answer is fine).

4. (4 points) Now ask four of your friends who are NOT taking 14.13 (or fell asleep during slides 60–61 of lecture 15/16) the question in (1). What is the average of the answers you receive? If your friends get the answer wrong, why do you think that is? What concept discussed in class can explain the mistake?

5. (4 points) Now suppose that the government is interested in allowing people to leave quarantine if they test positive for the antibody. In addition, the government wants to implement the policy soon, i.e., when their epidemiology models predict that 5% of the overall population has the antibody.

   The government will use the next test approved by the FDA if, when someone tests positive for the antibody, they have the antibody at least 99% of the time.

   Suppose that $\alpha = \beta = \overline{p}$ for this next test. What is the minimum $\overline{p}$ for which the FDA should certify the test to ensure the governmental assumption of 99% accuracy described above?