Recitation 1: Utility Maximization

Maddie McKelway & Will Rafey

Overview

1. Utility Maximization: The Basics
2. Utility Maximization over Two Goods
3. Utility Maximization over Two Periods
4. Utility Maximization over Three Periods

Utility and Diminishing Marginal Utility

- Utility: the satisfaction from consuming a good or service
  - Utility function $u : X \to \mathbb{R}$.
- Marginal utility: the additional satisfaction from consuming one more unit of the good or service
- Law of diminishing marginal utility
  - The more you consume, the less utility you get from the additional unit.

The Budget Constraint and Utility Maximization

- We live in a scarce world; we face constraints on what and how much goods and services we can have
- Economics assumes that people maximize their utility functions subject to their constraints
- Today we will review utility maximization in traditional economic theory
- Behavioral economics considers whether these models are realistic and, if not, how they can be extended to be more realistic

Start with a Model of Two Goods

Suppose you are trapped on an island. There are only two kinds of plants that can be planted on the island: oranges and potatoes. The island has a cultivated area of 4 acres. Each acre can produce 1 unit of oranges *or* 1 unit of potatoes. How should you allocate the land between oranges and potatoes?

- Need a measure to compare different combinations of oranges and potatoes - use a utility function!
- Assume your utility function is

$$u(o, p) = \ln(o) + 2\ln(p)$$

Model of Two Goods (Continued)

- The constraints can be constructed from the information provided:

$$o + p \leq 4$$

  as well as $o, p \in [0, 4]$
  - In this example, the prices of oranges and potatoes are the same – both require 1 acre of land for 1 unit of output
  - Can drop $o, p \in [0, 4]$
    - Lower bound implied by log utility and upper bound implied by $o + p \leq 4$
- The problem becomes

$$\begin{aligned} &\max_{o,p} \ln(o) + 2\ln(p) \\ s.t. \quad & o + p \leq 4 \end{aligned}$$

Solving the Math

- There are many ways to solve constrained maximization problems
- A common method used in economics is the Lagrangian method
- Another is to equate the ratios of marginal utilities to prices
  - If (x^*, y^*) is an interior solution to the maximization problem

$$\max_{x,y} u(x,y)$$
$$s.t. \qquad p_1 x + p_2 y \leq w$$

  then

$$\frac{MU_x}{p_1}\Big|_{(x^*,y^*)} = \frac{MU_y}{p_2}\Big|_{(x^*,y^*)}$$

  - Combining this with the requirement that the solution lie on the budget constraint gives (x^*, y^*)
    - "Don't leave money on the table"
    - i.e. $p_1 x^* + p_2 y^* = w$

Graphical Interpretation

Image removed due to copyright restrictions.
View Fig. 3.6 Utility Maximization.

Solving the Math in Our Example

- Equating the ratios of marginal utilities to prices gives

$$\frac{\partial u}{\partial o} = \frac{\partial u}{\partial p}$$

$$\frac{1}{o^*} = \frac{2}{p^*}$$

- Assuming the solution lies on the budget constraint gives

$$o^* + p^* = 4$$

- Combining gives the solution

$$o^* = \frac{4}{3}, \; p^* = \frac{8}{3}$$

Two Periods

Now suppose there are two periods and no oranges.

You begin period 1 with 4 units of potatoes. You cannot grow any more and you have no other source of food for the two periods.

This means that in period 1 you have to save food for period 2. You can store the potatoes in a basket between the periods, but in period 2, only 80% of the saved potatoes will remain (the rest will be eaten by mice!).

Discounting Future Consumption

- Two goods become consumption in period 1 (c_1) and consumption in period 2 (c_2)
- Now the utility function will include temporal discounting
  - Why? People may not value current and future consumption the same
- Utility becomes

$$u(c_1, c_2) = \ln(c_1) + \delta \ln(c_2)$$

  - δ is called the "discount factor" and captures intertemporal preferences
  - Generally we assume $\delta \leq 1$
  - Larger $\delta \Rightarrow$ more patient

Solving the Math

- The constraints are

$$c_2 \leq 0.8(4 - c_1)$$

  as well as $c_1 \in [0, 4]$, $c_2 \in [0, 3.2]$
  - $c_1 \in [0, 4]$, $c_2 \in [0, 3.2]$ again implied
  - We rewrite constraint as $0.8c_1 + c_2 \leq 3.2$
    - As if we have prices $(p_1, p_2) = (0.8, 1)$

- Equating the ratios of marginal utilities to prices gives

$$\frac{\partial u / \partial c_1}{p_1} = \frac{\partial u / \partial c_2}{p_2}$$

$$\frac{1}{c_1^*} = \frac{0.8\delta}{c_2^*}$$

Solving the Math (Continued)

- Combining with $c_2^* = 0.8(4 - c_1^*)$ gives

$$c_1^* = \frac{4}{1+\delta}, \;\; c_2^* = \frac{3.2\delta}{1+\delta}$$

- Comparative statics
  - More patient (i.e., higher δ) $\implies$ more c_2, less c_1

Three Periods

Bad news! Your Amazon Prime membership has lapsed and now you must rely on potatoes for three periods.

From period 2 to period 3, the mice will again eat 20% of the remaining potato stock.

Now at period 1, you also value period 3 consumption (c_3) but value it even less than you do period 2 consumption.

How should you allocate your consumption across periods?

Exponential Discounting

- Paul Samuelson (MIT) proposed using the same discount factor on future utility from each period to the next

$$U(c_1, c_2, ..., c_T) = u(c_1) + \delta u(c_2) + \delta^2 u(c_3) + ... + \delta^{T-1} u(c_T)$$

  - Here $u(c_t)$ is the per-period utility, and $U(\cdot)$ specifies how people value consumption into the future at $t = 1$

- In a three-period model, we consider

$$U(c_1, c_2, c_3) = \ln(c_1) + \delta \ln(c_2) + \delta^2 \ln(c_3)$$

- Is this a realistic assumption?

Solving the Math

- The problem becomes

$$\begin{aligned} &\max_{c_1,c_2,c_3\geq 0} \ln(c_1)+\delta\ln(c_2)+\delta^2\ln(c_3) \\ s.t. \quad (i) \quad & c_2 \leq 0.8(4-c_1) \\ (ii) \quad & c_3 \leq 0.8[0.8(4-c_1)-c_2] \end{aligned}$$

  as well as $c_t \in [0, 0.8^{t-1}\times 4]$

- $c_t \in [0, 0.8^{t-1}\times 4]$ is implied. Since $c_3 \geq 0$, (ii) implies (i). So we can use just (ii) and can rewrite it as

$$0.64c_1 + 0.8c_2 + c_3 \leq 2.56$$

Solving the Math

- The interior solution (c_1^*, c_2^*, c_3^*) satisfies

$$\frac{\frac{\partial U(c_1,c_2,c_3)}{\partial c_1}}{0.64} = \frac{\frac{\partial U(c_1,c_2,c_3)}{\partial c_2}}{0.8} = \frac{\frac{\partial U(c_1,c_2,c_3)}{\partial c_3}}{1}$$

$$\frac{1}{c_1^*} = \frac{0.8\delta}{c_2^*} = \frac{0.64\delta^2}{c_3^*}$$

- Combining with $0.64c_1^* + 0.8c_2^* + c_3^* = 2.56$, we get

$$c_1^* = \frac{4}{1+\delta+\delta^2},\ c_2^* = \frac{3.2\delta}{1+\delta+\delta^2},\ c_3^* = \frac{2.56\delta^2}{1+\delta+\delta^2}$$

- Note that the ratio of consumption across periods is the same! (As long as per-period utility and price ratios are the same.) This is the essence of exponential discounting.

Time Consistency

- Suppose you follow the optimal allocation plan and consume $c_1^* = \frac{4}{1+\delta+\delta^2}$ in period 1. Then in period 2, will you deviate from consuming $c_2^* = \frac{3.2\delta}{1+\delta+\delta^2}$?

- In period 2, there remains $0.8\frac{4(\delta+\delta^2)}{1+\delta+\delta^2}$ potatoes. Now you are facing a 2-period problem. As shown above, in a 2-period problem, you will consume fraction $\frac{1}{1+\delta}$ of the total in the first period and leave the rest for the second period.

$$\frac{1}{1+\delta} \cdot \frac{0.8 \cdot 4(\delta+\delta^2)}{1+\delta+\delta^2} = \frac{3.2\delta}{1+\delta+\delta^2}$$

- This is exactly c_2^*

- Is this a coincidence? No! Exponential discounting assumes the same discount factor between every future period to the next.

Further References

- *Microeconomics* by Pindyck and Rubinfeld
- 14.03 MIT OpenCourseware: see notes on class website

Recitation 2: Exponential vs. Quasi-Hyperbolic Discounting

Maddie McKelway & Will Rafey

Exponential Discounting Model

$$U_t \equiv \sum_{\tau=t}^{\infty} \delta^{\tau-t} u_\tau = u_t + \delta u_{t+1} + \delta^2 u_{t+2} + \delta^3 u_{t+3} + \dots$$

- What is the key assumption of this model?
    - Amount of patience between any two periods the same
- What does this assumption imply?
    - Same degree of patience in the short- and long-run
    - Time consistency
    - No demand for commitment
- Does this seem realistic?

Exponential discounting: calibration

- Assume exponential discounting and linear utility of consumption.
- A student is indifferent between \$100 today and \$120 in two weeks.
  - What is δ? $5/6$ for two weeks.

$$100 = \frac{5}{6} \cdot 120$$

  - So the student discounts one month by $(5/6)^2$.
  - Discounts one year by $(5/6)^{24}$.
- Implies indifference between \$100 today and \$7949.68 in one year!

$$100 = \left(\frac{5}{6}\right)^{24} \cdot 7949.68$$

Exponential discounting: calibration

- Assume exponential discounting and linear utility of consumption.
- Suppose $\delta = 0.9$ (over one month).
- Pick between \$50 today and \$100 in two months.
  - Will pick \$100 in two months. $100 \cdot 0.9^2 = 81 > 50$.
- Suppose $\delta = 0.7$.
- Pick between \$50 today and \$100 in two months.
  - Will pick \$50 today. $100 \cdot 0.7^2 = 49 < 50$.

Evidence against the Exponential Discounting Model

- Short-run impatience and long-run patience
- Time inconsistency
- Demand for commitment

Evidence against the Exponential Discounting Model

- **Short-run impatience and long-run patience**
- Time inconsistency
- Demand for commitment

What Does Patience being Constant over Time Mean?

- Question 1: would you like to
  (a) eat one piece of candy now, or
  (b) eat two pieces of candy in an hour?
- Question 2: would you like to
  (a) eat one piece of candy in a week, or
  (b) eat two pieces of candy in a week and an hour?
- Patience being constant over time means you'd either choose (a) for both or (b) for both
- Bonus question: why do the (a) options have one piece and the (b) options have two pieces?
  - The exponential discounting world does allow for impatience (i.e. $\delta < 1$)
- Lots of evidence of short-run impatience and long-run patience
  - which implies many individuals would choose (a) for question 1 and (b) for question 2

Frederick et al. (2002): Estimated δ increases by time horizon

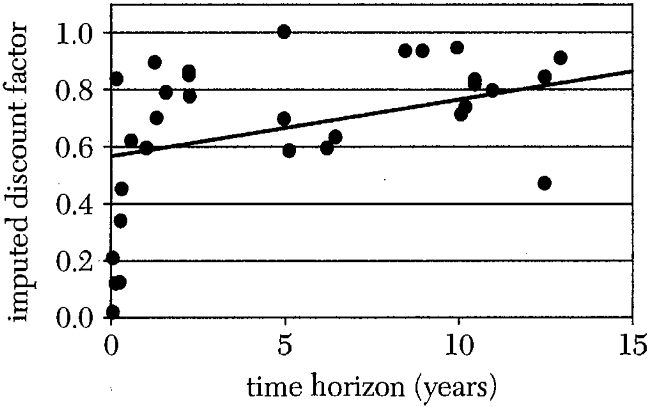

Figure: Frederick et al. (2002), Figure 1a

Evidence against the Exponential Discounting Model

- Short-run impatience and long-run patience
- **Time inconsistency**
- Demand for commitment

Time Consistency

- Time consistency (or dynamic consistency) = the action a person thinks she should take in the future always coincides with the action that she actually prefers to take once the time comes

- Time consistency an implication of the exponential discounting model
  - Consider the choice between two actions in period 1, A and B
  - At time $t = 0$, the individual prefers action A over B if and only if

$$u_0 + \delta u_1(A) + \delta^2 u_2(A) + \ldots \geq u_0 + \delta u_1(B) + \delta^2 u_2(B) + \ldots$$

  - Subtracting u_0 and dividing by δ gives

$$u_1(A) + \delta u_2(A) + \ldots \geq u_1(B) + \delta u_2(B) + \ldots$$

    which means the individual prefers A over B at time $t = 1$
  - That is, in the exponential discounting model, preferring A over B at $t = 0$ implies the individual will choose A over B at $t = 1$
    - i.e. the individual is time consistent

- Is time consistency realistic? Can you think of examples of time inconsistency?

Evidence against the Exponential Discounting Model

- Short-run impatience and long-run patience
- Time inconsistency
- **Demand for commitment**

Demand for Commitment

- Commitment device = a choice an individual makes in the present which restricts his set of choices in the future
- In the exponential discounting model, would the individual want a commitment device?
  - No. In this model, choices are time consistent so the future self will make whatever decision the present self prefers, whether or not choices are restricted.
- Can you think of examples of people demanding commitment devices?

Evidence against the Exponential Discounting Model

- Short-run impatience and long-run patience
- Time inconsistency
- Demand for commitment

Quasi-Hyperbolic Discounting Model

At time t, the person aims to maximize

$$u_t + \beta\delta u_{t+1} + \beta\delta^2 u_{t+2} + \beta\delta^3 u_{t+3} + \ldots,$$

- What's the key difference between this model and the exponential discounting model?
  - β, the short-term discount factor
  - β relaxes the assumption that the amount of patience between any two periods is the same; it allows for more impatience between today and tomorrow than between 7 and 8 days from now
- Why is the quasi-hyperbolic discounting model a better fit, at least in some situations?
  - Its two parameters allow for short-run impatience and long-run patience
  - It predicts time-inconsistent behavior and demand for commitment

Quasi-hyperbolic discounting

Algorithm

Utility is given for each t by

$$U_t = \delta^{t-1} u_t(x_t) + \beta \sum_{s \geq t}^{T} \delta^{s-1} u_s(x_s). \tag{1}$$

The algorithm to solve the optimal plan $(x_t^*)_{t=1}^{T}$ is by **backwards induction**.

1. Determine $x_T^*(\cdot)$, a function of $(x_s)_{s<T}$.
   - first, calculate payoffs for each possible choice of x_T, given $(x_s)_{s<T}$
   - second, choose the best choice; this is the function x_T^*
2. Then use this information to determine $x_{T-1}^*(\cdot)$, as function of $(x_s)_{s<T-1}$.
3. Continue until you reach $t = 1$.

Example

Actions $x_t \in \{0, 1\}$. Payoffs

$$u_t(x) = \begin{cases} 0 & \text{if } x_t = 0 \text{ and } t < T \\ -\theta_t & \text{if } x_t = 1 \\ -\infty & \text{if } t = T \text{ and } x_s = 0 \text{ for all } s \leq T. \end{cases} \tag{2}$$

I.e., at T, if you have not done x_t, you must do it!

At T, optimal policy is $x^*_T(x) = 0$ if $x_t > 0$ for any $t < T$, and 1 otherwise.

At $T - 1$, it is more interesting. If $x_t = 0$ for all $t < T - 1$, then the optimal x^*_{T-1} is to delay to T if and only if

$$\theta_{T-1} > \beta\delta\theta_T.$$

$\therefore$ incentives to delay increase as $\beta \to 0$. ✓

Three-period example ($T = 3$)

Backwards induction:

- If $x_1 = x_2 = 0$, then $x_3^* = 1$.
- If $x_1 = 0$, then $x_2^* = 1 \iff -\theta_2 > -\beta\delta\theta_3$.
- Then payoffs from x_1 are

$$\begin{cases} -\theta_1 & \text{if } x_1 = 1 \\ -\beta\left[\delta\theta_2 x_2^* + \delta^2\theta_3 x_3^*\right] & \text{if } x_1 = 0 \end{cases} \tag{3}$$

  so that $x_1^* = 1 \iff -\theta_1 > \beta\left[\delta\theta_2 x_2^* + \delta^2\theta_3 x_3^*\right]$.

- So we deduce that

$$x^* = \begin{cases} (0,0,1) & \text{if } \beta\delta\theta_3 < \theta_2 \text{ and } \beta\delta^2\theta_3 < \theta_1 \\ (0,1,0) & \text{if } \theta_2 < \beta\delta\theta_3 \text{ and } \beta\delta\theta_2 < \theta_1 \\ (1,0,0) & \text{otherwise.} \end{cases} \tag{4}$$

  E.g., as $\beta \to 0$, $x^* = (0,0,1)$.
  As $\beta, \delta \to 1$, then $x^* = (1,0,0)$ (when θ_t increases in t).

Three-period example, continued

The x^* is the optimal policy or the agent's **behavior**. ✓

Welfare (utility) at $t = 1$ is given by

$$u(\theta, \delta, \beta) = \begin{cases} -\beta\delta^2\theta_3 & \text{if } \beta\delta\theta_3 < \theta_2 \text{ and } \beta\delta^2\theta_3 < \theta_1 \\ -\beta\delta\theta_2 & \text{if } \theta_2 < \beta\delta\theta_3 \text{ and } \beta\delta\theta_2 < \theta_1 \\ -\theta_1 & \text{otherwise.} \end{cases} \tag{5}$$

Now suppose the parameters are such that $x_1^* = 0$.

Demand for commitment. At $t = 1$, would prefer to commit to $x_2 = 1$ if

$$\theta_2 < \delta\theta_3$$

but in reality, will **not** do $x_2 = 1$ at $t = 2$ unless

$$\theta_2 < \beta\delta\theta_3.$$

Hence commitment has value when $\theta_2 \in [\beta\delta\theta_3, \delta\theta_3]$. In this region, the willingness to pay for a commitment device at $t = 1$ is $-\beta\delta(\theta_2 - \delta\theta_3)$.

Numerical example

Let $(\theta_1, \theta_2, \theta_3) = (\frac{8}{9}, 1, 2)$.

Let $\delta = 0.9$ and $\beta = \frac{1}{2}$. Recall the optimal policy is

$$x^* = \begin{cases} (0,0,1) & \text{if } \beta\delta\theta_3 < \theta_2 \text{ and } \beta\delta^2\theta_3 < \theta_1 \\ (0,1,0) & \text{if } \theta_2 < \beta\delta\theta_3 \text{ and } \beta\delta\theta_2 < \theta_1 \\ (1,0,0) & \text{otherwise.} \end{cases} \tag{4}$$

Check:

- $\beta\delta\theta_3 = \frac{1}{2} \cdot \frac{9}{10} \cdot 2 < 1 = \theta_2$ ✓
- $\beta\delta^2\theta_3 = \frac{1}{2} \cdot \frac{81}{100} \cdot 2 < \frac{8}{9} = \theta$ ✓

$\therefore$ Agent does the action at $t = 3$ by equation (4).

Would the agent prefer to do it at $t = 2$, from the viewpoint of $t = 1$? I.e., check if $\theta_2 < \delta\theta_3$:

$$\theta_2 = 1 < \delta\theta_3 = \frac{9}{10} \cdot 2$$

Indeed! The agent would. And the value of the commitment device is $-\beta\delta(\theta_2 - \delta\theta_3) = -\frac{1}{2} \cdot \frac{9}{10} \cdot (1 - \frac{18}{10}) = \frac{1}{2} \cdot \frac{9}{10} \cdot \frac{4}{5} = \frac{9}{25}$.

Beliefs

Studying the model further.

Now, although the agent's true preferences are still given by (1) in each t, the agent **thinks** that it will behave in the future as if its β were some $\hat{\beta}$. Say

- "naïve" if $\hat{\beta} = 1$
- "sophisticated" if $\hat{\beta} = \beta$

This affects the calculation of the x^*'s, which depend on $\hat{\beta}$! In the example, use $\hat{\beta}$ in (4) rather than the true β.

Remark. Frank's shortcut. If $\hat{\beta} = 1$, then you can calculate all of the x^*'s as in a "standard" (exponential-discounting) dynamic optimization problem starting at each t.

But to evaluate payoffs, still use the true β. E.g.,

$$\tilde{u}(\theta, \delta, \beta, \hat{\beta}) = \begin{cases} -\beta\delta^2\theta_3 & \text{if } \hat{\beta}\delta\theta_3 < \theta_2 \text{ and } \hat{\beta}\delta^2\theta_3 < \theta_1 \\ -\beta\delta\theta_2 & \text{if } \theta_2 < \hat{\beta}\delta\theta_3 \text{ and } \hat{\beta}\delta\theta_2 < \theta_1 \\ -\theta_1 & \text{otherwise.} \end{cases} \tag{6}$$

Recitation 3

Alex Olssen and Aaron Goodman

Outline

1. Quasi–hyperbolic Savings

2. Risk Aversion

Outline

1. Quasi–hyperbolic Savings
2. Risk Aversion (also Autor's notes on Stellar: Review notes 3/3)

Solving Problems with (Quasi-)Hyperbolic Discounting

- **Fully naïve** decision-makers ($\hat{\beta} = 1$):
  1. Start at the beginning.
  2. Solve for the optimal plan, assuming future selves will follow the plan.
  3. The person takes the first step in that plan.
  4. Go to the next period, and keep doing the same.
- **Fully sophisticated** decision-makers ($\hat{\beta} = \beta$):
  1. Start at the end.
  2. Solve for optimal action.
  3. Go back to the previous period.
  4. Solve for the optimal action, taking into account what happens in the next period.
  5. Go back to the previous period, and keep doing the same.
- **Partially naïve** decision-makers ($\beta < \hat{\beta} < 1$):
  1. Start at the end. Solve for what the person *thinks* she will do (using $\hat{\beta}$).
     [This is like solving for a fully sophisticated decision maker with a true β of $\hat{\beta}$.]
  2. Work your way to the first period using backward induction until period 2 (using $\hat{\beta}$).
  3. Then solve for the optimal action in period 1 (using the true β and the already derived prediction on future behavior).
  4. Then move to the next period, repeat steps (1) to (3).

The Model: Illiquid savings, credit card debt, commitment

- Alex is a fully naive hyperbolic discounter with $\beta = 0.5$ and $\delta = 1$ and $\hat{\beta} = 1$
- Alex lives for three periods $t = 0$, 1, and 2
- His instantaneous utility from consuming an amount $c_t > 0$ at time t is

$$u(c_t) = ln(c_t) \text{ for } t = 0, 1, 2$$

Alex's discounted lifetime utility from the perspective of period 0 is given by

$$U_0(c_0, c_1, c_2) = ln(c_0) + \beta(ln(c_1) + ln(c_2))$$

Moving money across periods (Q1.1)

- Alex starts with wealth of \$60 at $t = 0$
- Several ways to move money across periods
  - Checking account: put $\$x$ in at time t, can withdraw up to $\$x$ at $t + 1$
  - Retirement account: deposit s at $t = 0$, can withdraw $(1 + r^r)s$ at $t = 2$ $(r^r = .2)$
  - Credit card for $t = 1$: borrow b at $t = 1$, must repay $(1 + r^c)b$ at $t = 2$ $(r^c = .5)$
- How will Alex move money to $t = 1$? How about $t = 2$? Why?
  - To move money to $t = 1$, use checking account because alternative (credit card paid off at $t = 2$) is expensive
  - To move money to $t = 2$, use retirement savings because get a good return!

Optimal plan at $t = 0$ (Q1.2)

- Show that the consumption plan Alex makes at $t = 0$ involves $c_1 = \beta c_0$
- Given the previous answer, interest rate of 0 between $t = 0$ and $t = 1$
- Accordingly, he will equalize marginal utilities at $t = 0$ and $t = 1$
- Direct implication $c_1 = \beta c_0$ (let's work through the FOCs)

Optimal plan at $t = 0$ (Q1.3)

- Use (1) and (2), write Alex's maximization problem in period 0 and solve for planned c_0, c_1, and c_2
- Part (2) means $c_1 = \beta c_0$ at the optimum. Part (1) means we can ignore b. Thus

$$\max_{c_0, c_1, c_2} \; u(c_0) + \beta u(c_1) + \beta u(c_2)$$
$$\text{s.t. } \; c_1 = \beta c_0 \text{ and } c_2 = (60 - c_0 - c_1)(1 + r^r)$$

- Solution: $c_0^* = 30$, $\hat{c}_1 = 15$, and $\hat{c}_2 = 18$ (Let's work through FOCs)

Present Bias (Q1.4)

- What does Alex end up doing at $t = 1$?
- Being naive, at $t = 1$ Alex solves

$$\max_{c_1, c_2} \; u(c_1) + \beta u(c_2) \text{ s.t. } c_2 = \hat{c}_2 - (c_1 - \hat{c}_1)(1 + r^c)$$

- Taking the FOC and simplifying gives

$$\frac{1}{c_1} = \frac{\beta(1+r^c)}{c_2}$$
$$c_2 = \beta(1 + r^c)c_1$$

- Solution: $c_1^* = 18$, $b^* = 3$, and $c_2^* == 13.5$

Full Sophistication (Q1.9)

- Suppose Alex becomes fully sophisticated.
  Argue that at $t = 0$, Alex anticipates that at $t = 1$ he will choose c_1 and c_2 such that $c_2 = \beta(1 + r^c)c_1$.
- Being sophisticated, Alex understands that he will solve his consumption-savings decision in exactly the same way as already determined in (Q1.4)
- Recall that (Q1.4) was $c_2 = \beta(1 + r^c)c_1$

Full Sophistication (Q1.10)

- Write down Alex's maximization problem at $t = 0$. Explain what is different from Alex's maximization problem in part (3) and why
- Alex solves the following maximization problem:

$$\max_{c_0, c_1, c_2} \; u(c_0) + \beta u(c_1) + \beta u(c_2)$$
$$\text{s.t. } \; c_2 = \beta(1 + r^c)c_1 \text{ and } c_2 = (60 - c_0 - c_1)(1 + r^r)$$

- Fully sophisticated Alex knows he lacks time consistency
- Thus he solves his $t = 0$ problem with constraints that reflect his knowledge that he will re-optimize in the future

Commitment devices (Q1.11)

- Aaron offers (fully sophisticated) Alex a commitment device
- Can Alex be worse off (using discounted utility at $t = 0$) by (voluntarily) choosing *any* commitment contract that Aaron offers to him at $t = 0$?
- Solution: No, it is impossible for fully sophisticated Alex to be worse off.
- A fully-sophisticated agent anticipates his/her future behaviors
- At $t = 0$ Alex makes plans that maximize his utility from the perspective of $t = 0$
- If Aaron's commitment contract would make Alex worse off, then he would never (voluntarily) choose it

Commitment devices (Q1.12)

- Suppose Alex is partially naive
- Can Aaron make Alex worse off by offering him a commitment device (using discounted life-time utility at $t = 0$)?
- Yes, partially-sophisticated Alex can be worse off even when (voluntarily) choosing.
- Suppose the commitment device raises r^c at $t = 1$ above 50%.
- Alex might (voluntarily) choose the commitment device, hoping it will help him avoid borrowing.
- However, if β turns out to be (much) lower than anticipated, then he might end up borrowing at high interest rates after all
- This would make him worse off than he would have been borrowing at a 50% interest rate

Outline

Expected Utility Theory

- Describes agents' preferences and behavior when faced with uncertainty
- General lottery setup:
  - Agent gets utility from wealth $u(.)$
  - Potential states of the world: $i \in \{1, ..., n\}$
  - Each state has associated probabilities p_i and monetary payout x_i
- Expected value of lottery: $EX = \sum_{i=1}^{n} p_i x_i$
- Expected utility of lottery: $EU = \sum_{i=1}^{n} p_i u(x_i)$
- Utility of the expected value: $UE = u(\sum_{i=1}^{n} p_i x_i)$

Risk Preferences

- Risk loving: $EU > UE$
  - Prefers taking the lottery to receiving the expected value with certainty
- Risk neutral: $EU = UE$
  - Indifferent between taking the lottery and receiving the expected value with certainty
- Risk averse: $EU < UE$
  - Prefers receiving the expected value with certainty to taking the lottery

Curvature of $u(.)$

- Jensen's inequality: $f(.)$ is concave iff $f(\sum_{i=1}^{n} w_i y_i) > \sum_{i=1}^{n} w_i f(y_i)$
- Risk preferences involve comparison between:
  - $EU = \sum_{i=1}^{n} p_i u(x_i)$
  - $UE = u(\sum_{i=1}^{n} p_i x_i)$
- This implies:
  - Risk loving ($EU > UE$) iff $u(.)$ is convex
  - Risk neutral ($EU = UE$) iff $u(.)$ is linear
  - Risk averse ($EU < UE$) iff $u(.)$ is concave

Risk Aversion and Certainty Equivalents

- Certainty equivalent: the level of x that would make the agent indifferent between taking x and participating in the lottery
- Formally:
  - $u(CE) = EU = \sum_{i=1}^{n} p_i u(x_i)$
  - $CE = u^{-1}(EU) = u^{-1}(\sum_{i=1}^{n} p_i u(x_i))$
- Equivalent definition of risk preferences:
  - Risk loving if $CE > EX$
  - Risk neutral if $CE = EX$
  - Risk averse if $CE < EX$

Risk Aversion in a Picture

Lottery with 2 outcomes: (1) $x_1 = x$, $p_1 = p$; (2) $x_2 = y$, $p_2 = (1 - p)$

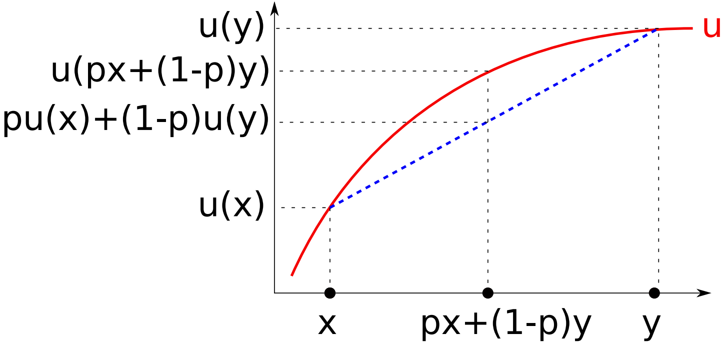

- Where is *EX*? *EU*? *UE*? *CE*?

CARA

- Coefficient of absolute risk aversion: $r = -\frac{u''(x)}{u'(x)}$
  - Normalized by $u'(x)$ (why?)
- Constant absolute risk aversion (CARA) utility: $u(x) = -\frac{e^{-rx}}{r}$
  - Absolute risk aversion is constant in x
- Problem: we typically believe wealthier people are riskier so risk aversion should be decreasing in x

CRRA

- Coefficient of relative risk aversion: $\gamma = -\frac{xu''(x)}{u'(x)}$
- Constant relative risk aversion (CRRA) utility: $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$
  - CRRA utility generates constant relative risk aversion
  - CRRA utility generates absolute risk aversion that is decreasing in wealth

Risk Aversion Takeaways

- Expected utility is (another) work horse model in economics
- Important distinction between the expected value of an uncertain lottery and the expected utility
- Risk aversion explains why people want insurance (some of the biggest markets in the economy are insurance markets)
- CARA and CRRA utility functions are common special cases (worth knowing)
- For further reading, see David Autor's notes on Stellar (Review notes (3/3) risk preferences)

Recitation 4

Aaron Goodman, Alex Olssen, Pierre-Luc Vautrey[1]

[1] These slides are partially based on notes from Drew Fudenberg. All errors are our own.

Outline

1. Rabin (2000)

2. Example problem on risk preferences

Outline

1. Rabin (2000)
2. Example problem on risk preferences

Recap: Expected Utility Theory

In recitation last week and lecture this week, we introduced expected utility theory:

- States of the world $i = \{1, .., n\}$, probabilities p_i, payoffs x_i
- Utility function $u(\cdot)$
- Expected utility is given by

$$EU = \sum_i p_i u(x_i) \tag{1}$$

- We generally assume that $u(\cdot)$ is concave, so agents are risk averse and

$$\sum_i p_i u(x_i) < u\left(\sum_i p_i x_i\right) \tag{2}$$

Rabin (2000)

- Rabin's paper is a very influential critique of expected utility theory
- Main idea: concavity of the utility function cannot be the only source of risk aversion. If it is, then we obtain some absurd results.
- Helpful to understand Rabin's argument, especially as we begin to consider deviations from expected utility theory (loss aversion, reference dependence, etc.) that address his critique
- The discussion today is only meant to be instructive - we won't ask you to prove Rabin's result!

Setup

- Consider an agent with utility function $u(\cdot)$ defined over wealth w
- Assume that at all wealth levels, the agent rejects a 50-50, lose \$100, gain \$110 gamble:

$$\frac{1}{2}u(w-100)+\frac{1}{2}u(w+110)\leq u(w) \tag{3}$$

$$\implies u(w+110)-u(w)\leq u(w)-u(w-100) \tag{4}$$

- Sounds like a reasonable assumption, but will see that it leads to unreasonable results!

First Step

- First, observe that:

$$\begin{aligned} 110u'(w+110) &\leq u(w+110) - u(w) && (5) \\ &\leq u(w) - u(w-100) && (6) \\ &\leq 100u'(w-100) && (7) \end{aligned}$$

- How do we justify each of these inequalities?
- Rearranging, we obtain

$$110u'(w+110) \leq 100u'(w-100) \quad (8)$$

$$\frac{u'(w+110)}{u'(w-100)} \leq \frac{10}{11} \quad (9)$$

Concavity

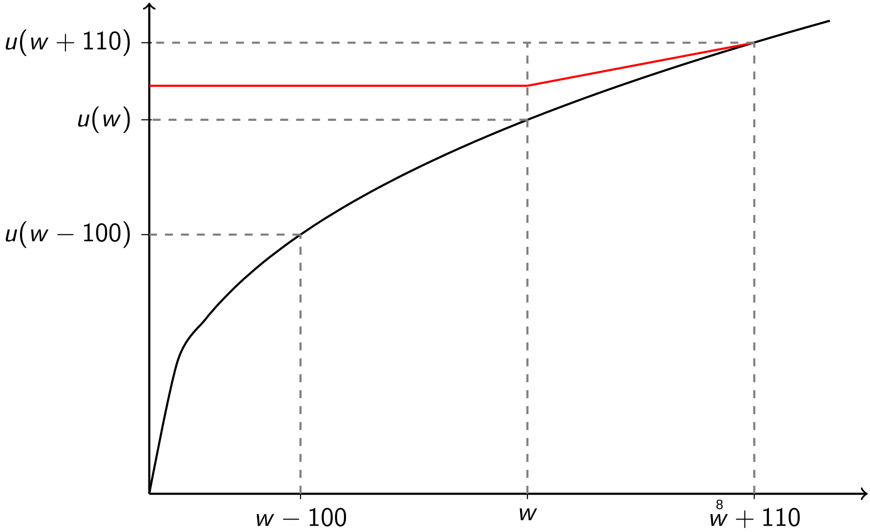

Concavity

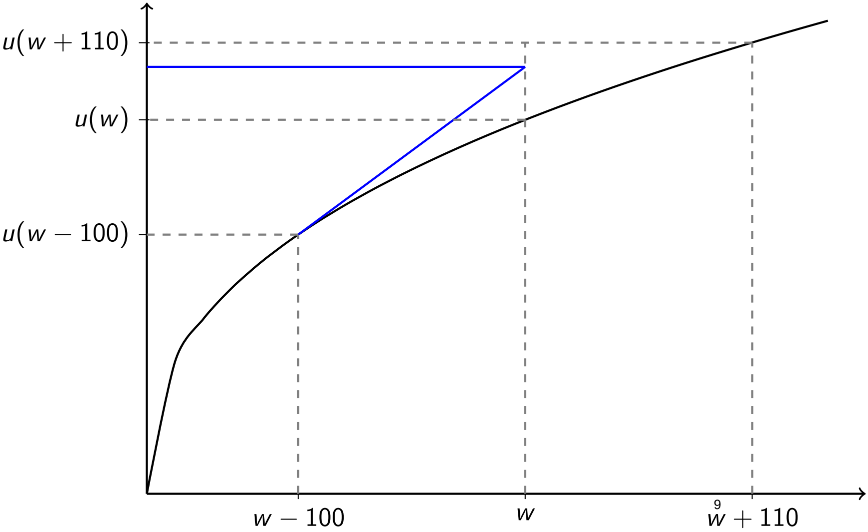

Iterating Forward

- Under our assumption, the agent also rejects the gamble when his wealth is $w + 210$. Applying the same logic, we obtain:

$$\frac{u'(w+210+110)}{u'(w+210-100)} = \frac{u'(w+320)}{u'(w+110)} \leq \frac{10}{11} \tag{10}$$

- This implies:

$$\frac{u'(w+320)}{u'(w-100)} = \frac{u'(w+320)u'(w+110)}{u'(w+110)u'(w-100)} \leq \left(\frac{10}{11}\right)^2 \tag{11}$$

- We can do this again:

$$\frac{u'(w+530)}{u'(w-100)} = \frac{u'(w+530)u'(w+320)}{u'(w+320)u'(w-100)} \leq \left(\frac{10}{11}\right)^3 \tag{12}$$

Keep Iterating Forward

- We can do this as many times as we want. In general:

$$\frac{u'(w+210k+110)}{u'(w-100)} \leq \left(\frac{10}{11}\right)^{k+1} \qquad k=1,2,\ldots \tag{13}$$

- Takeaway message: to justify seemingly reasonable risk aversion over small gambles (e.g., our lose \$100, gain \$110 bet), marginal utility must be diminishing very fast. If we iterate forward 100 times, then:

$$\frac{u'(w+210(100)+110)}{u'(w-100)} = \frac{u'(w+21110)}{u'(w-100)} \leq \left(\frac{10}{11}\right)^{101} \approx 0.00007 \tag{14}$$

Diminishing Marginal Utility

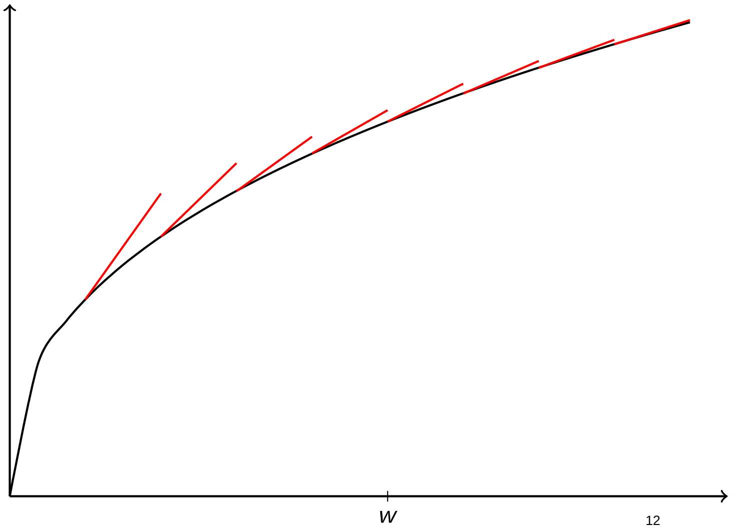

- Each slope is at most $\frac{10}{11}$ of the last

Implications

- Because marginal utility is diminishing so quickly, our agent turns down gambles with enormous upside
- In fact, there is no number x such that our agent will accept a 50-50, lose \$1,000, gain \$$x$ gamble. He refuses this offer even if $x = \infty$!
- The marginal utility of wealth becomes infinitesimally small at large dollar values, so the upside of any such gamble is outweighed by the downside:

$$u(w + x) - u(w) \leq u(w) - u(w - 1000) \;\; \forall x \tag{15}$$

Rabin's Corollary

TABLE I

IF AVERSE TO 50-50 LOSE \$100 / GAIN g BETS FOR ALL WEALTH LEVELS, WILL TURN DOWN 50-50 LOSE L / GAIN G BETS; G'S ENTERED IN TABLE.

		g		
L	\$101	\$105	\$110	\$125
\$400	400	420	550	1,250
\$600	600	730	990	∞
\$800	800	1,050	2,090	∞
\$1,000	1,010	1,570	∞	∞
\$2,000	2,320	∞	∞	∞
\$4,000	5,750	∞	∞	∞
\$6,000	11,810	∞	∞	∞
\$8,000	34,940	∞	∞	∞
\$10,000	∞	∞	∞	∞
\$20,000	∞	∞	∞	∞

Outline

1. Rabin (2000)
2. Example problem on risk preferences

Setup

From problem set 2 in 2017 (on course website):

- Alex is buying home insurance
- His current wealth is $w = \$100{,}000$
- He has CRRA utility with coefficient of relative risk aversion γ
- Damage occurs to his house next year with probability $\pi = .05$

Plan Choices

Alex is offered four plans by his insurance company

- Assume that not buying insurance is not an option
- Assume that if damage occurs, it always exceeds the deductible

Option	Deductible	Premium
1	1,000	757
2	500	885
3	250	999
4	100	1,171

Plan Choices

We can also represent the plans in terms of Alex's terminal wealth in each state of the world:

Option	Damage	No Damage
1	$w-1{,}757$	$w-757$
2	$w-1{,}385$	$w-885$
3	$w-1{,}249$	$w-999$
4	$w-1{,}271$	$w-1{,}171$

Is there a plan that Alex will *never* choose, regardless of his risk preferences?

Bounding Risk Aversion

Suppose Alex chooses plan 2. Calculate bounds on his risk aversion parameter γ.

What's the first step in answering this question?

Write down the expected utility of choosing plan j, with premium p_j and deductible d_j:

$$V_j = \pi u(w - p_j - d_j) + (1 - \pi) u(w - p_j) \tag{16}$$

$$= \pi \frac{(w - p_j - d_j)^{1-\gamma}}{1 - \gamma} + (1 - \pi) \frac{(w - p_j)^{1-\gamma}}{1 - \gamma} \tag{17}$$

Alex chooses the plan that maximizes his expected utility:

$$j^* = \underset{j \in \{1,2,3\}}{\text{argmax}} \; V_j \tag{18}$$

Bounding Risk Aversion

Since Alex chose plan 2, we have, for $k \in \{1, 3\}$:

$$V_2 \geq V_k \tag{19}$$

How do we use this to bound γ?

$$\pi u(w - p_2 - d_2) + (1 - \pi)u(w - p_2) \geq \pi u(w - p_k - d_k) + (1 - \pi)u(w - p_k) \tag{20}$$

Bounding Risk Aversion

We thus have:

$$0.05 \cdot (w - 1,385)^{1-\gamma} + 0.95 \cdot (w - 885)^{1-\gamma} \geq 0.05 \cdot (w - 1,757)^{1-\gamma} + 0.95 \cdot (w - 757)^{1-\gamma}$$
$$0.05 \cdot (w - 1,385)^{1-\gamma} + 0.95 \cdot (w - 885)^{1-\gamma} \geq 0.05 \cdot (w - 1,249)^{1-\gamma} + 0.95 \cdot (w - 999)^{1-\gamma}$$

Using a computer, we find that the first inequality implies

$$\gamma \geq 243.26$$

and the second inequality implies

$$\gamma \leq 726.50$$

Why does the first inequality place a lower bound on γ? Why does the second inequality place an upper bound on γ?

Note: these are implausibly high values for risk aversion!

Recitation 5: Reference Dependence

Pierre-Luc Vautrey

Reference Dependence: 3 Key Ingredients

1. Utility evaluate things (consumption, ...) **relative** to something, rather than in some absolute terms. What matters is **changes** rather than levels.
2. **Loss aversion**: losses hurt more than symmetric gains help
3. **Diminishing sensitivity**: Changes far away from the reference matter less than changes close to the reference

Utility is Evaluated Relative to a Reference Point

- Typically, evidence supporting this idea are behavioral patterns of bunching near some arbitrary level
  - Examples: Marathon runners bunch around salient times, taxi drivers vary labor supply to reach daily income targets
- What is the "reference point"?
  - An expectation
  - A goal or aspiration
  - A status quo
  - A starting point
  - An anchor
- General set-up:
  - Utility is over consumption, c, relative to a reference point, r; that is, utility is over $x = c - r$
  - The function $u(x)$ may take different shapes for $x > 0$ and $x < 0$

Loss Aversion

- The idea of evaluating changes rather than levels becomes much more concrete when we add in loss aversion
- A form of reference dependence with much empirical support
- Loss aversion: losses hurt more than gains help

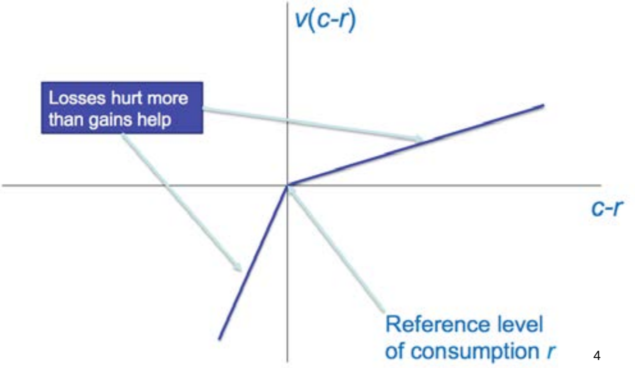

4

Loss Aversion

- Example: if I expect my birthday check to be \$100, the util cost of a \$90 check will be greater than the util gain of a \$110 check
  - What theory of the reference point is used in this example?
  - Expectation-based reference point
  - Another possibility: anchoring on 100?

5

Application: Endowment Effect

- Endowment effect: owning a good makes you less likely to trade it away
  - Widely documented empirical fact
  - Mugs experiments
  - Note: agent needs to think he owns the good - simply handing him the good if he thinks it will be taken away is not sufficient to generate endowment effects
- Loss aversion can explain it:
  - Your endowment becomes your reference point
  - The util gain you would get from the new item is less than the util cost of trading the endowment away
- Consequences e.g. for marketing:
  - Make potential consumers feel they own the good
  - Offer free returns

6

Diminishing Sensitivity

- Diminishing Sensitivity: sensitivity to additional changes in consumption should be smaller the further the changes are from the reference point
- Example: the utils cost from moving to a $100 to a $90 birthday check is greater than the cost of moving from $90 to $80

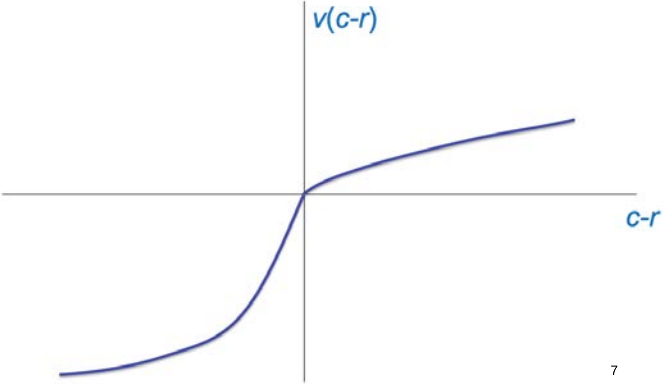

Implications: Convexity and Concavity

- Loss aversion and diminishing sensitivity implies usual concave utility in gains
- However, we have convex utility over losses
- In situations of risk, this implies risk aversion over gains and risk loving over losses
  - Helps explain empirical findings that people become risk-seeking after losses
  - Deal or No Deal? example
    - What is the reference point here?
  - Also from experiment with lotteries framed in loss domain or gain domain
    1. Here is $100. What do you think of: loosing 50 for sure vs loosing 100 with 50% probability?
    2. What do you think of winning 50 for sure vs winning 100 with 50% probability?
  - Diminishing sensitivity also means the risk-lovingness decreases once we get far in the loss domain

8

Implications: Better to Loose in Batch and Gain In Small Increments

- Diminishing sensitivity implies that marginal utility is highest near the reference point
  - Intuitively, this means small changes have the highest effect there
  - A small gain implies a utility increase greater than a tenth of the utility increase from a gain ten times larger
  - A small loss implies a utility decrease greater than a tenth of the utility decrease from a loss ten times larger
- If a large consumption gain can be broken down over time, and *if the reference point adapts* after each fraction of the gain, then it's better
- And conversely for losses: frequent small losses hurt much more than a comparable one-shot loss *if the reference point adapts*
- Implication: sellers should revise prices upwards unfrequently!

9

Question: Set Up

Maddie has wanted to buy several pairs of new pants, but she has not done so yet. Suppose she has reference-dependent utility over pants c_p and money c_M of the form

$$v(50c_p - 50r_p) + v(c_M - r_M), \tag{1}$$

where $v(x) = x$ for $x \geq 0$ and $v(x) = 2x$ for $x < 0$. Maddie has \$300 in cash.

- *What are r_p and r_M?*
- *What about the utility function makes her reference dependent?*

Part 1: Question

Suppose Maddie is expecting to buy two pairs of pants at \$40 each. That is, her reference point for pants is $r_P = 2$, and her reference point for money is $r_M = 220$. What is the maximum price, $p_{\max}$, at which she is willing to buy the first two pairs of pants?

- *Why do* $r_P = 2$ *and* $r_M = 220$ *become her reference points?*
- *Would we expect the price she is willing to pay for two pants to be above, below, or equal to* \$80*?*

Part 1: Question

Suppose Maddie is expecting to buy two pairs of pants at \$40 each. That is, her reference point for pants is $r_P = 2$, and her reference point for money is $r_M = 220$. What is the maximum price, $p_{\max}$, at which she is willing to buy the first two pairs of pants?

- *Why do $r_P = 2$ and $r_M = 220$ become her reference points?*
- *Would we expect the price she is willing to pay for two pants to be above, below, or equal to* \$80*?*

Part 1: Solution

- Let p_{max} be the maximum willingness to pay for 2 pairs
- By definition, p_{max} is the price such that she is indifferent between not buying, $(c_P, c_M) = (0, 300)$, and buying, $(c_P, c_M) = (2, 300 - p_{max})$
- Hence, p_{max} solves:

$$\begin{aligned}
v(50 \cdot 0 - 50 \cdot 2) + v(300 - 220) &= v(50 \cdot 2 - 50 \cdot 2) + v(300 - p_{max} - 220) \\
v(-100) + v(80) &= v(0) + v(80 - p_{max}) \\
-200 + 80 &= 2(80 - p_{max}) \\
-60 &= 80 - p_{max} \\
p_{max} &= 140.
\end{aligned}$$

Part 2: Question

Now suppose that after buying two pairs of pants at the price of \$40, unexpectedly, Maddie is contemplating buying another pair of pants. What is her maximum willingness to pay, p'_{max}, for a third pair of pants?

- *What does "unexpectedly" imply for her reference points?*
- *Should we expect p'_{max} to be above, below, or equal to \$40?*

Part 2: Solution

- $r_p = 2, r_m = 220$
- By definition, p'_{max} is the price such that Maddie is indifferent between not buying, $(c_P, c_M) = (2, 220)$, and buying $(c_P, c_M) = (3, 220 - p'_{max})$
- Hence, p'_{max} solves

$$\begin{aligned}
v(50 \cdot 2 - 50 \cdot 2) + v(220 - 220) &= v(50 \cdot 3 - 50 \cdot 2) + v((220 - p'_{max}) - 220) \\
v(0) + v(0) &= v(50) + v(-p'_{max}) \\
0 &= 50 - 2p'_{max} \\
p'_{max} &= 25.
\end{aligned}$$

Part 3: Question

Suppose the salesperson, Allan, exactly knows Maddie's preferences. To entice her, he offers her a bundle of three pants at

$$p_{\mathrm{b}} = p_{\max} + p'_{\max} - \varepsilon = 165 - \varepsilon, \tag{2}$$

where ε is very small, such that Maddie cannot resist and buys the three pairs of pants from him. If Allan had called Maddie ahead of time to tell her about the deal (i.e. p_{b}), her reference point would have adjusted to three pants. How (if at all) would her willingness to pay for the bundle upon arrival at the store have changed?

- *Should we expect her willingness to pay for the bundle to be less than, greater than, or equal to* 165*?*

Part 3: Solution

- If Allan called before the deal, then $r_P = 3, r_M = 300 - p_b$
- p_A is the price such that Maddie is indifferent between not buying $(c_P, c_M) = (0, 300)$ and buying $(c_P, c_M) = (3, 300 - p_A)$
- Hence, p_A solves

$$v(50 \cdot 0 - 50 \cdot 3) + v(300 - (300 - p_b)) = v(50 \cdot 3 - 50 \cdot 3) + v((300 - p_A) - (300 - p_b))$$

$$\begin{aligned} v(-150) + v(p_b) &= v(0) + v(p_b - p_A) \\ -300 + p_b &= 2(p_b - p_A) \\ p_A &= 0.5p_b + 150 = 232.5 - 0.5\varepsilon. \end{aligned}$$

Part 4: Question

Assume Allan can only sell one bundle of three pairs to Maddie and no other combinations. If Allan wants to maximize his revenues, what bundle price should he tell Maddie about on the phone? And what actual price should he charge upon her arrival at the store? (Assume the price over the phone and in the store can be different)

Part 4: Solution

- From part 3, we know Allan can announce p_b and charge $p_A = 0.5p_b + 150$ in the store
- To maximize revenue, Allan will announce the highest p_b he can
- He should tell Maddie the p_b that makes her indifferent between coming to the store and not
- When Allan calls, her reference point will still be for 2 pants; her maximum willingness to pay will be the sum of the prices from parts 1 and 2, \$165
- He will announce $p_b = \$165 - \varepsilon$ and charge $\$232.5 - 0.5\varepsilon$.

Part 5: Question

1. How much is Maddie willing to pay if she expects to buy 0 pants? (Assume the actual price would still be $ 40)
2. What is the lowest price would Maddie ask for if someone wanted to buy the pants after she bought them?
3. What is the difference between (1) and (2) reflect?

Part 5: Solution

5.1: From part 2, \$25

5.2:

- $r_p = 1, r_m = 260$
- Let p''_{max} be the price such that Maddie is indifferent between not selling, $(c_P, c_M) = (1, 260)$, and selling $(c_P, c_M) = (0, 260 + p''_{max})$
- Hence, p''_{max} solves

$$\begin{aligned} v(50 \cdot 1 - 50 \cdot 1) + v(260 - 260) &= v(50 \cdot 0 - 50 \cdot 1) + v((260 + p''_{max}) - 260) \\ v(0) + v(0) &= v(-50) + v(p''_{max}) \\ 0 &= -100 + p''_{max} \\ p''_{max} &= 100. \end{aligned}$$

5.3: an endowment effect

Pants, Mugs, and Pencils

Pants: key take-aways

- Reference dependence means willingness to pay changes based on expectations
- The specific type of reference dependence, loss aversion, means Maddie values the good more if she expects to have it

Pants, Mugs, and Pencils

- How are these results related to the mug-pencil example from class?
- Why do we expect people to exchange mugs and pencils?
- Why is a lack of trading consistent with an endowment effect?

Recitation 6: Midterm Review

Will Rafey

Outline

Midterm consists of three parts:

1. True/false
   - State true, false, or uncertain
   - Always explain answer carefully
   - Need to provide intuition
2. Multiple choice
3. Short answer (similar to problem set)

Most important resources:

- lecture + recitation slides
- problem sets and solutions.

True-False: Example 1

T/F. Consider individuals with "β, δ" preferences, who only differ by their present bias, $\beta \in [0, 1]$. Suppose there is a commitment savings device available. Willingness to pay for this commitment device strictly decreases in β.

False. Why?

- Individuals may be naïve
- Commitment device may not be effective
- Even if individuals are fully sophisticated and the device is effective, willingness to pay may not be strictly decreasing.
  - Individuals would be willing to pay 0 for $\beta = 0$ and for $\beta = 1$, but willing to pay a positive amount for $\beta \in (0, 1)$.

True-False: Example 2

T/F. Fully sophisticated individuals can experience large welfare losses from their present bias.

True. Why?

- Awareness of present bias (i.e. sophistication) does *not* remove present bias
- Sophisticates that lack commitment devices may still make suboptimal decisions

True-False: Example 3

T/F. Present-biased individuals will always have positive demand for commitment devices.

False. Why?

- Three conditions must be met for positive demand for commitment:
    1. Individuals must be present-biased.
    2. Individuals must be aware of their present-bias (i.e. they can't be fully naive).
    3. Individuals must perceive the commitment device as effective in helping overcome the self-control problem.
- When only the first is met, we cannot be sure there will be positive demand for commitment.

Multiple Choice: Example 1

Pierre-Luc is writing a problem set for 14.13. He gets utility $u(q)$ from the number of questions he writes. He has reference dependent preferences around his goal of writing 10 questions. Normalize $u(10) = 0$. Which of the following would be consistent with loss aversion?

(a) $u(8) = -2, u(12) = 2$
(b) $u(8) = -2, u(12) = 1$
(c) $u(8) = -1, u(12) = 2$

(b). Why?

- Loss aversion means losses hurt more than gains help
- With preferences in (b), Pierre-Luc would have a utility cost of 2 from falling short of his goal by 2 questions, but only gain 1 util from exceeding his goal by 2 questions.

Multiple Choice: Example 2

Q: Maddie is walking home and passes a bakery. She suddenly decides to buy a pastry. Prior to purchasing the pastry, her maximum willingness to pay for the pastry was p_0. She then runs into Pierre-Luc who asks to buy the pastry from her. She offers him the lowest price she is willing to accept, p_1. Which of the following comparisons between p_0 and p_1 is consistent with an endowment effect?

(a) $p_0 > p_1$
(b) $p_0 = p_1$
(c) $p_0 < p_1$

(c). Why?

- Consistent with an endowment effect, $p_0 < p_1$ implies Maddie values the pastry more after she has bought it than prior to buying it.

Q: Now suppose that Maddie first notices the pastry has gone stale, before she offers Pierre-Luc a price. Maddie always prefers fresher pastries. Which of (a)–(c) is consistent with the endowment effect?

Long Question: Example 1

Present Bias

Setup. Assume 14.13 students are present biased with $\beta < 1$ and $\delta = 1$. All students have the same $\beta < 1$ and $\delta = 1$ but differ in the value they derive from using laptops in class, L_i.

L_i is uniformly distributed across students i on the interval [0,1].

Each lecture generates no immediate utility, but does give a future benefit V. Using a laptop reduces the long-run benefit by D. Both V and D are the same for all students.

In summary, a student that uses a laptop in class gets immediate utility L_i and future (undiscounted) utility $V - D$. A student that does not use a laptop gets immediate utility 0 and future (undiscounted) utility V.

The social planner is not present biased and seeks to maximize the utility of 14.13 students.

Present bias

1(a). Show that a student i^* is just indifferent between using and not using their laptop in the current class if $L_{i^*} = \beta D$. Explain why students with lower values of L_i (i.e. $L_i < \beta D$) don't use laptops in class, while students with higher values of L_i (i.e. $L_i > \beta D$) do use laptops in class.

Present bias, cont'd

Utilities from the two choices are:

$$\begin{aligned} U(\text{laptop}) &= L_i + \beta\delta(V - D) \\ U(\text{nolaptop}) &= 0 + \beta\delta V \end{aligned}$$

For students that are indifferent, $U(laptop) = U(nolaptop)$. This gives:

$$\begin{aligned} L_{i^*} + \beta\delta(V - D) &= 0 + \beta\delta V \\ L_{i^*} &= \beta\delta D \end{aligned}$$

Students that choose not to use laptops will have low valuations of using laptops, while students that choose to use laptops will have higher valuations. Given the indifference condition and $\delta = 1$,

- Students i that do not use laptops: $L_i < \beta D$
- Students i that use laptops: $L_i > \beta D$

Present bias, cont'd

1(b). Now consider the policy that allows students to use laptops only if they sign up **in advance** to sit in a laptop section. Why is $L_i \geq D$, not $L_i \geq \beta D$, the threshold for opting into the laptop section?

Present bias, cont'd

Solution to 1(b)

Considered in advance, students evaluate:

$$\begin{aligned} U(laptop) &= 0 + \beta(\delta L_i + \delta^2(V - D)) \\ U(nolaptop) &= 0 + \beta\delta^2 V \end{aligned}$$

The threshold for opting in is defined by $U(laptop) \geq U(nolaptop)$. Using $\delta = 1$, this gives:

$$\begin{aligned} 0 + \beta(L_i + V - D) &\geq 0 + \beta V \\ L_i &\geq D \end{aligned}$$

The threshold changes from βD to D because when laptop use can only happen in the future, all benefits and costs are discounted at the same rate, β.

Present bias, cont'd

1(c). Assume there is no laptop policy. Show that if $\beta D < L_i < D$, the student i engages in preference reversals: she prefers not to use the laptop in future classes, but changes her mind when she's actually sitting in those future classes.

- When thinking about future laptop use, the student's problem is identical to the problem in part (b). Why?
  - Because she discounts time both one and two periods in advance by β
- We know from part (b) that if $L_i < D$, she would like to not use the laptop
- But from part (a), we know that if $\beta D < L_i$, she will end up using the laptop when she's actually sitting in the future class
- This implies a preference reversal! she prefers not to use the laptop in future classes, but switches her mind when she's actually sitting in those future classes.

Present bias, cont'd

1(d). Explain why fraction $1 - \beta D$ of the class uses a laptop in part 1, but fraction $1 - D$ of the class uses a laptop in part 2. Why does a smaller share of the class use their laptops in part 2?

Present bias, cont'd

Solution to 1(d)

In part 1, a student uses a laptop if $L_i > \beta D$. Define $F(\cdot)$ as the CDF of L_i. Given the uniform distribution:

$$\begin{aligned} P(L_i > \beta D) &= 1 - F(\beta D) \\ &= 1 - \beta D \end{aligned}$$

Likewise, in part 2, a student uses a laptop if $L_i > D$. We have:

$$\begin{aligned} P(L_i > D) &= 1 - F(D) \\ &= 1 - D \end{aligned}$$

A smaller share will use laptops in part 2 because the benefit of using a laptop is delayed and hence discounted by β.

Present bias, cont'd

1(e). Why would the social planner prefer the opt-in policy to both the policy of allowing students to choose whether to use their laptops and to banning laptops altogether?

- The planner is not present biased so would want only students with $L_i > D$ to use laptops; the opt-in policy achieves this
- Under the free choice policy, students with $\beta D < L_i < D$ will suboptimally use their laptops
- On the other hand, banning laptops altogether is suboptimal because welfare is gained by allowing the students with the highest valuations, $L_i > D$, use laptops

Long Question, Example 2

Reference dependence

Frank has reference-dependent preferences over donuts d and coffee k, which cost \$1 each. MIT gives him \$13 to spend at the coffee shop. His utility takes the form

$$u(d,k) = u_1(d-6) + u_2(k-2)$$

where

$$u_1(x) = \begin{cases} 2\sqrt{x} & \text{if } x \geq 0 \\ -4\sqrt{|x|} & \text{if } x < 0 \end{cases} \tag{1}$$

and

$$u_2(x) = \begin{cases} \sqrt{x} & \text{if } x \geq 0 \\ -2\sqrt{|x|} & \text{if } x < 0. \end{cases} \tag{2}$$

2(a). If Frank has six donuts, is Frank loss averse to changes in his donut supply?
Yes!

Reference dependence

2(b). Frank buys a positive number of donuts and a positive number of coffees. How many donuts and coffee should Frank buy?

Answer: the Lagrangian is

$$\mathcal{L}(d, k, \lambda) = u_1(d-6) + u_2(k-2) + \lambda \cdot (13 - d - k)$$

When $d, k > 0$, then

$$\frac{\partial u_1}{\partial d} = (d-6)^{-1/2} = \lambda$$

and

$$\frac{\partial u_2}{\partial k} = \frac{1}{2}(k-2)^{-1/2} = \lambda.$$

Then $d - 6 = \lambda^{-2}$ and $k - 2 = 2^{-2}\lambda^{-2}$, so $4(k-2) = d - 6$. And $k + d = 13$. So

$$4k - 8 = d - 6 = 13 - k - 6$$

so that $5k = 21 - 6$ or $\boxed{k = 3}$ and $\boxed{d = 10}$.

Frank's utility is $u(10, 3) = 2\sqrt{4} + \sqrt{1} = 5$.

Reference dependence, cont'd

2(c). Someone tells Frank that they eat fewer than six donuts per day; specifically, they eat two donuts. Frank decides he should cut back his reference point to two donuts, as a benchmark. His new preferences are

$$u(d, k) = u_1(d - 2) + u_2(k - 2).$$

Is Frank happier?

Yes! $u_1(d - 2) > u_1(d - 6)$ for all d.

Reference dependence, cont'd

2(d). Frank has bought his donuts and returned to his office. A doctor arrives from MIT Medical. Frank has a suspicion that the doctor will prescribe any desired level of donuts, $\overline{d} \geq 0$, that he asks. Frank's preferences then will become

$$u(d, k) = u_1(d - \overline{d}) + u_2(k - 2).$$

What does Frank ask the doctor to prescribe?

Frank's utility is always diminishing in $\overline{d}$, his reference level for donuts! He asks the doctor to prescribe $\overline{d} = 0$.

2(e). Now the doctor demands payment for his medical wisdom. How much is Frank's maximum willingness to pay the doctor for these new preferences?

Frank's utility rises to $2\sqrt{10}$ from $2\sqrt{10-6}$, so he is willing to pay $2(\sqrt{10} - \sqrt{4})$.

Reference dependence, cont'd

2(f). Suppose that the doctor is receiving payments from the donut industry and can only prescribe $\overline{d} = 1$, but will now also give Frank a machine that allows him to costlessly transform donuts into coffee and vice-versa. How much is Frank now willing to pay the doctor (in utils)?

If Frank can revise his consumption, his first-order conditions become

$$4(k-2) = d - 1$$

or

$$4k - 8 = 13 - k - 1$$

so that $5k = 20$, or $k = 4$ and $d = 9$.

$\therefore$ With the time machine and $\overline{d} = 1$, Frank will obtain $2\sqrt{9-1} + \sqrt{4-2} = 2\sqrt{8} + \sqrt{2} = 5\sqrt{2}$.

$\therefore$ Frank's WTP $\leq 5\sqrt{2} - 5$.

Psychology and Economics

14.13 Mid-term Review

Frank Schilbach & Will Rafey

MIT

Midterm: Overview

- Remember: the class is pass/fail. Try your best but do NOT stress about or lose sleep over this exam.
- You will do fine and pass as long as you answer all questions and get at least some of the questions almost right.
- Exam will be posted on on Monday (April 6) at 8:00 am EST.
    - You administer the exam online yourself.
    - You pick your own two-hour window to complete the exam.

What resources are you allowed to use while taking the exam?

- You can use slides and notes from lectures, recitations, psets, etc.
- You CANNOT consult or receive help from others while taking the exam (online, in person, or any other way).
- You CANNOT find try to answers to your the questions online other than the Learning Modules website (e.g. ask you CANNOT ask question on Piazza or try to google questions or answers).
- You CANNOT watch lecture videos during the exam.
- Support animals are fine!
- Honor code: we trust you to stick to those rules.

Midterm: Three types of questions

(I) True/false/uncertain

- State true, false, or uncertain
- Always explain answer carefully
- Need to provide intuition.
- Using math might be helpful but you always need to provide a verbal explanation.

(II) Multiple choice

- Pick correct answers, no further explanation needed

(III) Pset-style questions

- Similar to problem set questions
- Some algebra involved
- Always explain your answers carefully.

Midterm: How to best prepare?

- What materials are you responsible for?
  - Lectures up to and including lecture 12 (March 11) [up to slide 67 of lectures 11 to 13]
  - Recitations 1 to 5 [recitations 6 and 7 are just reviews that might be helpful for some of you]
  - Psets 1 to 3
  - Readings (starred or non-starred) cited in class are only relevant to the extent that they appear in lectures and/or recitation.

- How to get ready?
  - Study lecture and recitation slides carefully
  - Psets and solutions: make sure you understand and are able to solve psets on your own.
  - Great resource to practice: previous psets and exams
  - Readings (starred or non-starred) may help you deepen your understanding of the material but we won't ask about details of those readings that beyond what was covered in class.

Time preferences: Exponential discounting model

- What is the exponential discounting model?
- What is δ? What does it measure? How can we estimate it?
- What are the main assumptions of this model?
- What evidence do we have against those assumptions?

Time preferences: Quasi-hyperbolic discounting model

- What is the quasi-hyperbolic discounting model? How is it different from the exponential discounting model?
- What empirical evidence can the quasi-hyperbolic model explain better than the exponential discounting model? Why?
- Sophistication and (partial) naivete
  - What does β measure? What does $\hat{\beta}$ measure?
  - Full sophistication, full naivete, partial naivete
  - Does sophistication make people always better off? Why (not)?
- Demand for commitment
  - What is demand for commitment? Who demands commitment, who doesn't?
  - What kinds of people do (not) demand commitment?
  - Can people be worse off from being offered a commitment device? Why (not)?

Time preferences: Empirical applications and solving problems

- Empirical applications
  - Be familiar with the empirical applications from lectures 5 and 6
  - Understand why the quasi-hyperbolic model can explain (some of) the empirical evidence better than the exponential discounting model.

- You need to be able to solve problems:
  - for exponential discounters
  - for quasi-hyperbolic discounters
  - for fully naive, fully sophisticated, and partially naive agents

- How does one solve such problems?
  - Solving problems forward and backward (depending on the case)
  - See slide 62 of lectures 3 & 4, and slide 37 of lectures 5 & 6
  - Plenty of pset and mid-term examples to practice with

Time preferences: Example of True/False/Uncertain Question 1

Statement: *Consider individuals with "β, δ" preferences, who only differ by their present bias, $\beta \in [0, 1]$. Suppose there is a commitment savings device available. Willingness to pay for this commitment device strictly decreases in β.*

False. Why?

- Individuals may be naive
- Commitment device may not be effective
- Even if individuals are fully sophisticated and the device is effective, willingness to pay may not be strictly decreasing.
  - Individuals would be willing to pay 0 for $\beta = 0$ and for $\beta = 1$, but willing to pay a positive amount for $\beta \in (0, 1)$.

Time preferences: Example of True/False/Uncertain Question 2

Statement: *Fully sophisticated individuals can experience large welfare losses from their present bias.*

True. Why?

- Awareness of present bias (i.e. sophistication) does *not* remove present bias
- Sophisticates that lack commitment devices may still make suboptimal decisions

Time preferences: Example of True/False/Uncertain Question 3

Statement: *Present-biased individuals always have positive demand for commitment devices.*

False. Why?

- Three conditions must be met for positive demand for commitment
  i. Individuals must be present-biased.
  ii. Individuals must be aware of their present-bias (i.e. they can't be fully naive).
  iii. Individuals must perceive the commitment device as effective in helping overcome the self-control problem.
- When only the first is met, we cannot be sure there will be positive demand for commitment.

Risk preferences: expected utility

- What is the expected utility model?
  - What is risk aversion? Why are people risk averse?
  - How is risk aversion modeled in the expected utility model?
  - What is the expected monetary value?
- How can we measure risk aversion within the expected utility model?
  - Certainty equivalents
  - Choices from gambles
  - Insurance choices
- What is problematic about the estimates of risk aversion in the expected utility model?
  - Substantial small-scale risk aversion (high γ)
  - Relatively low large-scale risk aversion (low γ)
  - Expected utility model only has one parameter, can thus not explain both of those features.
  - See Rabin (2000), Rabin & Thaler (2001), and recitation 4

Kahneman and Tversky (1979): Prospect Theory

- What evidence in Kahneman and Tversky (1979) is inconsistent with expected utility?
  - Risk aversion in the gain domain, risk loving in the loss domain
- What are the most important points in Kahneman and Tversky's Prospect Theory (slide 3 of 51 of lecture 9):
  (1) Changes rather than levels
  (2) Loss aversion
  (3) Diminishing sensitivity
- What does the proposed alternative utility (value) function look like? How does it incorporate the three above features?
- How is the reference point determined? What are some candidate reference points? See discussion in recitation 5.

Risk preferences: reference-dependent preferences

- What empirical evidence of loss aversion do we have?
  - Small-scale gambles
  - Endowment effect
  - Applications (lecture 9)

- Applications
  - Labor supply, housing market, stocks, marathon running, golf
  - Be familiar with the empirical applications from lecture 9
  - Understand why reference-dependent preferences can explain (some of) the empirical evidence better than the expected utility model
  - NOT relevant: ~~Deal or No Deal; Pierce et al. (2020)~~ (we did not cover this)

- Solving problems with reference-dependent preferences
  - See pset 3
  - Additional pset and exam questions to practice with

Reference-dependent preferences: Example of Multiple Choice Question 1

Question: *Maddie is writing a problem set for 14.13. She gets utility* $u(q)$ *from the number of questions she writes. She has reference-dependent preferences around her goal of writing 10 questions (her reference point). Normalize* $u(10) = 0$. *Which of the following would be consistent with loss aversion?*

(a) $u(8) = -2, u(12) = 1$
(b) $u(8) = -2, u(12) = 2$
(c) $u(8) = -1, u(12) = 2$

Answer: (a). Why?

- Loss aversion means losses hurt more than gains help
- With preferences in (a), Maddie would have a utility cost of 2 from falling short of her goal by 2 questions, but only gain 1 util from exceeding her goal by 2 questions.

Reference-dependent preferences: Example of Multiple Choice Question 2

Question: *Maddie is walking home and passes a bakery. Unexpectedly, she decides to buy a pastry. Prior to purchasing the pastry, her maximum willingness to pay for the pastry was p_0. She then runs into Allan who asks to buy the pastry from her. She offers him the lowest price she is willing to accept, p_1. Which of the following comparisons between p_0 and p_1 is consistent with an endowment effect?*

(a) $p_0 > p_1$
(b) $p_0 = p_1$
(c) $p_0 < p_1$

Answer: (c). Why?

- Consistent with an endowment effect, $p_0 < p_1$ implies Maddie values the pastry more after she has bought it than prior to buying it.

Social preferences

- What are social preferences?

- How can we measure social preferences?
  - Dictator Game
  - Ultimatum Game
  - Trust Game

- What evidence do we typically find in dictator and ultimatum games?

- Are people genuinely nice to others (because of pure altruism)? Why not?
  - Costly exit in dictator games
  - Hiding behind the computer
  - Moral wiggle room

- We will NOT ask you about models that estimate social preferences (this will be in pset 4).

Social preferences: Example of True/False/Uncertain Question

Statement: *if a person gives 0 in a dictator game, this is evidence that this person is selfish.*

Uncertain. Why?

- The person might give 0 to the other person in the dictator game and then donate the money to someone in greater need.
- The person might be very poor (relative to the other person in the game), so her marginal utility is very high.

Time preferences: Example of Long Question: Laptop Policies

- Assume 14.13 students are present biased with $\beta < 1$ and $\delta = 1$. All students have the same $\beta < 1$ and $\delta = 1$ but differ in the value they derive from using laptops in class, L.

- L is constant for each student from class to class but uniformly distributed across students on the interval [0,1].

- Each lecture generates no immediate utility, but does give a future benefit V. Using a laptop reduces the long-run benefit by D. Both V and D are the same for all students.

- In summary, a student that uses a laptop in class gets immediate utility L and future (undiscounted) utility $V - D$. A student that does not use a laptop gets immediate utility 0 and future (undiscounted) utility V.

- The social planner is not present biased and seeks to maximize the utility of 14.13 students.

Long Question: Part 1

Show that students are just indifferent between using and not using their laptop in the current class if $L = \beta D$. Explain why students with lower values of L (i.e. $L < \beta D$) don't use laptops in class, while students with higher values of L (i.e. $L > \beta D$) do use laptops in class.

Long Question: Solution, Part 1

- Utilities from the two choices are:

$$\begin{aligned} U(\text{laptop}) &= L + \beta(V - D) \\ U(\text{no laptop}) &= 0 + \beta V \end{aligned}$$

- For students that are indifferent, $U(\text{laptop}) = U(\text{no laptop})$. This gives:

$$\begin{aligned} L + \beta(V - D) &= 0 + \beta V \\ L &= \beta D \end{aligned}$$

- Students that choose not to use laptops will have low valuations, L, of using laptops, while students that choose to use laptops will have higher L. Given the indifference condition:
    - Students that do not use laptops: $L < \beta D$
    - Students that use laptops: $L > \beta D$

Long Question: Part 2

Now consider the policy that allows students to use laptops only if they sign up in advance to sit in a laptop section. Why is $L \geq D$, not $L \geq \beta D$, the threshold for opting into the laptop section?

Long Question: Solution, Part 2

- Students now compare:

$$
\begin{aligned}
U(\text{laptop}) &= 0 + \beta(L + V - D) \\
U(\text{no laptop}) &= 0 + \beta V
\end{aligned}
$$

- The threshold for opting in is defined by $U(\text{laptop}) \geq U(\text{no laptop})$. This gives:

$$
\begin{aligned}
0 + \beta(L + V - D) &\geq 0 + \beta V \\
L &\geq D
\end{aligned}
$$

- The threshold changes from βD to D because when laptop use can only happen in the future, all benefits and costs are discounted at the same rate, β.

Long Question: Part 3

Assume there is no laptop policy. Show that if $\beta D < L < D$, the student engages in preference reversals: she prefers not to use the laptop in future classes, but changes her mind when she's actually sitting in those future classes.

Long Question: Solution, Part 3

- When thinking about future laptop use, the student's problem is identical to the problem in part (2). Why?
    - Because she discounts time both one and two periods in advance by β
- We know from part (2) that if $L < D$, she would like to not use the laptop
- But from part (1), we know that if $\beta D < L$, she will end up using the laptop when she's actually sitting in the future class
- This implies a preference reversal: she prefers not to use the laptop in future classes, but switches her mind when she's actually sitting in those future classes.

Long Question: Part 4

Explain why fraction $1 - \beta D$ of the class uses a laptop in part 1, but fraction $1 - D$ of the class uses a laptop in part 2. Why does a smaller share of the class use their laptops in part 2?

Long Question: Solution, Part 4

- In part 1, a student uses a laptop if $L > \beta D$. Define $F(.)$ as the CDF of L.
- Given the uniform distribution:

$$\begin{aligned} P(L > \beta D) &= 1 - F(\beta D) \\ &= 1 - \beta D \end{aligned}$$

- Likewise, in part 2, a student uses a laptop if $L > D$. We have:

$$\begin{aligned} P(L > D) &= 1 - F(D) \\ &= 1 - D \end{aligned}$$

- A smaller share will use laptops in part 2 because the benefit of using a laptop is delayed and hence discounted by β

Long Question: Part 5

Why would the social planner prefer the opt-in policy to both the policy of allowing students to choose whether to use their laptops and to banning lap tops altogether?

Long Question: Solution, Part 5

- The planner is not present biased so would want only students with $L > D$ to use laptops; the opt-in policy achieves this.
- Under the free choice policy, students with $\beta D < L < D$ will sub-optimally use their laptops.
- On the other hand, banning laptops altogether is suboptimal because welfare is gained by allowing the students with the highest valuations, $L > D$, to use laptops

The End

- Don't worry too much about the exam – try your best and you'll do great!
- And even if you don't do great, you'll be fine!

Recitation 8: Bayesian Learning

Maddie McKelway

Plan for Recitation

1. Review Chetty et al. (2009) derivation from lecture
2. Bayesian Learning
3. Deviations from Bayesian Learning

Inattention to taxes: Chetty et al. (2009)

- Taxes not featured in price are likely to be ignored
  - Sales tax only added at the register
- Demand $D(\hat{V})$ is a function of perceived value $\hat{V}$
  - Visible part of the value $v = x - p$, where x reflects how much you like the good and p is its price
  - Less visible (opaque) part $o = -tp$, where t is the tax rate
  - $\hat{V} = v + (1 - \theta)o = x - p - (1 - \theta)\, tp$
  - Note that $\frac{dD}{d\hat{V}} > 0$ (and therefore $\frac{dD}{dp} < 0$)
  - Below focus on opaque part of $\hat{V}$ so write $\hat{V} = v - (1 - \theta)\, tp$

Effect of making the tax fully salient

- Would like to compute the change in log demand when θ falls to 0

$$\Delta \log D(\hat{V}) = \log D\left[v - tp\right] - \log D\left[v - (1-\theta)\, tp\right]$$

- Note that for any $f(x)$, $f(x+\alpha) \approx f(x) + \alpha f'(x)$
  - Equivalently, $f(x+\alpha) - f(x) \approx \alpha f'(x)$
  - Let $f(.) = \log D(.)$, $x = v - (1-\theta)\, tp$, and $\alpha = -\theta tp$
  - Then right-hand side above is $f(x+\alpha) - f(x)$, which is $\approx \alpha f'(x)$
- This gives:

$$\begin{aligned} \Delta \log D(\hat{V}) &= \log D\left[v - tp\right] - \log D\left[v - (1-\theta)\, tp\right] \\ &\approx -\theta tp \cdot \frac{d \log D[v - (1-\theta)tp]}{d\theta} \end{aligned}$$

Effect of making the tax fully salient

- Next, note that: $\frac{d \log Y(t)}{dt} = \frac{dY(t)}{Y(t)}$
- This means:

$$
\begin{aligned}
\Delta \log D(\hat{V}) &= \log D\left[v - tp\right] - \log D\left[v - (1-\theta)\, tp\right] \\
&\approx -\theta tp \cdot \frac{d \log D[v-(1-\theta)tp]}{d\theta} \\
&= -\theta tp * \frac{D'\left[v-(1-\theta)\, tp\right]}{D\left[v-(1-\theta)\, tp\right]}
\end{aligned}
$$

Effect of making the tax fully salient

- Finally, define the price elasticity of demand $\eta_{D,p}$ as $-\frac{p}{D} \cdot \frac{dD}{dp}$
- This gives:

$$\begin{aligned}
\Delta \log D(\hat{V}) &= \log D\left[v - tp\right] - \log D\left[v - (1-\theta)\, tp\right] \\
&\approx -\theta tp \cdot \frac{d \log D[v - (1-\theta) tp]}{d\theta} \\
&= -\theta tp * \frac{D'\left[v - (1-\theta)\, tp\right]}{D\left[v - (1-\theta)\, tp\right]} \\
&= -\theta t * \eta_{D,p}
\end{aligned}$$

- This implies $\theta = \frac{-\Delta \log D(\hat{V})}{t * \eta_{D,p}}$
  - Chetty et al. (2009) try to measure this

Bayesian Learning: Overview

- Almost all economic decisions are undertaken with some degree of uncertainty
- Individuals must make decisions based on perceived likelihoods of outcomes
- How do individuals form beliefs about statistical likelihoods?
    - Bayesian learning: the "statistically correct" way to form beliefs
    - In reality, we see systematic deviations from Bayesian learning
        - Base rate neglect
        - Gamblers' fallacy

Bayesian Learning: Overview

- Set-up:
  1. Individual has a *prior* belief of the likelihood that something is true
  2. Individual observes a *signal* in the world that is indicative of whether it's true
  3. Individual combines her prior and signal to form a *posterior* belief of the likelihood that it's true
- How should (in a statistical sense) the individual combine her prior and signal to form a posterior?
  - Use Bayes' Rule!

Bayes' Rule

Notation

- Individual has some hypothesis, h
- Her *prior* belief is that h is true with probability $P(h)$
- She observes *signal*, D, that provides information about the likelihood that h is true
- She forms a *posterior* belief about the probability h is true given D: $P(h|D)$

How does she form $P(h|D)$?

- Bayes' Rule: $P(h|D) = \frac{P(D|h) \cdot P(h)}{P(D)}$
- Where does Bayes' Rule come from?
  - We know $P(h|D) = \frac{P(h \cap D)}{P(D)}$, which implies $P(h \cap D) = P(h|D) \cdot P(D)$
  - Similarly, $P(D|h) = \frac{P(h \cap D)}{P(h)}$, implying $P(h \cap D) = P(D|h) \cdot P(h)$
  - Equating the two expressions for $P(h \cap D)$ gives $P(h|D) \cdot P(D) = P(D|h) \cdot P(h)$ or $P(h|D) = \frac{P(D|h) \cdot P(h)}{P(D)}$

Bayes' Rule Example

Suppose there are two urns: (a) one with an equal number of black and white balls, and (b) one with 75% black balls and 25% white balls. We pick one urn at random and draw a ball at random. The ball drawn in black. What is the probability that we were drawing from urn (a)?

- Should it be greater than, equal to, or less than 0.5?

Notation

- $h =$ ball is from urn (a)
- $D =$ black ball drawn
- We would like $P(h|D) = \frac{P(D|h) \cdot P(h)}{P(D)}$

Bayes' Rule Example

We would like $P(h|D) = \frac{P(D|h) \cdot P(h)}{P(D)}$

- $P(h)$ = the probability the ball comes from urn (a) before we observe the ball's color (the prior probability)
  - What value does $P(h)$ take? 0.5
- $P(D|h)$ = the probability a black ball is drawn if drawing from urn (a)
  - What value does $P(D|h)$ take? 0.5
- $P(D)$ = the probability a black ball is drawn
  - Law of total probability = the probability of an outcome occuring is equal to the sum of probabilities of every distinct way it can occur
  - $P(D) = P(D \cap h) + P(D \cap h') = P(D|h)P(h) + P(D|h')P(h') = (0.50)(0.50) + (0.75)(0.50)$
- Combining: $P(h|D) = \frac{P(D|h) \cdot P(h)}{P(D)} = \frac{(0.5)(0.5)}{(0.50)(0.50)+(0.75)(0.50)} = 0.4$

Base-Rate Neglect

- A common behavioral deviation from Bayesian learning
- *One in a hundred people have HIV, and we have a test for HIV that is 99% accurate. If a person tested positive, what's the probability that she has HIV?*
  - Most people answer 99%
  - Bayes' Rule provides a different answer

Bayes' Rule

- Notation: P=HIV-positive; N=HIV-negative; p=tested positive
- We would like to know $P(P|p) = \frac{P(p|P)P(P)}{P(p)}$
  - $P(p|P)P(P) = (0.99)(0.01)$
  - $P(p) = P(p \cap P) + P(p \cap N) = P(p|P)P(P) + P(p|N)P(N) = (0.99)(0.01) + (0.01)(0.99)$
- This implies: $P(P|p) = \frac{(0.99)(0.01)}{(0.99)(0.01)+(0.01)(0.99)} = 0.5 \neq 0.99$

Base Rate Neglect

- Base Rate Neglect: when given base rate information (i.e. information pertaining to everyone) and specific information (i.e. information pertaining to a particular individual), people focus on the latter and ignore the former
- In the HIV example, people see positive test results (specific information) and forget to account for the fact that HIV is unlikely in the first place (base rate information)
- Implies a deviation from Bayes' Rule

The Gambler's Fallacy

- Another common behavioral deviation from Bayesian learning
- *You toss a coin 20 times. The first 19 times are tails. What's the probability that the final toss is also tails?*
  - Some people might say the probability is very low, reasoning that you've just seen a lot of tails so it would be very unlikely to see another
  - Bayes' Rule gives probability of $\frac{1}{2}$

Bayes' Rule

- Notation: T = the final toss is tails, T_{19} = the first 19 tosses were tails
- Bayes' Rule gives $P(T|T_{19}) = \frac{P(T_{19}|T)P(T)}{P(T_{19})}$
  - Start with $P(T_{19})$, the probability the first 19 draws are tails
    - What value does $P(T_{19})$ take? $\frac{1}{2^{19}}$
  - $P(T_{19}|T)$: the probability that the first 19 draws are tails given the last one is tails
    - Tricky: the last draw being tails tells us nothing about the likelihood that the first 19 were tails
    - The outcomes are independent so $P(T_{19}|T) = P(T_{19}) = \frac{1}{2^{19}}$
  - $P(T)$: probability that the last toss is tails prior to observing the first 19 tosses
    - What value does $P(T)$ take? 0.5
- The $\frac{1}{2^{19}}$'s cancel and we are left with $P(T|T_{19}) = P(T) = 0.5$
  - Intuitively: the signal contains no information so we should stick with our prior
- We didn't have to use Bayes' rule to get this (though going through it is good practice!): could instead have noted that independence means $P(T|T_{19}) = P(T) = 0.5$

The Gambler's Fallacy

- The Gambler's Fallacy: the belief that an event occuring frequently in the past means it's less likely to occur in the future when in fact the occurrences of the event in the past and in the future are independent
- In the coin toss example, many people don't internalize the independence between the last coin toss and the first 19; after they see 19 tails, they think it's very unlikely the 20th would also be tails
- Chen, Moskowitz, & Shue (2016): evidence of the Gambler's Fallacy in high-stakes, real-world decisions
    - Study decisions of asylum judges, loan officers, and baseball umpires
    - Find negative autocorrelation (what does this mean?) of decisions that is unrelated to case quality

Recitation 9: Social Preferences

Aaron Goodman

Outline

1. Hjort (2014)
2. Problem Set 4

Hjort (2014)

- Another good example of field research on social preferences in the workplace
- Complements our discussion in lecture of Bandiera et al. (2005), Beza et al. (2018), Rao (2019), Lowe (2019)
- Highlights the importance of employers' compensation and personnel policies when workers have social preferences

Setting

- Flower packaging plant in Kenya
- Workers are drawn from two rival tribes (Kikuyu and Luo)
- Workers must collaborate in teams of three to produce packages of flowers
- One "supplier" prepares roses and passes them to two downstream "processors"

4

Production Teams

Figure 1a: Organization of team production

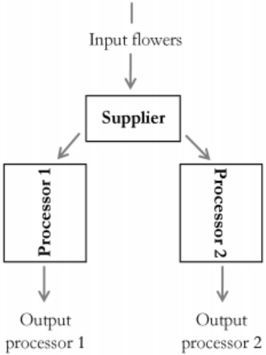

5

Possible Team Configurations

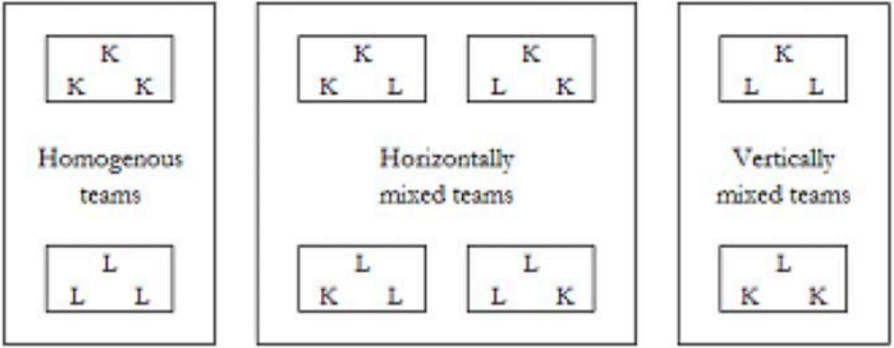

Compensation Policy and Timeline of Events

Initial compensation policy at beginning of sample period:

- Suppliers are paid a piece rate w
- Processors are paid a piece rate $2w$

December 2007:

- Presidential election takes place
- Leads to political and violent conflict between the tribes
- Firm continues to operate

February 2008:

- Firm changes its compensation policy for processors
- Processors are now paid w per package produced by the team, rather than $2w$ per package produced individually

Simple Model

- Let y denote income and e denote effort
- Let s denote the supplier, p_1 denote the first processor, and p_2 denote the second processor
- Allow the supplier have social preferences:
  - Attaches weight α_y to utility of processors from the same tribe
  - Attaches weight α_n to utility of processors from a different tribe
- Assume for simplicity that the processors do not have social preferences

Supplier Utility

Supplier's utility given by:

$$u(y_s, e_s) + \alpha_1 u(y_{p_1}, e_{p_1}) + \alpha_2 u(y_{p_2}, e_{p_2}),$$

where

$$\alpha_i = \begin{cases} \alpha_y & \text{if processor } i \text{ is from same tribe} \\ \alpha_n & \text{if processor } i \text{ is from different tribe} \end{cases}$$

Effects of the Election and Compensation Change

- Within the model, how might we account for the heightened conflict caused by the presidential election?
- How do we think the presidential election would affect the productivity of:
  - Homogenous teams?
  - Horizontally mixed teams?
  - Vertically mixed teams?
- How do we expect the ensuing compensation change to affect:
  - Homogenous teams?
  - Horizontally mixed teams?
  - Vertically mixed teams?

Observed Effects

Figure 2: Output in homogeneous and mixed teams across time

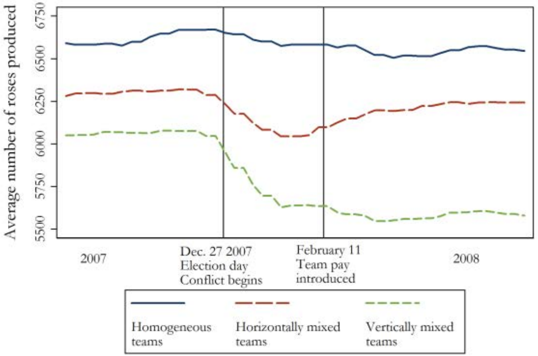

Hjort (2014): What Did We Learn?

- Workers have social preferences
- Compensation policies interact with social preferences; employers' optimal compensation policies depend on their workers' preferences
- Employers can also affect productivity with non-compensation personnel policies:
  - What if the firm reassigned its workers so that all teams were homogenous?
  - Short-run vs. long-run effects of worker segregation?

Outline

Problem Set 4

- With just one paper to cover in recitation this week, we thought it would be helpful to address any questions and talk through the general approach to each part
- Any particular questions?

Part 1: General Approach

- Workers have utility

$$u_i(y_i, q_i) = y_i - c(q_i) + \alpha \sum_{j \neq i} u_j(y_j, q_j)$$

- Workers can be:
  - Selfish: $\alpha = 0$
  - Altruistic: $\alpha > 0$
- Compensation can be:
  - Piece-rate: $y_i = pq_i$
  - Relative: $y_i = pq_i - \gamma \sum_{j \neq i} \frac{q_j}{N-1}$
- So four possible cases. Before doing any math:
  - Should we expect the workers' optimal effort to be different in each of the four cases?
  - If not, which subset(s) of the four cases have the same solutions?

Part 2: Setup

- Alex's payoff is x_1 and Aaron' payoff is x_2. Aaron's utility is:

$$u_2(x_1, x_2) = \begin{cases} \rho x_1 + (1-\rho)x_2 & \text{if } x_2 \geq x_1 \\ \sigma x_1 + (1-\sigma)x_2 & \text{if } x_2 < x_1 \end{cases}$$

Part 2: General Approach

How do we interpret ρ and σ?

- $\sigma \leq \rho < 0$
  - Simple competitive preferences; Aaron's utility always increasing in his own payoff and always decreasing in Alex's payoff.
  - Aaron becomes more competitive when his own payoff is smaller than Alex's.
- $\sigma < 0 < \rho < 1$
  - Aaron becomes altruistic only when his own payoff is larger than Alex's.
- $0 < \sigma \leq \rho \leq 1$
  - "Social-welfare preferences" (Charness and Rabin 2002): Aaron's utility is always increasing in both his and Alex's payoff.
  - Aaron cares more about Alex's payoff when his own payoff is larger than Alex's.
- $\sigma = \rho = 0$
  - Simple self-interest; Alex's payoff never matters to Aaron.

14.13: Psychology and Economics
Spring 2020

Recitation 10: Projection Bias and Attribution Bias

Alex Olssen and Will Rafey

Outline

1. Projection Bias

2. Attribution Bias

Overview

- Understand the influence of future vs. past states
  - Projection bias: mis–prediction of influence of **future** states
  - Attribution bias: mis–prediction of influence of **past** states

Outline

1. Projection Bias
2. Attribution Bias

Projection Bias

Image by Jeff on flickr. CC BY-NC-SA

Image by harry_nl on flickr. CC BY-NC-SA

Projection Bias: Model

- True utility at time t is $u(c_t, s_t)$
- True utility depends on consumption c_t and the state s_t at time t
  - State could be anything that affects utility from consumption, e.g., level of hunger or addiction
- At time t, predict future utility from consuming c_τ in state s_τ at time $\tau > t$:

$$\hat{u}(c_\tau, s_\tau) = (1 - \alpha)u(c_\tau, s_\tau) + \alpha u(c_\tau, s_t)$$

- $\alpha \in [0, 1]$ is the degree of projection bias.

Projection Bias: Example Problem

- Maddie's dog, Emma, suffers from severe projection bias. Emma's true utility at hour t, $u(c_t, s_t)$

$$\begin{aligned} u(c_t = whine, s_t = hungry) &= -3, \\ u(c_t = don't\ whine, s_t = hungry) &= -10, \\ u(c_t = whine, s_t = not\ hungry) &= -5, \\ u(c_t = don't\ whine, s_t = not\ hungry) &= 0 \end{aligned}$$

- At time t, Emma predicts her future utility of whining at time $\tau > t$ to be

$$\hat{u}_t(c_\tau, s_\tau) = (1-\alpha)u(c_\tau, s_\tau) + \alpha u(c_\tau, s_t),$$

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

Projection Bias: α Parameter

- What does α measure?
  - α measures the degree of projection bias
- What does it mean for α to equal 0?
  - $\alpha = 0$ means there is no projection bias (rational expectations)
- What does it mean for α to equal 1?
  - $\alpha = 1$ means there is full projection bias

Projection Bias: Predicted Utility

- Suppose $\alpha = \frac{3}{4}$. Consider time period $\tau > t$ when Emma will be hungry $s_\tau = \textit{hungry}$
- How much utility will Emma expect to get from whining if she is hungry in period t?

$$\begin{aligned}
\hat{u}_t(c_\tau, s_\tau) &= (1-\alpha)u(c_\tau, s_\tau) + \alpha u(c_\tau, s_t) \\
\hat{u}_t(\textit{whine}, \textit{hungry}) &= (1-\alpha)u(\textit{whine}, \textit{hungry}) + \alpha u(\textit{whine}, \textit{hungry}) \\
\hat{u}_t(\textit{whine}, \textit{hungry}) &= -3(\frac{1}{4}) + -3(\frac{3}{4}) = -3
\end{aligned}$$

- How much utility will Emma expect to get from whining if she is not hungry in period t?

$$\begin{aligned}
\hat{u}_t(c_\tau, s_\tau) &= (1-\alpha)u(c_\tau, s_\tau) + \alpha u(c_\tau, s_t) \\
\hat{u}_t(\textit{whine}, \textit{nothungry}) &= (1-\alpha)u(\textit{whine}, \textit{hungry}) + \alpha u(\textit{whine}, \textit{not\ hungry}) \\
\hat{u}_t(\textit{whine}, \textit{nothungry}) &= -3(\frac{1}{4}) + -5(\frac{3}{4}) = -4.5
\end{aligned}$$

- Do the utilities differ? Why?
- Yes. If she is hungry at t, she correctly predict future utility. If she is not hungry at t, she incorrectly “projects” her lo... utilit... from whining ... to the future

Outline

1 Projection Bias

2 Attribution Bias

Attribution Bias

- Definition
  - When judging the value of a good, people are overly influenced by the state in which they previously consumed it
- Examples
  - More likely to return to a restaurant first tried when hungry
  - More likely to negatively rate a movie seen while tired
  - Less likely to recommend a zoo to a friend if it rained during last visit

Attribution Bias: Model

- Predict utility of consuming c_t while in state s_t, given prior consumption experience of c_t was in state $s_{t-1}(c_t)$:

$$\hat{u}(c_t, s_t) = (1-\gamma)u(c_t, s_t) + \gamma u(c_t, s_{t-1}(c_t))$$

  where $\gamma \in [0, 1]$ is the degree of attribution bias.
- Recall in the Projection bias model $\tau > t$ and

$$\hat{u}(c_\tau, s_\tau) = (1-\alpha)u(c_\tau, s_\tau) + \alpha u(c_\tau, s_t)$$

- $\alpha \in [0, 1]$ is the degree of projection bias.

Attribution Bias: Experimental Evidence (Haggag et al. 2019)

- Basic structure
  - Randomly manipulate the thirst levels of participants
  - Have participants drink a new mixed drink
  - Elicit preference measures while in a later neutral state
- Suppose participants who were thirsty say the like the drink better in the neutral state
- Why can we take this as evidence of attribution bias? I.e., how do we know the state of prior consumption explains the difference?

Experiment outline

1. Manipulation of thirst
   - Treatment group: drink 3 cups of water
   - Control group: drink 1/2 cup of water
2. Answer demographic questions, measure of current thirst (7-pt scale)
3. Stir together ingredients for new mixed drink:
   - 1 cup milk
   - 1/3 cup orange juice
   - 1 tablespoon sugar
4. Consume the mixed drink
5. Answer "how enjoyable was drinking the mixed drink?" (7-pt scale)

Ingredient photo upload

15

Summary of experiment

- Experiment manipulated state (thirst) during new experience (drink)
- Thirsty individuals like the drink better.
- Since the two groups were randomized, we know that any difference must be due to the differential thirst across groups.
- In the absence of attribution bias, people's willingness to have the same drink again shouldn't depend on thirst when first consuming it.
- But people who were thirsty during their first experience are more willing to have the drink again.

Attribution Bias Example Problem

- Maddie's true utility from exercising on day t, $u(c_t, s_t)$ depends not only on the type of exercise c_t (run or gym), but also on the state s_t (hot or cool):

$$
\begin{aligned}
u(c_t = run, s_t = hot) &= 5, \\
u(c_t = gym, s_t = hot) &= 6, \\
u(c_t = run, s_t = cool) &= 10, \\
u(c_t = gym, s_t = cool) &= 7.
\end{aligned}
$$

Attribution Bias: Predicted Utility

- Fall arrives and the weather in Cambridge today is $s_t = cool$
- Suppose Maddie most recently went to the gym when it was hot and ran when it was hot
- Will she choose to run or go to the gym?
  - Maddie will choose the form of exercise that gives the higher predicted utility.

$$\begin{aligned}
\hat{u}_t(c_t, s_t) &= (1-\gamma)u(c_t, s_t) + \gamma u(c_t, s_{t-1}(c_t)) \\
\hat{u}_t(run, cool) &= (1-\gamma)u(run, cool) + \gamma u(run, hot) \\
\hat{u}_t(run, cool) &= 10(1-\gamma) + 5\gamma \\
\hat{u}_t(run, cool) &= 10 - 5\gamma
\end{aligned}$$

$$\begin{aligned}
\hat{u}_t(c_t, s_t) &= (1-\gamma)u(c_t, s_t) + \gamma u(c_t, s_{t-1}(c_t)) \\
\hat{u}_t(gym, cool) &= (1-\gamma)u(gym, cool) + \gamma u(gym, hot) \\
\hat{u}_t(gym, cool) &= 7(1-\gamma) + 6\gamma \\
\hat{u}_t(gym, cool) &= 7 - \gamma
\end{aligned}$$

18

  - Maddie will run if $10 - 5\gamma > 7 - \gamma$ or if $\gamma < 3/4$

Attribution Bias: Predicted Utility II

- Suppose Maddie most recently went to the gym when it was hot and ran when it was cool
- Will she choose to run or go to the gym?
  - Maddie will choose the form of exercise that gives the higher predicted utility.

$$\begin{aligned}
\hat{u}_t(c_t, s_t) &= (1-\gamma)u(c_t, s_t) + \gamma u(c_t, s_{t-1}(c_t)) \\
\hat{u}_t(run, cool) &= (1-\gamma)u(run, cool) + \gamma u(run, cool) = u(run, cool) = 10
\end{aligned}$$

$$\begin{aligned}
\hat{u}_t(c_t, s_t) &= (1-\gamma)u(c_t, s_t) + \gamma u(c_t, s_{t-1}(c_t)) \\
\hat{u}_t(gym, cool) &= (1-\gamma)u(gym, cool) + \gamma u(gym, hot) \\
\hat{u}_t(gym, cool) &= 7(1-\gamma) + 6\gamma \\
\hat{u}_t(gym, cool) &= 7 - \gamma
\end{aligned}$$

  - Maddie will run if $10 > 7 - \gamma$ or if $\gamma > -3$. This is always true.
  - Intuition: Maddie last ran when it was cool (the best time to run). She last went to the gym when it was hot (the worst time to go to the gym). Her attribution bias works in favor of running (which is her preferred choice when it is cool under her true utility function anyway)

Attribution Bias: Predicted Utility III

- Suppose Maddie most recently went to the gym when it was cool and ran when it was cool
- Will she choose to run or go to the gym?
  - Maddie will choose the form of exercise that gives the higher predicted utility.

$$\begin{aligned}\hat{u}_t(c_t, s_t) &= (1-\gamma)u(c_t, s_t) + \gamma u(c_t, s_{t-1}(c_t)) \\ \hat{u}_t(run, cool) &= (1-\gamma)u(run, cool) + \gamma u(run, cool) = u(run, cool) = 10\end{aligned}$$

$$\begin{aligned}\hat{u}_t(c_t, s_t) &= (1-\gamma)u(c_t, s_t) + \gamma u(c_t, s_{t-1}(c_t)) \\ \hat{u}_t(gym, cool) &= (1-\gamma)u(gym, cool) + \gamma u(gym, cool) = u(gym, cool) = 7\end{aligned}$$

  - In this case, attribution bias has no effect because the current state is the same as the state for prior consumption
  - Thus γ is irrelevant

Attribution Bias: Wrong Choices

- In which of the three cases above and for which values of γ might we say Maddie makes the "wrong" choice?
- When it is cool, best choice is to run
- First case
    - Her prior run and gym visit were both when it was hot
    - Her predicted utility is biased because of attribution bias
    - If $\gamma \geq 3/4$, then she will go to the gym (i.e., when she is close to full attribution bias)
- Second case
    - Her predicted utility makes running even more preferable (because her past experiences running were good and at the gym were bad)
    - She runs for any $\gamma \in [0, 1]$
- Third case
    - Her predicted utility happens to have no bias (because the current state is the same as the prior ones)
    - She runs for any $\gamma \in [0, 1]$