ECN 190A: Behavioral Economics
A Study Guide

Christina Sun

This version: April 2023

1 Basic Probability Theory

I have uploaded some resources on probability theory and I would highly recommend you to review these. They are under `Christina'sSections/probability_resources`.

You might find that you fit into one of the following cases:

1. If you have already taken a course on probability theory and just need a quick refresher, you can check out `Introprobability_appendix_simon_blume.pdf`.

2. If you know probability theory, want a referesher, and also some practice problems, I recommend `probabilitylecturenotes_gravner.pdf`, which is the lecture notes of a professor at UC Davis that I got from the internet and seems to have a lot of good examples.

3. Finally, **if you have never learned probability theory**, or do not remember anything from it (*cue: you cannot remember what a set or a partition is, or you cannot remember how to calculate the conditional probability*), then I would **strongly recommend** `Probability_Hogg.pdf` for a more comprehensive introduction to basic probability theory. This was the first chapter from the textbook I used to learn probability in college, and I find it very helpful. **Please make sure you understand probability, as it is the bread and butter of this course!**

1.1 Probability Axioms

Let set C be a sample space of all possible outcomes. An event A is a subset of C. Let the set B be a collection of events. The probability P of every event A satisfies the following 3 axioms:

1. $P(A) \geq 0 \ \forall A \in B$
2. $P(C) = 1$
3. If $A_1, A_2, ...$is a sequence of pairwise disjoint events (what does pairwise disjoint mean? $A_i \cap A_j = \emptyset$ for all $i \neq j$), then $P(\cup_{n=1}^{\infty} A_i) = \sum_{i=1}^{\infty} P(A_i)$

Example: the sample space for my midterm score is $C = \{0, 1, 2, 3, ..., 100\}$

1.2 Useful Theorems

Theorem 1. *If A_1 and A_2 are events such that $A_1 \subset A_2$, then $P(A_1) \leq P(A2)$*

Theorem 2. *For events $A_1 \in B$ and $A_2 \in B$, $P(A_1 \cup A_2) = P(A_1) + P(A_2) - P(A_1 \cap A_2)$*

Theorem 3. *For mutually exclusive events A_1, A_2 and $A_1 \cap A_2 = \emptyset$, $P(A_1 \cup A_2) = P(A_1) + P(A_2)$*

Definition 1 (Conditional Probability)**.** *The conditional probability of A_1 given A_2 is $P(A_1|A_2) = \dfrac{P(A_1 \cap A_2)}{P(A_2)}$*

Definition 2 (Independence)**.** *Events A_1 and A_2 are independent if and only if $P(A_{12}) = P(A_1) \times P(A_2)$*

Theorem 4 (Law of Total Probability)**.** *If $B_1, B_2, B_3, ...,$ is a partition of the sample space S, then for any event A we have*

$$P(A) = \sum_i P(A \cap B) = \sum_i P(A|B_i)P(B_i) \tag{1}$$

1.3 Continuous Random Variables

An understanding of random variable and probability distributions is needed to fully understand what is going on in the sequential search model. If you do not remember what a probability density function is, please review page 32-56 of Probability_Hogg.pdf on Canvas in my sections folder under handouts!!! It gives a rigorous treatment of probability distributions, whereas the brief review I provide here is not comprehensive and leaves out some details.

A **random variable** X is a function that assigns to each element c in the sample space $\mathcal{C}$ one and only one number X(c).

The **range** of X is the set of real numbers that it can take, $\mathcal{D} = \{x : x = X(c), c \in \mathcal{C}\}$.

Example 1. *Suppose the midterm has 100 multiple choice questions and each question is worth 1 point. The sample space of the outcome of the midterm is the set {Get 0 question correct, get 1 question correct, ..., get 100 questions correct}. The* <u>*score*</u> *of my midterm is a random variable X which assigns a number each simple event in my sample space. So the range of X is the set* $\mathcal{D} = \{, 1, 2, ..., 100\}$.

Given a random variable X, its range $\mathcal{D}$ becomes the sample space of interest, and X induces a probability distribution of X. Intuitively speaking, a random variable is just a function that gives a number to each possible outcome that we are interested in, so that it's easier to study them.

In the above example, the X is a **discrete random variable**. The range of a discrete random variable is a finite or countable set. In this case of the midterm score, X has a finite range. And it gives a probability distribution of X on $\mathcal{D}$, which can be described by $p_X(x) = P(\{c : X(c) = x\})$. The function $p_X(\cdot)$ is called the **probability mass function** of X.

1.3.1 Expectations of Discrete Random Variables

Expectations of discrete probability distributions is calculated by summing the product of the value of the random variable and its associated probability, taken over all of the values of the random variable, i.e. a weighted average of all possible outcomes weighted by the probability associated with each outcome.

$E(X) = \sum P(X = x) \times x$

E.g. the expected value of my midterm grade is $\sum_{x=1}^{100} P(x)x$

1.3.2 Continuous Random Variable

Definition 3. *A random variable X is a* ***continuous random variable*** *if its cumulative distribution function* $F_X(x)$ *is a continuous function for all* $x \in R$

We can think of a continuous random variable as one which takes an infinite number of possible values (the range is uncountable). Continuous random variables are usually measurements. Examples include height, weight, the amount of sugar in an orange, the time required to run a mile.

1.3.3 CDF

In general, given a random variable, we can define a **Cumulative Distribution Function** for it.

Definition 4 (Cumulative Distribution Function)**.** *Let X be a random variable. Then its **cumulative distribution function** (cdf) is defined by $F_X(x)$, where*

$$F_X(x_= P_X((-\infty, x]) = P(\{c \in \mathcal{C} : X(c) \leq x\}) \tag{2}$$

We usually shorten $P(\{c \in \mathcal{C} : X(c) \leq x\})$ to $P(X \leq x)$

We can think about the CDF as accumulating the probability of X being less than or equal to x.

What are the properties of the CDF?

Theorem 5. *Let X be a random variable with cumulative distribution function F(x). Then*

1. *For all a and b, if $a < b$, then $F(a) \leq F(b)$*
2. $\lim_{x\to-\infty} F(x) = 0$ *(intuition: the probability of nothing is 0)*
3. $\lim_{x\to\infty} F(x) = 1$ *(intuition: when you accumulate the probability of everything together, it should add up to 1)*

It follows from the definition of the cdf that:

Theorem 6. *Let X be a random variable with cdf F_X. Then for $a < b$, $P(a < X \leq b = F(b) - F(a)$.*

1.3.4 PDF

Given a continuous random variable, the **Probability Density Function** is $\dfrac{d}{dx}F_X(x) = f_X(x)$. By the Fundamental Theorem of Calculus, we have $F_X(x) = \int_{-\infty}^{x} f_X(t)dt$.

We have:

$$P(a \leq x \leq b) = F_X(b) - F_X(a) = \int_a^b f_X(t)dt \tag{3}$$

And for continuous random variables,

$$P(a < x \leq b) = P(a \leq x < b) = P(a \leq x \leq b) = P(a < x < b) \tag{4}$$

$P(X = a) = 0$ for all $a \in R$. This is because a continuous random variable can take an infinite number of values, so if we assign positive probability for a point, then when you add all of them together they will be greater than 1!

The **support** of a continuous random variable X consists of all points x such that $f_x(x) > 0$.

What are the properties of the PDF?

1. $f_X(x) \geq 0$
2. $\int_{-\infty}^{\infty} f_X(t)dt = 1$

1.3.5 Expectation

We already know the expectation (aka expected value, mean) of a discrete random variable X is $E(X) = \sum_x xp(x)$.

The **expectation** of a continuous random variable X is

$$E(X) = \int_{-\infty}^{\infty} xf(x)dx \tag{5}$$

1.3.6 Uniform Distribution

Any continuous or discrete random variable X whose pdf or pmf is constant on the support of X is said to have a **uniform** distribution.

The probability density function of a continuous uniform distribution is

$$f_X(x) = \begin{cases} \frac{1}{b-a} & \text{for } a \leq x \leq b \\ 0 & \text{for } x < a \text{ or } x > b \end{cases} \tag{6}$$

The mean of random variable $X \sim unif(a, b)$ is $\frac{1}{2}(a + b)$.

2 Risk Preferences

Definition 5 (Risk Aversion). *When offered a choice between a lottery and the expected value of the lottery for sure, the decision maker is:*

- *Risk averse if* $EV(L) \succ L$
- *Risk neutral if* $EV(L) \sim L$
- *Risk loving if* $EV(L) \prec L$

Definition 6 (Certainty Equivalent)**.** *The certainty equivalent of a lottery is the amount of guaranteed money that would make the decision maker indifferent between taking the money and accepting the lottery. In other words,* $CE \sim L$*, and under expected utility,* $u(CE) = EU(L)$

Example 2. *What is the certainty equivalent of the lottery L = (10, 0.5; 20, 0.5) if the utility function is* $u(x) = x^2$*?*

Answer: Solve $10^2 \times 0.5 + 20^2 \times 0.5 = x^2 \implies 250 = x^2 \implies x = 5\sqrt{10}$

Example 3. 1. *Assume expected utility with* $u(x) = x^2$*. Is it concave or convex? What are the EU and the CE of the lottery L = (10, 0.5; 0, 0.5)?*

- *Visual trick for convex and concave: shape of function. Concave function increases at a decreasing rate, convex function increases at an increasing rate.*
- *Solution: Convex because* $u'' > 0$*.* $EU(L) = u(10) \times 0.5 + u(0) \times 0.5 = 50$. $u(CE) = EU(L) \implies CE = \sqrt{50} = 5\sqrt{2}$

2. *Assume expected utility with* $u(x) = \sqrt{x}$*. Is it concave or convex? What is the CE of L now?*

- *Solution: Concave because* $u'' < 0$*.* $u(CE) = EU(L) = \sqrt{10} \times 0.5 \implies CE = 2.5$

General result: Under EU, when u is concave, $CE(L) < EV(L)$; when u is convex, $CE(L) > EV(L)$ (See section 1 annotated slides for graph representation of this result)

3 Expected Utility Theory

3.1 From Preferences to Utility

1. Preferences: they describe how a decision maker chooses between different objects Start with a set of possible choices: choice set. A person's preferences tell us how they will choose between objects in this choice set.

2. Utility: a utility function is a representation of the preference structure. It maps each choice object to a utility number.
   Expected Utility is just one (albeit a widely used) way of representing preferences under uncertainty.

3.2 Expected Utility Theory

3.2.1 Preferences Under Expected Utility

Q: What does EUT say about preferences?
A: A preference relation can be represented by EU if there is a utility function u such that for two lotteries $L_1 = \{x_1, p_1; x_2, p_2; ...; x_k, p_k\}$ and $L_2 = \{x_1, q_1; x_2, q_2; ...; x_k, q_k\}$, $L_1 \succeq L_2 \iff \sum_{i=1}^{k} u(x_i)p_i \geq \sum_{i=1}^{k} u(x_i)q_i$

Intuition: preferences can be represented by EU if we can calculate utility of a lottery by calculating the *expected value* over the utilities of each outcome in the lottery. Hence *Expected* Utility

Remember how we calculate the expectations.

3.2.2 Expected Utility Theorem

Theorem 7 (Expected Utility Theorem)**.** *Preferences satisfying Axioms A1, A2, A3* $\iff$ *Preferences can be represented by EUT*

Axiom 1 (A1: Completeness and Transitivity (aka Rationality))**.** *Intuition: can compare any two options, and preferences cannot be circular (this is easy to think about in terms of being rational)*

Axiom 2 (A2: Continuity)**.** *Intuition: this basically says that you can always mix a good lottery and a bad lottery to get something in the middle. i.e. Nothing is so bad that you won't risk it a teeny tiny bit for a high enough chance of a good outcome. Or, there are no jumps or holes in your preferences, hence continuity.*

Axiom 3 (A3: Independence)**.** *For any* $L_1, L_2, L_3 \in \mathcal{L}$, $L_1 \succ L_2 \iff \alpha L_1 + (1-\alpha)L_3 \succ \alpha L_2 + (1-\alpha)L_3$ $, \forall \alpha \in (0, 1]$. *Same holds for* $\sim$.

This is also called the independence of irrelevant alternatives. **Intuition**: if I prefer an apple over a banana, I should still prefer half an apple + half an orange over half a banana + half an orange. The orange, or L_3 here is the irrelevant alternative.

Example 4 (Independence Axiom Practice, HW1 Q3)**.** $L_1 = \{10, 0.5; 0, 0.5\}, L_2 = \{4, 1\}. L_1 \sim L_2$. *Use IA to show* $L_1 \sim L_3 = \{10, 0.25; 4, 0.5; 0, 0.25\}$.

Solution: Notice that $L_3 = 0.5L_1 + 0.5L_2$ so we can use the IA to get $0.5L_1 + 0.5L_1 \sim 0.5L_2 + 0.5L_1 = L_3$. The trick is to first write the relationship between two lotteries and combine them with a third lottery. In this case you are combining a lottery with itself.

We talked about 2 cases where preferences violate the EUT (both by violating the Independence Axiom).

1. Ellsberg Paradox
   Ambiguity aversion is an explanation.

2. Allais Paradox
   Rank Dependent Utility is an explanation.

A lot of behavioral economics involves finding cases where existing theories don't work and proposing alternative theories that explain these cases!

3.2.3 Ellsberg Paradox Revisited

An urn with 1 R, Y+B=2 (could be 0)

Choice 1:

- Bet 1: \$10 if R, 0 otherwise
- Bet 2: \$10 if Y, 0 otherwise

Choice 2:

- Bet 3: \$10 if R or B, 0 otherwise
- Bet 4: \$10 if Y or B, 0 otherwise

A decision maker who assumes the worst would choose $Bet1 \succ Bet2$ and $Bet3 \prec Bet4$

Suppose $u(x) = x$. Suppose the DM is an EU maximizer. What can we say about them? The first choice implies $P(R)(10) + (1 - P(R)) \times > P(Y) \times 10 + (1 - P(Y)) \times 0$, which implies $P(R) > P(Y)$. The second choice implies $P(R) < P(Y)$. Contradiction! So the DM violates EU.

3.2.4 Allais Paradox Revisited

Allais presented experimental subjects with two choices.

Choice 1 is between L_1 and L_2, where $L_1 = \{1m, 1\}$ and $L_2 = \{5m, 0.1; 1m, 0.89; 0, 0.01\}$. Choice 2 is between L_1' and L_2', where $L_1' = \{5m, 0.1; 0, 0.9\}$ and $L_2' = \{1m, 0.11; 0, 0.89\}$. Allais found that most people choose $L_1 \succ L_2$ and $L_1' \succ L_2'$. Now this is a problem because it violates Expected Utility Theory by violating the Independence Axiom.

How do we show it?

To show it violates EUT is easy, you just assume EUT and calculate the expected utility of the 4 lotteries and you will see that the preferences displayed here imply a contradiction.

To show it violates IA specifically, we have to do some mental and mathematical gymnastics. The basic idea is this: we assume that IA holds, derive some implications, and then eventually arrive at a contradiction. Then IA should not hold, and it is thus violated. One way we can do this is by showing that under IA, $L_1 \succ L_2$ implies $L_2' \succ L_1'$ (which would be a contradiction).

Proof. Assume the IA holds. Then we have $L_1 \succ L_2 \iff \{1m, 1\} \succ 5m, 0.1; 1m, 0.89; 0, 0.01\}$. Now we see that L_2 has the prizes from L_1 and L_1', so we want to break L_2 apart into a combination of a lottery with only 1m as the prize and another lottery with both 5m and 0 as possible prizes. To do this, we can see that $L_2 = 0.11(5m, p_1; 0, p_2) + 0.89(1m, 1)$, and from this equation we can get $p_1 = \frac{0.1}{0.11}$ and $p_2 = \frac{0.01}{0.11}$. We see also that $(1m, 1)$ is just L_1. Now, $L_1 = 0.11L_1 + 0.89L_1$, so we have $0.11L_1 + 0.89L_1 \succ 0.11(5m, \frac{0.1}{0.11}; 0, \frac{0.01}{0.11}) + 0.89L_1$. By the IA (we are using the "$\Longleftarrow$" direction of the IA), we have $0.11L_1 \succ 0.11(5m, \frac{0.1}{0.11}; 0, \frac{0.01}{0.11})$. Now we can see that $0.11L_1 + 0.89(0, 1) = L_2'$ and $0.11(5m, \frac{0.1}{0.11}; 0, \frac{0.01}{0.11}) + 0.89(0, 1) = L_1'$, so using IA again (this time the "$\Longrightarrow$" direction), we have $L_2' \succ L_1'$, which is a contradiction! Therefore, the assumption that IA holds must not be true.

□

4 Rank Dependent Utility

Steps for calculating the RDU of a lottery:

1. Rank the prizes of the lottery $x_1 \succ x_2 \succ x_3 \succ ... \succ x_n$

2. For lottery $L = \{x_1, p_1; x_2, p_2; ...; x_n, p_n\}$, we calculate the RDU as $U(L) = \sum_{i=1}^{n} \pi(i)u(x_i)$, where $\pi_i = g(p_1 + p_2 + ... + p_i) - g(p_1 + p_2 + ... + p_{i-1})$
   Intuition: g(probability of getting at least as good as x_i) - g(probability of getting strictly better than x_i)

$g : [0, 1] \to [0, 1]$ is increasing, $g(0) = 0$ and $g(1) = 1$

Intuition for RDU:

We take EUT and introduce something extra: the **probability weighting function** g. The RDU is basically a *weighted average* of the utilities from each prize, weighted using the function g (strictly speaking weighted using a function of g). Compare with EU, which is a weighted average with weighting function $p(x)$. We are essentially multiplying the utility of each prize $u(x_i)$ by the term (g(probability of getting at least as good as x_i) - g(probability of getting strictly better than x_i)) and adding this together for all prizes in the lottery.

Example 5 (RDU: HW1 Q2)**.** *In the Allais paradox, most people choose* $L_1 = \{1m, 1\} \succ L_2 = \{5m, 0.1; 1m, 0.89; 0, 0.01\}$ *and* $L_3 = \{5m, 0.1; 0, 0.9\} \succ L_4 = \{1m, 0.11; 0, 0.89\}$ *Let* $u(0) = 0$*.*

1. *Show that under EU, one would get* $0.1u(1m) > 0.11u(5m)$ *and* $0.1u(1m) < 0,11u(5m)$*, which cannot hold simultaneously.*

2. *Show that under RDU with* $g(p) = p^2$*, these two conditions become* $\frac{u(5m)}{u(1m)} < \frac{1 - 0.99^2 + 0.1^2}{0.1^2}$ *and* $\frac{u(5m)}{u(1m)} > \frac{0.11^2}{0.1^2}$*. Can this hold simultaneously?*

Example 6 (RDU and Risk Aversion: HW1 Q4)**.** *Assume RDU with* $u(x) = x^2$ *and* $g(p) = p^2$*. Is* u *convex or concave? What is the certainty equivalent of* $\{10, 0.5; 0, 0.5\}$*? Does the individual prefer the lottery or its expected value for certain? Is this preference consistent with risk aversion or with risk seeking?*

See annotated section 2 slides for step-by-step solution for these two RDU examples.

5 Reference Dependent Utility

Intuition: before you make a choice, you have a reference point. To calculate the reference dependent utility of a choice, you add up 2 parts of the utility together. For example, the reference point is 0 mug and 1 pen and the outcome is 1 mug and 0 pen.:

1. Consumption utility. This is same as usual, just calculate the utility of the outcome from the choice, which is $u(1mug, 0pen)$

2. Gain/loss utility. Here you are comparing the **outcome** with the **reference point**, and see what you have gained and what you have lost relative to the reference. Remember that *gains and losses are always determined relative to the reference point!!* For some problems (like in this example) we might find ourselves in the situation where we gain one thing but lose another. In this case, we separate them and calculate the gain utility for the things we gained, and calculate the loss utility for the things we lost.

   - For gains, the gain utility is $\eta(u(outcome) - u(reference)) = \eta(u(1pen) - u(0pen))$.
   - For losses, the loss utility is $\eta\lambda(u(outcome) - u(reference)) = \eta\lambda(u(0mug) - u(1mug))$
   - then you add the gain utility and loss utility together to get the total gain/loss utility of $\eta(u(1pen) - u(0pen)) + \eta\lambda(u(0mug) - u(1mug))$

3. Finally, add up consumption utility and gain/loss utility to get the overall reference dependent utility. $U((1mug, 0pen)|(0mug, 1pen)) = \eta(u(1pen) - u(0pen)) + \eta\lambda(u(0mug) - u(1mug))$

Note: Here $\lambda \geq 1$ so the losses factor into your utility more than the gains. This is loss aversion. Also, we are always going to assume that $u(x) = x$ for simplicity unless told otherwise.

A relevant definition to know is personal equilibrium.

Definition 7 (Personal Equilibrium)**.** *A choice is called a personal equilibrium for a decision maker if whenever she expects to make that choice, it is indeed the optimal choice.*

What does this mean?

Personal Equilibrium describes a **choice** of the decision maker. Notice that the definition for personal equilibrium already prescribes the *reference point* which we will use, which is that same choice itself. And it being the optimal choice means that the utility from making this choice has to be greater than or equal to all other choices (a choice is better than some other choice if it has a higher utility, this is how we can rank choices), given that she expects to make that choice. Since we are in a reference dependent utility world, translating this into math, we have that a choice x is a personal equilibrium if $U(x|x) \geq U(y|x)\ \forall y \neq x$.

Armed with the mathematical translation, now you should know what to do when you are asked to show that something is a personal equilibrium: simply by showing that $U(\textit{that something}|\textit{that something}) \geq U(\textit{everything else}|\textit{that something})$.

5.1 Example: Calculating Reference Dependent Utility

Example 7 (Reference Dependent Utility). *Let $\lambda = 2, \eta = 2$. p is price for a mug. Take any $\frac{3}{5} < p < \frac{5}{3}$. Show that at such a price, there are two Personal Equilibria.*

1. *Equilibrium 1: If you start by expecting to buy at price p, then it is your best response to indeed buy the mugs.*
2. *Equilibrium 2: If you start by not expecting to buy at price p, then it is your best response to not buy the mugs.*
3. *What is the utility at each equilibrium outcome? Which outcome has higher utility? (Your answer should depend on whether $p \geq 1$ or $p \leq 1$)*

Solution:

Equilibrium 1:

We want to compare $U(buy|buy)$ with $U(notbuy|buy)$. Reference point $r = (1, -p)$

$$\begin{aligned} U(buy|buy) &= m(1,-p) + \eta(m(1,-p) - m(1,-p)) \\ &= m(1,-p) \\ &= 1 - p \end{aligned} \tag{7}$$

$$\begin{aligned} U(notbuy|buy) &= m(0,0) + \eta(m(0) - m(-p)) + \eta\lambda(0-1) \\ &= -\eta\lambda + \eta p \\ &= \eta(p - \lambda) \\ &= 2(p-2) \end{aligned} \tag{8}$$

Now since $\frac{3}{5} < p < \frac{5}{3}$, we have

$$\frac{-2}{3} < 1 - p < \frac{2}{5} \tag{9}$$

and

$$\frac{-14}{5} < 2(p-2) < \frac{-2}{3} \tag{10}$$

Therefore we have $1-p > 2(p-2)$ so buying the mug when you expect to buy is a personal equilibrium.

Equilibrium 2:

Now we want to compare $U(buy|notbuy)$ with $U(notbuy|notbuy)$. Reference point $r = (0,0)$

$$\begin{aligned} U(buy|notbuy) &= m(1,-p) + \eta(1-0) + \eta\lambda(-p-0) \\ &= 1 - p + \eta - \eta\lambda p \\ &= 3 - 5p \end{aligned} \tag{11}$$

$$\begin{aligned} U(notbuy|notbuy) &= m(0,0) + \eta(m(0,0) - m(0,0)) \\ &= 0 \end{aligned} \tag{12}$$

Since $\frac{3}{5} < p < \frac{5}{3}$, we have

$$-16/3 < 3 - 5p < 0 \tag{13}$$

So $U(notbuy|notbuy) > U(notbuy|notbuy)$ so not buying when you expect not to buy is a personal equilibrium.

Which equilibrium has higher utility:

At personal eqm 1: $U_1 = U(buy|buy) = 1 - p$.

At personal eqm 2: $U_2 = U(notbuy|notbuy) = 0$

When $p < 1$, $U_1 > U_2$; when $p > 1$, $U_1 < U_2$; when $p = 1$, $U_1 = U_2$.

5.2 Example: Reference Dependent Utility under Uncertainty

3) Suppose in an experiment, a subject is randomly assigned a mug. She is told that with probability, q, she can choose to freely exchange the mug for a pen, and with probability $1-q$ she will be stuck with the default allocation of the mug.
If the subject chooses to exchange when allowed, her stochastic reference point is $r_1 = \underbrace{q}_{\text{choose}} \underbrace{(0,1)}_{\text{(mug,pen)}} + \underbrace{(1-q)}_{\text{default}} \underbrace{(1,0)}_{\text{(mug,pen)}}$. If the the subject chooses not to exchange when allowed, her stochastic reference point is $r_2 = \underbrace{(1,0)}_{\text{(mug,pen)}}$. Let the utility of mug be u_m and the utility of pen be u_p.
Take $\eta = 2, \lambda = 2$. Suppose she expects to not exchange, that is, her reference point is r_2.
[10 points] Show that the reference-dependent utility of exchanging is

$$qu_p + (1-q)u_m + 2q(u_p - 2u_m)$$

[5 points] Show that the reference-dependent utility of not exchanging is u_m.
[5 points] Finally, show that not exchanging is a personal equilibrium if $u_m \geq \frac{3}{5}u_p$.

Here onwards, suppose that she expects to exchange, and hence, her reference point is r_1.
[10 points] Show that the reference-dependent utility of exchanging is

$$qu_p + (1-q)u_m - 2q(1-q)(u_p + u_m)$$

[10 points] Show that the reference-dependent utility of not exchanging is $u_m + 2q(u_m - 2u_p)$.
[5 points] Finally, show that exchanging is a personal equilibrium if $q \geq \frac{5u_m - 3u_p}{2(u_m+u_p)}$.

Solution: (this is going to involve a lot of words, I am trying to explain the reasoning as clearly as possible, at the risk of being verbose)

5.2.1 Part 1

Here we are assuming the reference point is not exchanging, i.e. $r_2 = (1,0)$

Now we want $U(exchange|not\ exchanging)$. The reference point is easy, we already have it. But what about the outcome from making the choice of exchanging? The problem is that even if the Decision Maker (DM) wants to exchange, she will not always be able to. In this world, with probability q she is allowed to exchange, so since she wants to exchange, she will. With probability 1-q, she is not allowed to exchange, so even if she wants to exchange, she cannot. So our final outcome from choosing to exchange is not a certain outcome, but a probabilistic one. Let's try to think about the possible outcomes:

1. With probability q, the DM can exchange, thus she will exchange. Outcome is (0,1) lose 1 mug, gain 1 pen
2. With probability (1-q), the DM cannot exchange, so she will not exchange. Outcome is (1,0)
   no loss, no gain

So, the final outcome is $q(0,1) + (1-q)(1,0)$, and so

$$\begin{aligned}
U(exchange|not\ exchange) &= U(q(0,1) + (1-q)(1,0)|(1,0)) \\
&= qu_p + (1-q)u_m \ \ \text{consumption utility} \\
&+ q((u_p - 0) + \eta\lambda(0 - u_m)) \ \ \text{gain-loss utility from case 1} \\
&+ (1-q)0 \ \ \text{gain-loss utility from case 2} \\
&= qu_p + (1-q)u_m + 2q(u_p - 2u_m)
\end{aligned}$$

What is important to understand here is that with reference dependent utility, there are always two things at play, and you need both of them in order to calculate the utility:

1. The reference point (expectation): this is what's in your head
2. The actual outcome (what you actually do): this is what actually happens in the real world

Here, you expect to not exchange, so whether or not you actually are allowed to exchange doesn't matter for your reference point: you expect always to not exchange, so your reference point is (1,0). But given that reference point, the real world can turn out one of two ways:

1. Our DM actually are allowed to exchange, and here since we are calculating U(exchange — not exchange), that means our DM will exchange if given the chance, so she will exchange. The actual outcome is (0,1). This outcome happens with probability q
2. Our DM is not allowed to exchange. When this is the case, even though she wants to exchange, she cannot. So the actual outcome is (1,0). This outcome happens with probability (1-q).

For each of these outcomes, we compare against the reference point to get the gains and losses. and the overall reference dependent utility of exchanging when we expect not to exchange is q(reference dependent utility of situation 1) + (1-q)(reference dependent utility of situation 2) = q(consumption utility of situation 1 + gain/loss utility of situation 1) + (1-q)(consumption utilitiy of situation 2 + gain/loss utility of situation 2)

5.2.2 Part 2

If the DM chooses not to exchange, the outcome is always (1,0), so the gain-loss utility here is 0.

So $U((1,0)|(1,0)) = u_m$

5.2.3 Part 3

To show that not exchange is a personal equilibrium, we have to show that

$U(\textit{not exchange}|\textit{not exchange}) \geq U(\textit{exchange}|\textit{not exchange})$

We already have both sides from part 1 and part 2, so we can just plug in the expressions for them, simplify, and get $u_m \geq \frac{3}{5}u_p$

You should try to do this yourself!

5.2.4 Part 4

Here we are switching our reference point to r_1, that is, now the DM will expect to exchange. We want $U(\textit{exchange}|\textit{exchange})$. Let's ignore the complicated expression of r_1 for now and think about is actually going on.

Here the problem is more complicated because our reference point is no longer certain. With probability q, our DM expects to be allowed to exchange so she expects to exchange.

With probability 1-q, she expects not to be able to exchange, so she expects not to exchange. This is why the reference point is this probabilistic thing given in the problem.

Now, again we have to think about the possible outcomes, and here we have to think about **two layers:** the first layer of what the **expectation (reference point)** is, and the second layer of what the **actual outcome is under the given expectation**. We can only know what the loss and the gain is once we know both the reference point and the actual outcome.

Recall from part 1 where we talked about how the reference point is what happens in your head (because it's your expectation), versus the actual outcome is what happens in the real world. And what we need to do is to find the possible outcomes that can actually happen in the real world *given* our expectation in our heads.

How do we do this?

First, let's think about the first layer, our expectation/reference point. The difference between this part and part 1 is that in part 1, our reference point is fixed, i.e. not exchanging. But here when we expect to exchange, our reference point is no longer fixed, because our DM knows that she doesn't always get to exchange. So our DM's expectation will be probabilistic and is no longer certain. So she will expect to exchange with probability q and expect not to exchange with probability (1-q). So with probability q, the reference point is (0,1), and with probability (1-q), (1,0).

Now, given the possible situations in our DM's head, we have found the first layer, the reference point. Next we will find what are the possible outcomes that can actually happen in the real world. Since our expectation is in our head, it doesn't have anything to do with the real world. So given each expectation, there is q chance we are actually allowed the exchange, and (1-q) chance we are not allowed. We are calculating the utility of U(exchange—exchange), so our DM will want to exchange when she can.

Here's what all the possible situations of the world looks like when she wants to exchange under the expectation of exchanging (here the first level is the expectation, the situation of the world in her head; the second level is the actual outcome, the actual situation of the world in the real world. The situation of the world in the head does not impact the actual situation of the world out in the real world: our DM does not have telekinesis after all!):

1. With probability q, in her head she expects to be able to exchange, so she expects she will exchange. i.e. with probability q, the reference point is (0,1)

(a) With probability q, in the real world she actually is allowed to exchange, and since she wants to exchange, she will. That is, with probability q, the actual outcome is (0,1). Call this situation A

(b) With probability 1-q, in the real world she is not allowed to exchange, so even though she wants to exchange, she will not exchange. That is, with probability 1-q, the actual outcome is (1,0). Call this situation B

2. with probability (1-q), in her head she expect to not be able to exchange, so she expects not to exchange. i.e. with probability (1-q), the reference point is (1,0)

   (a) With probability q, in the real world she actually is allowed to exchange, and since she wants to exchange, she will. That is, with probability q, the actual outcome is (0,1). Call this situation C.

   (b) With probability 1-q, in the real world she is not allowed to exchange, so even though she wants to exchange, she will not exchange. That is, with probability 1-q, the actual outcome is (1,0). Call this situation D.

Here we have the complete picture of what might happen when our DM wants to exchange under the expectation of exchanging. Notice in each situation, we have a reference point and an actual outcome, which is exactly what we need to calculate reference dependent utility. So we have in essence listed out all the possible combinations of reference point and actual outcome, and each of these combinations has an associated probability. All we need to do to calculate the overall reference dependent utility is to find the reference dependent utility of each situation, multiply it by the probability of that situation happening, and add them together.

Again, let's look at each of these situations and find the reference dependent utility for them:

1. Situation A: the probability of this situation is $q \times q$, because it is the probability of the combination of expectation (0,1) and actual outcome (0,1), and since expectation in the head does not affect real outcome, they are independent. So the probability of getting this combination is the product of their probabilities (see Definition of Independence in the probability section). Here there is no loss, no gain. So the reference dependent utility is just the consumption utility $u((0,1))$

2. Situation B: the probability of this situation is $q \times (1-q)$. Here we gain 1 mug and lose 1 pen, so the reference dependent utility is U((2,0)—(0,1)). I will let you calculate this.

3. Situation C: the probability of this situation is $q \times q$, and we gain 1 pen and lose 1 mug. The reference dependent utility of this situation is $U(0,1)|(1,0))$.

4. Situation D: the probability of this situation is $q \times (1-q)$. Here no loss no gain, so the reference dependent utility of this situation is just its consumption utility u((1,0).

I will leave you to calculate the overall reference dependent utility of U(exchange—exchange) asked by the question. The blueprint is already laid out for you.

5.2.5 Part 5 and 6

Hopefully by now you understand how to tackle this sort of problem. I encourage you to try solving these parts yourself! Please come to my office hour if you run into any difficulties, I would be happy to go over them in detail again.

5.3 Example: Homework 2 Q1

Example 8 (Reference Dependence: HW2 Q1)**.** *The utility of consuming* c_1 *mugs and* c_2 *money is* $m(c_1, c_2) = c_1 + c_2$*, and the consumer expects to buy at price p with probability* $q \in [0, 1]$

5.3.1 Part i: utility from buying and not buying

The key in this problem is that the reference point is probabilistic, so there are two possible cases:

1. With q probability, she expects to buy. So with q probability the reference point is (1,-p).

2. With 1-q probability, she expects not to buy. So with 1-q probability the reference point is (0,0)

First, to calculate the utility of buying, we want

$U(buying|$expecting to buy with probability q).

This is because we are dealing with reference dependent utility, so when we say the utility of buying, we actually mean the utility of buying under the expectation of whatever our reference point is. Here the consumer has full autonomy, so despite the reference point (expectation) being probabilistic, the actual outcome is always going to be buying, i.e. (1,-p). So the consumption utility is $1 - p$. To get the gain/loss utility, we need to compare the actual outcome (1, -p)against the reference point (2 situations above). In situation 1, we gain nothing and lose nothing. In situation 2, we gain 1 mug and lose p money. So the gain/loss utility is $(1 - q)(\eta - \eta\lambda p$

So utility from buying is $1 - p + (1 - q)\eta(1 - \lambda p)$

Next, to calculate the utility of not buying, we want

U(not buying|expecting to buy with probability q).

The consumption utility is 0. We run through the same process as above to get the gain/loss utility for each situation, and that's how we get the answer.

5.3.2 Part ii

A decision maker is indifferent between A and B if U(A) = U(B).

So, she is indifferent between buying and not buying if

$U(buying|$expecting to buy with probability q) $= U$(not buying|expecting to buy with probability q)

(14)

We already have both sides from part i, the rest is algebra.

6 Time Preferences

6.1 Detour: Constrained Optimization

For problems involving time preferences, we need to know how to solve constrained optimization problems, i.e. find the best option given some constrains. For example, how many cakes do I eat each day over a 3 day period given that I only have 6 cakes available so that I can maximize my utility (I do not endorse eating 6 cakes over a 3 day period by the way :D).

In this course, we will only deal with problems involving one linear constraint. So we can write the general form of the constrained optimization problem as:

$$\max_{x_1,x_2,...,x_k} f(x_1, x_2, ..., x_k) \text{ subject to } x_1 + x_2 + ... + x_k = c \tag{15}$$

Here the function $f(x_1, x_2, ..., x_k)$ is called the **objective function**. It is the function we are trying to maximize, e.g. a utility function. The **solution** to the problem is the *values of the variables* $x_1, x_2, ..., x_k$ such that the objective function f is maximized (subject to the constraint $x_1 + x_2 + ... + x_k = c$). That's why we write $\max_{x_1,x_2,...,x_k}$: it means we are choosing $x_1, x_2, ..., x_k$ to maximize f. We usually write $x_1*, x_2*, ..., x_k*$ to denote the optimal solutions.

We use the Lagrange multiplier method to solve constrained optimization problems. We rewrite the problem using the Lagrangian:

$$\mathcal{L} = f(x_1, x_2, ..., x_k) + \lambda(c - x_1 - x_2 - ... - x_k) \tag{16}$$

Here the λ is called the Lagrange multiplier. We can solve the original problem by solving the transformed problem

$$\max_{x_1,x_2,...,x_k} \mathcal{L} = f(x_1, x_2, ..., x_k) + \lambda(c - x_1 - x_2 - ... - x_k) \tag{17}$$

With this transformed problem, it should be obvious what to do: we just take the first order conditions and set them to 0. The solutions to these first order conditions will be the solution to the problem.

First order conditions:

$$\frac{\partial \mathcal{L}}{\partial x_1} = \frac{\partial \mathcal{L}}{\partial x_2} = ... = \frac{\partial \mathcal{L}}{\partial x_k} = \frac{\partial \mathcal{L}}{\partial \lambda} = 0 \tag{18}$$

i.e.,

$$\begin{aligned} \frac{\partial f}{\partial x_1} - \lambda &= 0 \\ \frac{\partial f}{\partial x_2} - \lambda &= 0 \\ &... \\ \frac{\partial f}{\partial x_k} - \lambda &= 0 \\ c - x_1 - x_2 - ... - x_k &= 0 \end{aligned}$$

The set of FOC gives us a system of $k+1$ equations with $k+1$ variables ($x_1, x_2, ..., x_k, \lambda$) so there is a unique solution.

How do we solve for the system of equations? We can use the first k equations in the FOC to write the $x_1, x_2, ..., x_k$ variables as a function of λ, then substitute them into the linear constraint (last equation in the FOC). That lets us solve for λ. Once we have that we have all the x variables.

6.2 Classical Case: Exponential Discounting

Traditionally, in economics we assume people follow **exponential discounting** when considering inter-temporal choices, i.e. you discount future utility back to the present using a fixed discount rate δ, and therefore one dollar tomorrow is only worth δ dollar today, and one dollar the day after tomorrow is worth δ^2 dollars today, etc. The intuition for this is that delaying utility until tomorrow has a cost, for example, you can invest money and get interest tomorrow. So money today is worth more than the same amount of money tomorrow.

Formally, the preferences of an exponential discounter can be represented as

$$U_0 = u(c_0) + \sum_{t=1}^{T} \delta^t u(c_t) \tag{19}$$

Once this decision maker has made a plan for the future (at time 0, for example), if they are an exponential discounter, they will always follows through with their consumption plan in the future. We can this feature of preferences **time consistency**. Your preferences are time consistent if when tomorrow arrives, you do as you had planned for it today.

6.3 Present Bias

6.3.1 Overview

Obviously, many people do not have time consistent preferences in the real world. On Friday, I tell myself that I will exercise on Saturday by going for a bike ride to Winters. When Saturday arrives, I am too lazy to go out and end up eating banana bread. I don't follow through with my plan! My preferences here are time inconsistent.

We can model time-inconsistent preferences with a model of **present bias**. A DM that suffers from present bias follows **$\boldsymbol{\beta}-\boldsymbol{\delta}$ discounting**. The idea is that I am very impatient, so today, I discount my future utility by β in addition to the regular exponential discounting with discount factor δ. But I think tomorrow and the day after tomorrow and the day after that... I will always discount by a factor of δ. So essentially I am impatient today but think I will be patient tomorrow. But surprise! When tomorrow arrives I am impatient again so I discount the day after tomorrow by $\beta\delta$ again. So that's where the time inconsistency comes from.

Formally, a decision maker with $\beta\delta$ discounting can be represented by the following utility function:

$$U_0 = u(c_0) + \beta \sum_{t=1}^{T} \delta^t u(c_t) \tag{20}$$

With a present biased DM, we can look at 3 cases: the DM can be naive (they don't know they are present biased), sophisticated (they know they are present biased), and commitment (they know they are present biased and they tie their hands at the get go by committing to certain actions).

We will illustrate these cases with the following example:

Example 9 (Present Bias). *a) You have a cake and are trying to decide how to eat it over 3 days.* $u(c) = ln(c)$. *You have beta-delta preferences and log utility. What is your plan on Day 1? When arrive at Day 2, you re-optimize. What is your plan on Day 2? Do you follow through with your plan on Day 1? How much do you end up eating on Day 3?*

c) You are a beta-delta discounting consumer but you are sophisticated. You expect yourself to deviate from your plan on Day 2, so you solve the problem backwards. How much cake do you plan to eat each day?

6.3.2 Naivete

When the DM is naive, she first optimizes on Day 1. She solves the problem:

$$\max_{c_1,c_2,c_3} ln(c_1) + \beta\delta ln(c_2) + \beta\delta^2 ln(c_3) \text{ subject to } c_1 + c_2 + c_3 = 1 \tag{21}$$

The Lagrangian is

$$\mathcal{L} = ln(c_1) + \beta\delta ln(c_2) + \beta\delta^2 ln(c_3) + \lambda(1 - c_1 - c_2 - c_3) \tag{22}$$

Solving this problem gives us

$$c_1* = \frac{1}{1 + \beta\delta + \beta\delta^2} \tag{23}$$

$$c_2* = \frac{\beta\delta}{1 + \beta\delta + \beta\delta^2} \tag{24}$$

$$c_3* = \frac{\beta\delta^2}{1 + \beta\delta + \beta\delta^2} \tag{25}$$

On Day 2, she takes a look at how much cake is left over and re-optimizes. Let the leftover cake be M (we technically already know the exactly number $(1 - c_1*)$ from previously, but writing M is way more manageable than writing the whole big fraction).

Your Day 2 problem is

$$\max_{c_2,c_3} ln(c_2) + \beta\delta ln(c_3) \text{ subject to } c_2 + c_3 = M \tag{26}$$

Solving it we get:

$$c_2' = \frac{M}{1+\beta\delta}$$
$$c_3' = \frac{\beta\delta M}{1+\beta\delta}$$

Now since we know $M = 1 - c_1* = 1 - \frac{1}{1+\beta\delta+\beta\delta^2} = \frac{\beta\delta+\beta\delta^2}{1+\beta\delta+\beta\delta^2}$, substituting for M we have

$$c_2' = \frac{\beta\delta+\beta\delta^2}{(1+\beta\delta)(1+\beta\delta+\beta\delta^2)}$$
$$c_3' = \frac{\beta\delta(\beta\delta+\beta\delta^2)}{(1+\beta\delta+\beta\delta^2)(1+\beta\delta)}$$

So our DM does not follow through with her plan on Day 1! Further, we can see that $c_3' < c_3$ and $c_2' > c2$. Intuitively this makes sense, because on Day 2 the DM is impatient again, and she eats more cake than she originally planned for on Day 1, and on Day 3 she ends up with less cake than what she thought she would have.

6.3.3 Sophistication

A sophisticated DM knows that she is present biased and that she will get impatient again on Day 2. So she will first consider what she will do on Day 2, and given her actions on Day 2, she formulates a plan on Day 1 that takes into account her impatience on Day 2. That is, our DM solves the problem backwards.

Let's denote the leftover cake by M. Here we don't know what M is so it's a place-holder. Her day 2 self's problem is

$$\max_{c_2,c_3} ln(c_2) + \beta\delta ln(c_3) \text{ subject to } c_2 + c_3 = M \tag{27}$$

Solving it we get:

$$c_2* = \frac{M}{1+\beta\delta}$$
$$c_3* = \frac{\beta\delta M}{1+\beta\delta}$$

Her day 1 self's problem is therefore to optimize utility subject to what her day 2 self will do:

$$\max_{c_1,c_2,c_3} ln(c_1) + \beta\delta ln(\frac{M}{1+\beta\delta}) + \beta\delta^2 ln(\frac{\beta\delta M}{1+\beta\delta}) \text{ subject to } c_1 + M = 1 \tag{28}$$

So we can rewrite the problem as

$$\max_{c_1,c_2,c_3} ln(c_1) + \beta\delta ln(\frac{1-c_1}{1+\beta\delta}) + \beta\delta^2 ln(\frac{\beta\delta(1-c_1)}{1+\beta\delta}) \tag{29}$$

This is an unconstrained optimization of only one variable. So we can set FOC to 0 and get the solution for c_1*:

$$c_1* = \frac{1}{1+\beta\delta+\beta\delta^2} \tag{30}$$

Once we have c_1*, we can get

$$M = 1 - c_1* = \frac{\beta\delta+\beta\delta^2}{1+\beta\delta+\beta\delta^2} \tag{31}$$

So we have

$$c_2* = \frac{\beta\delta+\beta\delta^2}{(1+\beta\delta)(1+\beta\delta+\beta\delta^2)}$$
$$c_3* = \frac{\beta\delta(\beta\delta+\beta\delta^2)}{(1+\beta\delta+\beta\delta^2)(1+\beta\delta)}$$

Here the sophisticated agent's plan is time consistent.

Example 10. *Your local cinema theater offers a mediocre movie this week, a good movie next week, a great movie in two weeks, and a fantastic movie in three weeks. The utility of watching these movies are 4, 7, 10, 18 respectively. You can only watch one of these movies. Suppose $\delta = 1$.*

a) Which movie would you watch if you were an exponential discounting person? ($\beta = 1$)

b) What if you were naive and present biased? ($\beta = 0.5$)

c) What if you were sophisticated and present biased? ($\beta = 0.5$)

For a sophisticated agent, in week 2, you know you will watch in week 3, so you compare 7 with utility from watching week 3 movie, which is 5. So you will watch week 2. In week 1 is the same story - you know you will watch week two's movie if you get to it, so you compare 4 against 3.5. I leave the rest for you as practice :)

	Week 1	Week 2	Week 3	Week 4
	4	7	10	18
Exponential Discounting $\beta = 1$ $\delta = 1$	4	7	10	18
	N	N	N	Y
Naivete: $\beta = 0.5$ $\delta = 1$	4	7 vs 5 vs 9	10 vs 9	18
	N	N	Y	Y
Sophisticated: $\beta = 0.5$ $\delta = 1$	4 vs 3.5	7 vs 5	10 vs 9	18
	Y	Y	Y	Y

6.3.4 Commitment

In the commitment case, the DM commits to a plan of action on Day 1 by removing the possibility of falling into temptation and deviating from her plan. To discuss commitment, we need to discuss preferences for menus, because committing to a plan of action is essentially forcing yourself to choose from a menu and removing menus with temptations. This leads us to a discussion of Gul-Pesendofer preferences

6.4 Preferences over Menus and the Gul-Pesendorfer Model

6.4.1 The Model

Under the Gul-Pesendorfer Model, the utility of menu A is

$$U(A) = \max_{p \in A}[u(p) + v(p)] - \max_{q \in A} v(q) \tag{32}$$

Intuition:

- Find the p gives you the maximum combined commitment utility and temptation utility, so you will pick that out of the menu
- But you have to avoid the maximum temptation from the things in the menu, so that's a cost v(q). q is the item in the menu that gives you the most temptation utility

- When the menu only has one item, the utility of the menu is the commitment utility of the item

e.g. I really like mochi and donuts (who doesnt?!). Now I am planning to get boba with a friend. I also know that I am weak, so if I go to T% for boba, I will fall into the temptation of Mochinut. So I will go to sharetea instead. I am taking away the option of Mochinut from myself by committing to the menu that does not have mochi donuts.

6.4.2 An Example

3)

Object	u	v
Salad	9	0
Fish	5	3
Burger	3	9

For the utility gives above, what is the relation between $\{s\}, \{f\}, \{s, f\}$ under Gul-Pesendorfer preferences? What is the relation between $\{s\}, \{b\}, \{s, b\}$?

Object	u	v	u+v
Salad	9	0	9
Fish	5	3	8
Burger	3	9	12

$U(\{s\}) = u(s) = 9$

$U(\{f\}) = u(f) = 5$

$U(s, f) = \max_{p\in\{s,f\}}[u(p) + v(p)] - \max_{q\in\{s,f\}} v(q) = u(s) + v(s) - v(f) = 9 - 3 = 6$

So $\{s\} \succ \{s, f\} \succ \{f\}$

$U(\{b\}) = u(b) = 3$

$U(s, b) = \max_{p\in\{s,b\}}[u(p) + v(p)] - \max_{q\in\{s,b\}} v(q) = u(b) + v(b) - v(b) = 3$

So $\{s\} \succ \{s, b\} \sim \{b\}$

7 Bounded Rationality and Sequential Search

7.1 Sequential Search Model

7.1.1 General Case

To successfully understand all the math behind this model, you need to understand the properties of probability density functions of continuous random variables, as well as the properties of integrals and how to evaluate them.

Setup:

- Choice set A containing M items
- Utility function u(x) for $x \in A$
- Belief about the value of each option ex ante is a probability distribution f
- Search cost k for each alternative you look at

Lets say you have already looked at n items ($n < M$), and the best thing you have seen among these n items has the utility $\bar{u}$. To decide if we want to continue searching, we need to compare the utility of stopping and keep searching.

The utility of stopping now is $\bar{u} - nk$.

The (expected) utility of looking at the next alternative is:

utility of the next item conditional on getting something better *(expected utility since we don't know what we will get)* + P(getting something worse than or equal to $\bar{u}$) $\times \bar{u}$ *(this is the expected utility of getting something worse)*

Writing it in math, we have the utility of searching is

$$\begin{aligned}
&P(u \leq \bar{u}) \times \bar{u} + E[u|u \geq \bar{u}] - (n+1)k \\
&= \int_{-\infty}^{\bar{u}} f(u)du \times \bar{u} + \int_{\bar{u}}^{\infty} uf(u)du - (n+1)k \\
&= \int_{-\infty}^{\bar{u}} \bar{u}f(u)du + \int_{\bar{u}}^{\infty} uf(u)du - (n+1)k
\end{aligned}$$

We should keep searching if the utility from searching is higher than the utility from stopping, that is,

$$\int_{-\infty}^{\bar{u}} \bar{u}f(u)du + \int_{\bar{u}}^{\infty} uf(u)du - (n+1)k \geq \bar{u} - nk \tag{33}$$

which gives us

$$\int_{-\infty}^{\bar{u}} \bar{u}f(u)du + \int_{\bar{u}}^{\infty} uf(u)du - k \geq \bar{u} \tag{34}$$

Notice that **the strategy does not depend on n!** So the strategy for each period is exactly the same. And solving for $\bar{u}$ using this equation with equality $\int_{-\infty}^{\bar{u}} \bar{u}f(u)du + \int_{\bar{u}}^{\infty} uf(u)du - k = \bar{u}$, we can get the threshold level of u*.

Now since by definition we know

$$1 = \int_{-\infty}^{\infty} f(u)du = \int_{-\infty}^{\bar{u}} f(u)du + \int_{\bar{u}}^{\infty} f(u)du \tag{35}$$

We can rewrite the previous equation as

$$\int_{-\infty}^{\bar{u}} \bar{u}f(u)du + \int_{\bar{u}}^{\infty} uf(u)du - k \geq \bar{u} \times [\int_{-\infty}^{\bar{u}} f(u)du + \int_{\bar{u}}^{\infty} f(u)du] \tag{36}$$

$\iff$

$$\int_{-\infty}^{\bar{u}} \bar{u}f(u)du + \int_{\bar{u}}^{\infty} uf(u)du - k \geq \int_{-\infty}^{\bar{u}} \bar{u}f(u)du + \int_{\bar{u}}^{\infty} \bar{u}f(u)du \tag{37}$$

So after cancelling on both sides we have keep searching is better if

$$k \leq \int_{\bar{u}}^{\infty} (u - \bar{u})f(u)du \equiv F(\bar{u}) \tag{38}$$

Here $F(\bar{u})$ is decreasing in $\bar{u}$ which means that the higher your search cost is, the lower is the threshold of previous best option that's required to stop searching

Equipped with this knowledge of what's going on in the background of the model, you should be able to handle any problem being thrown your way!

Example 11. *Suppose the utility of items is distributed uniformly randomly over [1,2]. Search cost is k=0.5. You have already searched M-1 out of M items, and the best out of those M-1 items give a utility of $\bar{u}$. What is the threshold value of baru, that makes you indifferent between searching once more and not searching? What would be the value of $\bar{u}$ if k= 0.1 instead?*

I am indifferent if

$$\int_{-\infty}^{\bar{u}} \bar{u} f(u)du + \int_{\bar{u}}^{\infty} uf(u)du - Mk = \bar{u} - (M-1)k \tag{39}$$

Here, f is a uniform distribution, so $f(u) = 1$ if $u \in [1, 2]$ and 0 otherwise. So the equation becomes

$$\int_{1}^{\bar{u}} \bar{u} du + \int_{\bar{u}}^{2} udu - Mk = \bar{u} - (M-1)k \tag{40}$$

The left hand side becomes

$$\begin{aligned}
&\int_{1}^{\bar{u}} \bar{u} du + \int_{\bar{u}}^{2} udu - Mk \\
&= [\bar{u}u]_{u=1}^{u=\bar{u}} + [\frac{u^2}{u=2}]_{u=\bar{u}}^{2} - Mk \\
&= \bar{u}(\bar{u}-1) + \frac{2^2}{2} - \frac{(\bar{u})^2}{2} - Mk
\end{aligned}$$

You can simplify, set it equal to the right hand side, plug in the value for k, and solve for $\bar{u}$ using the formula for the roots of the quadratic equation. Remember that only the root between 1 and 2 is the one we want because u needs to be inside the interval $[1, 2]$! I leave the rest as an exercise for you :)

8 Bayes' Rule

For a relatively accessible and fun introduction to Bayesian reasoning, please see Bayes_Pinker.pdf in the handouts folder on Canvas. It is a chapter from the book Rationality by Steven Pinker.

It is a rule that tells you how to updated your beliefs given information. For example, you start with a belief on the probability of event A happening, and you receive a signal S about event A, then your updated belief on A happening is:

$$P(A|S) = \frac{P(S|A)P(A)}{P(S)} \tag{41}$$

It can be derived from the formula for conditional probabilities:

$$P(A|S) = \frac{P(A \cap S)}{P(S)} \tag{42}$$

and

$$P(S|A) = \frac{P(A \cap S)}{P(A)} \tag{43}$$

We call the probability of the event before receiving information the "prior", and the probability after receiving information the "posterior". We call $P(S|A)$ the accuracy of the information. The more accurate the information, the more likely it is that you receive it given that the event A actually happened.

Notice that the posterior probability is increasing in the prior and also the accuracy of the information $P(S|A)$.

8.1 Taxicab Problem and Base Rate Neglect

Example 12. *A cab was involved in a hit and run accident at night. Two cab companies, the Green and the Blue, operate in the city. 85% of the cabs in the city are Green and 15% are Blue.*

A witness identified the cab as Blue. The court tested the reliability of the witness under the same circumstances that existed on the night of the accident and concluded that the witness correctly identified each one of the two colors 80% of the time and failed 20% of the time.

What is the probability that the cab involved in the accident was Blue rather than Green knowing that this witness identified it as Blue?

A lot of people will say something above 50%, some even 80% which is incorrect. The problem is that there are very few blue cabs to start with, so even with the information that witness saw blue, the posterior probability still won't be very high. Saying it is 80% probability is throwing the prior (base rate) out of the window completely - this is a classic case of **base rate neglect**. Why is it not 80%? This is because 80% is $P(Signal|Event)$ whereas we are interested in $P(Event|Signal)$. These two are not the same! Think about this: whenever it rains I always carry an umbrella. You see me carrying an umbrella today, what is the probability that it rains? It is not 100%!! Because I could just be carrying an umbrella everyday, then me carrying an umbrella doesn't necessarily mean that it rains. This is because we know $P(umbrella|rain) = 1$, but our event of interest is $P(rain|umbrella)$, and these two are not the same! Knowing that I always carry an umbrella whenever it rains does not tell you anything about whether I carry an umbrella or not when it's not raining.

Here, we are interested in $P(Cab\ is\ blue|Witness\ saw\ blue)$. Denote the event that the cab is actually blue by B, and witness saw a blue cab by b, the cab being green by

G, and witness saw green by g. By Bayes' Rule, we have

$$\begin{aligned} P(B|b) &= \frac{P(b|B)P(B)}{P(b)} \\ &= \frac{P(b|B)P(B)}{P(b|B)P(B) + P(b|G)P(G)} \end{aligned}$$

We used the Law of Total Probability to calculate the denominator (the *unconditional* probability of the signal), because Green and Blue form a partition of the sample space.

So that gives us

$$\begin{aligned} P(B|b) &= \frac{P(b|B)P(B)}{P(b)} \\ &= \frac{P(b|B)P(B)}{P(b|B)P(B) + P(b|G)P(G)} \\ &= \frac{0.8 \times 0.15}{0.8 \times 0.15 + 0.2 \times 0.85} \\ &= \frac{12}{29} \\ &\simeq 41\% \end{aligned}$$

Now let's look at a frequency diagram for intuition. Let's say there are 400 total cabs, and 60 of them are blue, and 340 of them are green.

What is the probability that the cab is actually blue when the witness identified it as blue? There are two cases:

1. The witness identified it as blue, and the cab is actually blue

2. The witness identified it as blue, but the cab is actually green

For the first case, the witness can correct identify 80% of all blue cabs, so that's 48 blue cabs. For the second case, that's around 68 green cabs the witness identified as blue. So out of the $48 + 68 = 116$ cabs the witness identified as blue, only 48 are actually blue, which gives us $48/116 = \frac{12}{29} \simeq 41\%$

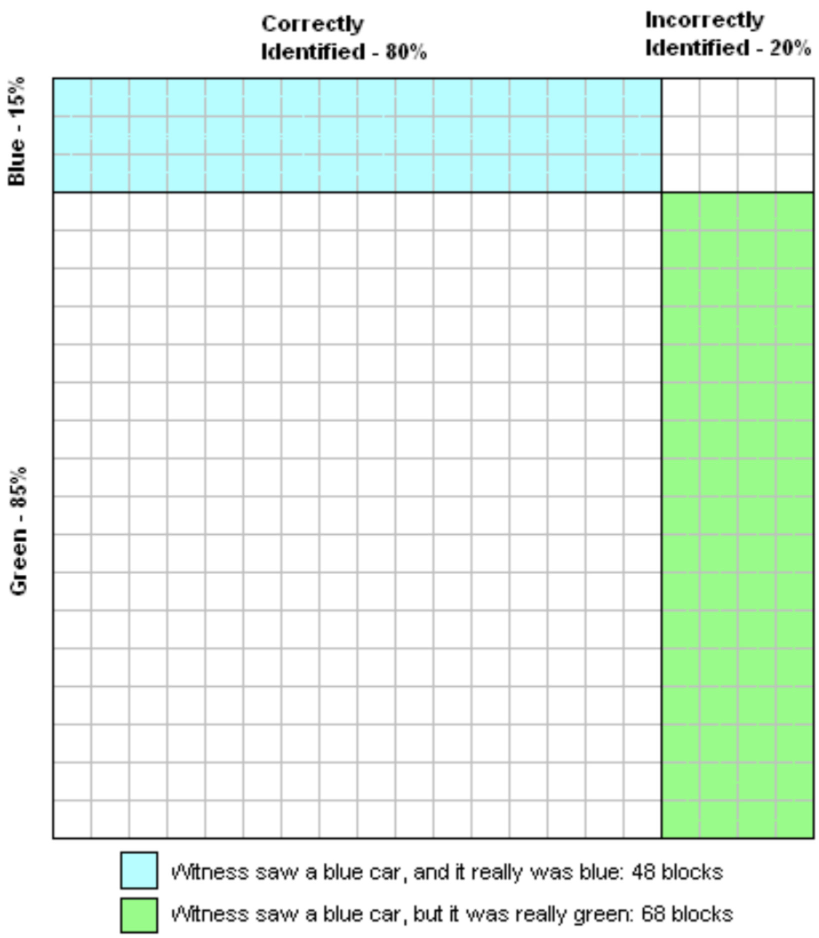

The key to solving Bayes' Rule problem is to find the event of interest, and to find the signal accuracy.

8.2 Monty Hall Problem: A Counter-Intuitive Example

Example 13. *Suppose you're on a game show, and you're given the choice of three doors: Behind one door is a car; behind the others, rocks. You pick a door, say No. 1, and the host, who knows what's behind the doors, opens another door, say No. 2, which has a rock. He then says to you, "Do you want to pick door No. 3?" Is it to your advantage to switch your choice?*

We need to calculate $P(car = 1|open = 2)$ and $P(car = 3|open = 2)$.

We know $P(car = 1|open = 2) = \dfrac{P(open = 2|car = 1)P(car = 1)}{P(open = 2)}$ and $P(car =$

$3|open = 2) = \dfrac{P(open = 2|car = 3)P(car = 3)}{P(open = 2)}$.

Now we know $P(open = 2|car = 1) = 1/2$, $P(open = 2|car = 3) = 1$, and $P(open = 2) = P(open = 2|car = 1)P(car = 1) + P(open = 2|car = 3)P(car = 3) = \frac{1}{2} \times \frac{1}{3} + 1 \times \frac{1}{3} = \frac{1}{2}$ so

$$P(car = 1|open = 2) = \frac{1/6}{1/2} = 1/3 \tag{44}$$

and

$$P(car = 3|open = 2) = \frac{1/3}{1/2} = 2/3 \tag{45}$$

So you should switch! This seems counter-intuitive, but the key thing here is that you are actually getting information about where the car is by the host revealing a door with a rock. This is because the host knows where the cars and rocks are, and cannot open the door to the car, so this action carries information content.

One way to think about the intuition behind this problem is to think about the possible cases: there are 3 possible cases, car behind door 1, door 2, and door 3. Let's list them:

1. C, R, R
2. R, C, R
3. R, R, C

Now, you chose door 1, and the host revealed a rock behind another door. If you switch, you win in case 2 and 3, but if you don't switch, you only win in case 1! So you should switch.

8.3 A Practice Problem

Christina comes to section with a boba. You are craving some boba right now and the one Christina has looks soooooo good, so you want to figure out where she got it from. Unfortunately, Christina is suffering from fuzzy memory because she is so exhausted from working so hard. The only thing she can tell you is that she got it from either ShareTea or Teaspoon. Fortunately, all is not lost and you have the following information at your disposal:

- You know Christina likes ShareTea and Teaspoon equally

- Christina likes milk tea at ShareTea better and fruit tea at Teaspoon better. So she has a 70% probability of getting a fruit tea if she goes to Teaspoon, and has a 40% chance of getting a fruit tea if she goes to ShareTea. Assume that milk tea and fruit tea are the only types of boba.

- You notice Christina has a fruit tea in her hand.

Given the above information, what do you think is the probability that Christina went to ShareTea?

9 BCGS Stereotype Model

In the BCGS stereotype model, the key idea is that the representativeness of a characteristic in a group depends on its relative frequency against the comparison group. The higher the relative frequency, the more representative it is.

9.1 The Model

Define the group of interest G, comparison group $-G$, with types $t \in \{t_1, t_2, ..., t_k\}$. The representativeness of attribute t in group G with respect to comparison group $-G$ is formally defined as:

$$R(t, G, -G) = \frac{P(t|G)}{P(t|-G)} \tag{46}$$

That is, the representativeness is the ratio of the conditional likelihood of type t in group G against that of t in comparison group $-G$. Notice that even though an attribute could be the most likely in its group G, it still might not be the most representative attribute, because representativeness depends on the conditional likelihood of this attribute in the comparison group $-G$.

Take the example of the stereotype that there are a lot of old people in Florida. Below is the distribution of people in different age ranges in Florida vs. the rest of the country.

P(t|G)		Types (t)			
		0-19	20-44	45-64	65+
Groups	Florida	27%	27%	23%	23%
	Rest	28%	29%	22%	21%

Here the likelihood of 65+ people in Florida is $P(65 + |Florida) = 23\%$, and the likelihood of 65+ people in the rest of the country is $P(65 + |Rest) = 32\%$, so the representativeness of 65+ people in Florida is

$$R(65+, Florida, Rest) = \frac{P(65 + |Florida)}{P(65 + |Rest)} = \frac{23\%}{21\%} \simeq 1.10 \tag{47}$$

Now the representativeness of the rest of the age groups is:

$$R(0-19, F, R) = \frac{27\%}{28\%} \simeq 0.96 \tag{48}$$

$$R(20-44, F, R) = \frac{27\%}{29\%} \simeq 0.93 \tag{49}$$

$$R(45-64, F, R) = \frac{23\%}{22\%} \simeq 1.05 \tag{50}$$

So as we can see, despite 65+ people not being the most likely within Florida, it is the most representative age group for Florida!

9.2 Rank-Based Truncation

Since highly representative attributes are easier to recall and easier to come to mind, when trying to form a subjective assessment of type distribution, the decision maker distorts the probabilities because it's hard to recall non-representative attributes. One way to describe this is by **rank-based truncation**.

When the DM engages in rank-based truncation, they can only recall attributes that have representativeness rankings of at most d. And they form subjective probabilities as if the recalled attributes are all that there are. This gives rise to probability distortions.

For example, let's say $d = 2$, so the DM can only recall the top 2 most representative characteristics. In our Florida example, that means the DM can only recall 65+ and 45-64 age groups. Then, the DM acts as if there are only these two types, and the stereotyped probabilities are

$$\pi^{st}_{65+,F} = \frac{\pi_{65+,F}}{\pi_{65+,F} + \pi_{45-64,F}} = \frac{23\%}{23\% + 23\%} = 0.5 \tag{51}$$

$$\pi^{st}_{45-64,F} = \frac{\pi_{45-64,F}}{\pi_{65+,F} + \pi_{45-64,F}} = \frac{23\%}{23\% + 23\%} = 0.5 \tag{52}$$

So clearly, our DM has a stereotype that there are lots of 65+ people in Florida!

9.3 Example: HW4 Q3

3) Based on Coffman et al: As an outside observer, you know the following joint distribution over (G, t). Groups are Florida and Rest of USA. Types are age categories.

$P(t, G)$		Types (t)			
		0-19	20-44	45-64	65+
Groups (G)	Florida	.02	.02	.02	.04
	Rest	.26	.24	.20	.20

Suppose the DM only recalls types that have representativeness rankings of at most $d = 3$. What would be the stereotyped probabilities $\pi^{st}_{t,G}$ that the DM would recall?

First, we need to find the conditional distribution of types within Florida and Rest of the country.

In Florida,

$$P(0-19|F) = \frac{0.02}{0.02+0.02+0.02+0.04} = 0.2 \tag{53}$$

$$P(20-44|F) = \frac{0.02}{0.02+0.02+0.02+0.04} = 0.2 \tag{54}$$

$$P(45-64|F) = \frac{0.02}{0.02+0.02+0.02+0.04} = 0.2 \tag{55}$$

$$P(65+|F) = \frac{0.04}{0.02+0.02+0.02+0.04} = 0.4 \tag{56}$$

In the rest of the country,

$$P(0-19|R) = \frac{0.26}{0.26+0.24+0.20+0.20} \simeq 0.29 \tag{57}$$

$$P(20-44|R) = \frac{0.24}{0.26+0.24+0.20+0.20} \simeq 0.27 \tag{58}$$

$$P(45-64|R) = \frac{0.20}{0.26+0.24+0.20+0.20} \simeq 0.22 \tag{59}$$

$$P(65+|R) = \frac{0.20}{0.26+0.24+0.20+0.20} \simeq 0.22 \tag{60}$$

So the representativeness of each type in Florida is:

$$R(0-19, F, R) = \frac{P(0-19|F)}{P(0-19|R)} = \frac{0.2}{0.29} \simeq 0.69 \tag{61}$$

$$R(20-44, F, R) = \frac{P(20-44|F)}{P(20-44|R)} = \frac{0.2}{0.27} \simeq 0.75 \tag{62}$$

$$R(45-64, F, R) = \frac{P(45-64|F)}{P(45-64|R)} = \frac{0.2}{0.22} \simeq 0.9 \tag{63}$$

$$R(65+, F, R) = \frac{P(65+|F)}{P(65+|R)} = \frac{0.4}{0.22} \simeq 1.8 \tag{64}$$

The DM only recalls age groups 65+, 45-64, and 20-44. I leave the stereotyped probabilities as an exercise for you.