9 Intertemporal Tradeoffs

9.1 Motivating Example: Saving for Retirement

Consider an individual who lives for two periods: $t = 0$ and $t = 1$. Period $t = 0$ might be youth and period $t = 1$ might be old age. In period $t = 0$, the individual receives income $w > 0$; she needs to decide how much to consume in period $t = 0$ and how much to save for (consumption in) period $t = 1$.

Funds not spent in Period 0 earn interest at rate $r > 0$. Thus, her budget constraint limiting consumption in period 1 is:

$$c_1 \leq (w - c_o)(1 + r) \tag{1}$$

where c_0 indicates consumption at $t = 0$ and c_1 denotes consumption at $t = 1$. Rearranging the terms, we can also express the budget constraint as:

$$c_0 + \left(\frac{1}{1+r}\right) c_1 \leq w. \tag{2}$$

This formulation highlights the fact that the implicit price of a dollar of consumption in Period 1 is $1/(1 + r)$ dollars of foregone consumption in Period 0.

Suppose individual utility from consumption in any period is $u(c_t)$ where $u'(c_t) > 0$ and $u''(c_t) < 0$. We will refer to this utility function, which characterizes the utility from consumption in Period t, as the **instantaneous utility function**. We will assume that this utility function is the same in all periods. If we further assume that overall utility is the sum of instantaneous utility in each period, then we can write down our decision-maker's utility function as

$$U(c_0, c_1) = u(c_0) + u(c_1). \tag{3}$$

To find our decision-maker's optimal consumption and savings plan, we can set up a Lagrangian:

$$\mathcal{L} = u(c_0) + u(c_1) + \lambda \left[w - c_0 - \left(\frac{1}{1+r}\right) c_1\right] \tag{4}$$

The resulting first-order conditions are:

1. $\frac{\partial \mathcal{L}}{\partial c_0} = 0 \Leftrightarrow u'(c_0) = \lambda$
2. $\frac{\partial \mathcal{L}}{\partial c_1} = 0 \Leftrightarrow u'(c_1) = \lambda \left(\frac{1}{1+r}\right)$

3. $\frac{\partial \mathcal{L}}{\partial \lambda} = 0 \Leftrightarrow c_0 + \left(\frac{1}{1+r}\right) c_1 = w$

Combining (1) and (2), we see that at the optimum

$$u'(c_0) = (1+r)u'(c_1) \tag{5}$$

which means that $c_0^* < c_1^*$. A utility-maximizing decision-maker who weights periods $t = 0$ and $t = 1$ equally and can save at a positive real interest rate will shift consumption toward later periods. This pattern would be even more pronounced if we extended the model to include additional periods.

Is such a model realistic? Do people really save most of their income to realize the future benefits of earning interest? Conventional wisdom suggests that this is not the case: for instance, more than half of all people who win the lottery spend all of their winnings within five years.[1] How can we extend our model of intertemporal tradeoffs to account for the fact that people do not tend to save very much?

9.2 Why Do People Save?

- Reasons to care about the future:
  - We have the psychological ability to anticipate future needs and wants
  - We can increase our overall wellbeing by making forward-looking investments and avoiding activities that will have large future costs
  - We know that we will need to eat in the future, so it may make sense to engage in precautionary saving and other types of planning behavior
  - Anticipatory utility: we can be happy about the future in the present
- Reasons to care **more** about the present:
  - The future is uncertain: will we even be here?
    - An expected utility maximizer should weight future payoffs by the probability that she will be around to enjoy them
  - Consumption now is pleasant, and self control can be psychologically costly

[1] `https://www.rd.com/list/13-things-lottery-winners/`

- We experience **affective states** in the present (e.g. hunger, fatigue, excitement) that can make immediate wants feel especially important

9.3 Discounting

- A **discount factor** is a weight, $\delta \leq 1$, placed on future consumption or utility. In our simple, two-period model, we might think of overall utility from consumption in periods $t = 0$ and $t = 1$ as:

$$U(c_0, c_1) = u(c_0) + \delta u(c_1) \tag{6}$$

  which only differs from Equation 4 only because we have (down)weighted future utility by the factor δ.

- Including δ changes the derivative of the Lagrangian with respect to c_1 to:

$$\delta u'(c_1) = \lambda \left(\frac{1}{1+r} \right) \tag{7}$$

- The consequence of this is that, at the optimum, it is now the case that:

$$u'(c_0) = \delta(1+r)u'(c_1) \tag{8}$$

  so if $\delta = 1/(1+r)$, a utility-maximizing individual would set $c_0 = c_1$.

- This is one motivation for including a discount factor in our models of intertemporal tradeoffs: it counteracts the impact that a positive real interest rate should have on savings behavior, leading to a more satisfying prediction that an optimizing an individual should seek to equalize consumption across periods (rather than consuming more and more over time because of excessive savings)

- Reflects uncertainty about the future

- $\delta < 1$ is also more consistent with many people's lived experience: people tend to save "too little" rather than "too much" (though such normative statements should probably be treated with some caution).

9.4 The Discounted Utility Model

- How can we extend this model ot more than two periods? Consider extending our simple example to three periods: $t = 0$, $t = 1$, and $t = 2$.

- How much weight should a forward-looking decision maker who is making a consumption decision in Period 0 put on utility in Period 2? What discount factor should they use to weight utility in Period 2?

	$t = 0$	$t = 1$	$t = 2$
discount factor at $t = 0$	1	δ	?
discount factor at $t = 1$	–	1	δ

- What if we extend our parallel between interest rates and the discount factor? How would we view the tradeoff between immediate consumption and consumption **two periods in the future** through that lens?

	$t = 0$	$t = 1$	$t = 2$
value of dollar saved at $t = 0$	1	$1 + r$	$(1 + r)^2$
required to have one dollar at t	1	$\frac{1}{1+r}$	$\frac{1}{(1+r)^2}$

- Budget constraint:

$$c_2 \leq (w - c_o)(1 + r)^2 \Leftrightarrow c_0 + \left(\frac{1}{1 + r}\right)^2 c_2 \leq w \tag{9}$$

- Extending further into the future: if you save x dollars today and your savings earns interest rate r, t periods in the future, you will have $x(1 + r)^t$; if you want to have x dollars t periods in the future, you need to save $x/(1 + r)^t$ today

- Reasons to discount utility t periods in the future using discount factor δ^t:

- The parallel with interest rates: if discounting did not compound the way interest compounds (that is: exponentially), people would end up putting excessive or near-zero weight on the distant future
- It is **dynamically consistent**: the tradeoff I make **today** (at $t = 0$) between tomorrow ($t = 1$) and the next day ($t = 2$) is the same as the tradeoff I will make tomorrow (at $t = 1$) between that day ($t = 1$) and the next day ($t = 2$)
- Any other way of discounting the future would mean that what consumption path is optimal (i.e. preferred, utility-maximizing) depends on when you ask someone to make the decision

◦ In the **discounted utility model** first proposed by Paul Samuelson in 1937, the (total) utility of a stream consumption stream $c = (c_0, c_1, c_2, \ldots, c_t, \ldots, c_T)$ is:

$$\begin{aligned} U(c_0, c_1, c_2, \ldots, c_t, \ldots, c_T) &= u(c_0) + \delta u(c_1) + \delta^2 u(c_2) + \ldots + \delta^T u(c_T) \\ &= \sum_{t=0}^{T} \delta^t u(c_t) \end{aligned} \tag{10}$$

where $u(c_t)$ is an **instantaneous utility function** characterizing realized utility in each period

- δ is the **discount factor**
- $u(c_t)$ does not change over time
- We typically normalize $u(0) = 0$
- Utility is additively separable across periods
- Does the length of the time period matter? Why or why not?

9.4.1 Examples

◦ **Example 1:** discounted utility 100 dollars in 3 days if $u(c_t) = \sqrt{c_t}$:

$$\begin{aligned} U(c_0, c_1, c_2, c_3, c_4, \ldots) &= U(0, 0, 0, 100, 0, \ldots) \\ &= \sqrt{0} + \delta\sqrt{0} + \delta^2\sqrt{0} + \delta^3\sqrt{100} + \delta^4\sqrt{0} + \ldots \\ &= \delta^3 10 \end{aligned} \tag{11}$$

- **Example 2:** z dollars per week until T weeks from now ($T+1$ weeks total)

$$\begin{aligned} U\left(z, z, z, z, z, \ldots, z\right) &= u(z) + \delta u(z) + \delta^2 u(z) + \ldots + \delta^T u(z) \\ &= u(z)\left(1 + \delta + \delta^2 + \ldots + \delta^T\right) \\ &= u(z)\left(\frac{1-\delta^{T+1}}{1-\delta}\right) \end{aligned} \tag{12}$$

  because for $r \neq 1$, the sum of the first $T+1$ terms of the geometric series

$$a + ar + ar^2 + ar^3 + \ldots + ar^T \tag{13}$$

  is given by

$$a\left(\frac{1-r^{T+1}}{1-r}\right) \tag{14}$$

- **Example 3:** z dollars per week forever

$$\begin{aligned} U\left(z, z, z, z, z, \ldots\right) &= u(z) + \delta u(z) + \delta^2 u(z) + \ldots \\ &= u(z)\left(1 + \delta + \delta^2 + \ldots\right) \\ &= u(z)\left(\frac{1}{1-\delta}\right) \end{aligned} \tag{15}$$

  because for $r < 1$, the sum of the (infinite) geometric series

$$a + ar + ar^2 + ar^3 + \ldots \tag{16}$$

  is given by

$$a\left(\frac{1}{1-r}\right) \tag{17}$$

- **Aside:** even though we don't live forever, we often model an series of repeating payments as an infinite (rather than a finite) stream

  - The math is easier!
  - In many situations, T is not particularly relevant (all models are wrong)
  - We may not even know exactly how large T is: e.g. Williams College pays me my wage w twice a month, and I am not usually thinking about how I am going

to retire approximately 250 paychecks from now

- Of course, sometimes T ("cap t") is particularly relevant or salient

– One of our justifications for discounting the future is uncertainty about how long we will be around

– δ^t gets quite low as t gets high

- A **very** patient, risk neutral person who is indifferent between 100 dollars today and 144 dollars in a year will discount payouts ten years from now by more than 97 percent (using a discount factor of approximately 0.026)

– **Example 4:** consider a Williams student deciding whether to go to graduate school. If she does not go to graduate school, she will begin earning $a > 0$ immediately, and she will earn a every period "forever" (at least as far as we are concerned). If she goes to graduate school, she earns nothing for four years. Normalize the utility of earning nothing to zero. Then, starting four years from now, she will earn $b > a$ every year forever. For what values of δ will she prefer to go to graduate school?

$$\begin{aligned} u\left(a, a, a, a, a, a, \ldots\right) &< u\left(0, 0, 0, 0, b, b, \ldots\right) \\ \Leftrightarrow u(a) + \delta u(a) + \delta^2 u(a) + \ldots &< 0 + 0 + 0 + 0 + \delta^4 u(b) + \delta^5 u(b) + \ldots \\ \Leftrightarrow \frac{u(a)}{1-\delta} &< \frac{\delta^4 u(b)}{1-\delta} \\ \Leftrightarrow \delta &> \left[\frac{u(a)}{u(b)}\right]^{1/4} \end{aligned} \tag{18}$$

9.5 Criticisms of the Discounted Utility Model: Present Bias

◦ Samuelson's discounted utility model is **dynamically consistent** in that the optimal (i.e. utility-maximizing) tradeoff between period t and period $t+k$ does not depend on when (i.e. how far in advance) you make your decision – the tradeoff you would make between today and tomorrow is the same as the tradeoff you would make between 998 days from now and 999 days from now.

– Samuelson's discounted utility model is the only model of intertemporal tradeoffs that has this property – a fact which is mathematically desirable though perhaps not psychologically realistic

- Self control is hard, and humans (and even animals including both monkeys and pigeons) often behave in a time-inconsistent or impatient manner
- One way to model this is to discount **all** future payoffs by an additional factor $\beta < 1$

	$t = 0$	$t = 1$	$t = 2$	$t = 3$
discount factor at $t = 0$	1	$\beta\delta$	$\beta\delta^2$	$\beta\delta^3$
discount factor at $t = 1$	–	1	$\beta\delta$	$\beta\delta^2$

  - The tradeoff that you make between (say) $t = 1$ and $t = 2$ now depends on when you make your decision: **present-biased preferences** are not dynamically consistent
  - Consider an action that involves immediate cost $-c$ and future benefit b
  - If you decide at time $t = 1$ whether to take the action immediately, you would do it whenever $c \leq \beta\delta b$
  - What if you have to commit to a future course of action at time $t = 0$? You would prefer to take the action as long as $c \leq \delta b$
  - Whether you do the action depends on when you have to decide
  - Exponential discounting along the lines of he discounted utility model is the only way of discounting the future that is dynamically consistent

- When there are more than two periods, individual beliefs matter: we need to model not only what an individual wants at different points in time (utility), but what she believes – specifically, what she believes she will do (given her tastes) in the future
  - Let $\tilde{\beta}$ represent an individual's belief about her degree of present bias
  - A present-biased individual is **sophisticated** if they fully anticipate their present bias: $\tilde{\beta} = \beta$
  - A present-biased individual is **naive** if they do not anticipate their present bias: $\tilde{\beta} = 1$

  - A present-biased individual is **partially sophisticated** if $\beta < \tilde{\beta} < 1$

- Considering the investment decision (immediate costs, delayed benefits) discussed above: a present-biased individual expects to invest whenever $c \le \tilde{\beta}\delta b$

  - $\beta < \tilde{\beta} \Rightarrow$ they over-estimate how often they will invest
  - $\tilde{\beta} < 1 \Rightarrow$ they don't invest as often as they would like

- **Practice Problem.** What is the most a present-biased individual would be willing to pay at time $t = 0$ to commit to making the investment at time $t = 1$?

- **Practice Problem.** Consider a present-biased individual deciding whether to take an action at time $t = 1$ that yields benefit $b > 0$ (at time $t = 1$) and future cost $c > 0$ (at time $t = 2$)

  - When will the present-biased individual choose to take the action?
  - If they were allowed to commit to a plan at time $t = 0$, when would they like to take the action?
  - When reflecting on future behavior, when do they expect to take the action?

9.5.1 Paying Not to Go to the Gym

- Consider an individual with the discount factor $\delta < 1$ who is deciding whether to purchase a gym membership.

$t = 0$	$t = 1$	$t = 2$
chooses contract	goes to the gym?	health benefits
	$-c$	b

- Costs are uncertain *ex ante*:

  - The decision-maker doesn't know at time $t = 0$ whether she will feel like going to the gym at time $t = 1$

- She only knows the probability distribution of c (which is a random variable)
- If going to the gym were free, she'd go if $\delta b - c \geq 0 \Leftrightarrow c \leq \delta b$
- The probability that she would go to the gym would be $\Pr(c \leq \delta b)$

◦ Two types of gym contracts available:

- Pay fixed fee L up-front at $t = 1$, irrespective of whether you go to the gym
- Pay p if you go to the gym, but nothing otherwise

◦ Structure of the **fixed fee contract**:

	$t = 0$	$t = 1$	$t = 2$
definitely pay	0	$-L$	0
if you go to the gym	0	$-c$	b

◦ Under the fixed fee contract, you go to the gym if $\delta b - c \geq 0 \Leftrightarrow c \leq \delta b$

◦ Probability of gym attendance: $\Pr(c \leq \delta b)$

◦ Since c is a random variable, you don't know how happy you'll be going to the gym

- Expected value of c: $E[c]$
- $E[c|c \leq \delta b]$ is the conditional expectation of c; essentially the average value of c that you'd observe if you only looked at the time when $c \leq \delta b$
- Also useful: $E[a + bx] = a + bE[x]$ if a and b are constants

◦ The fixed-fee contract as a lottery:

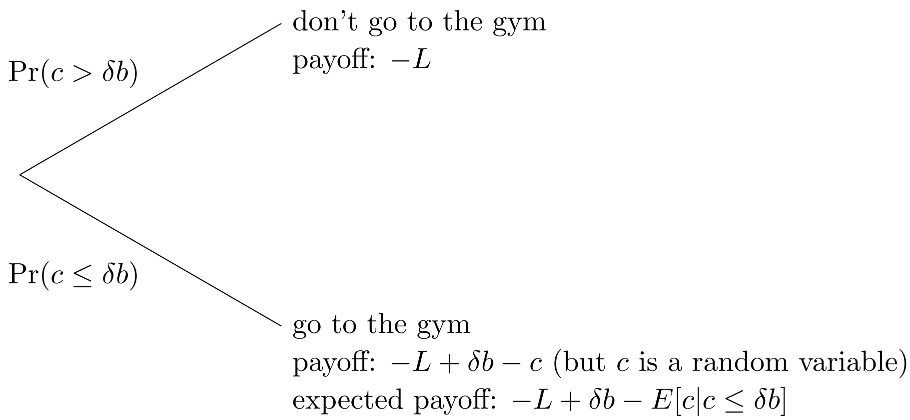

- Expected (discounted) utility of fixed-fee contract (evaluated at time $t = 0$):

$$EU_{FF} = \delta \left\{ -L + \Pr(c \leq \delta b) \left(\delta b - E[c|c \leq \delta b] \right) \right\}$$

- Structure of the **pay-per-visit contract**:

	$t = 0$	$t = 1$	$t = 2$
definitely pay	0	0	0
if you go to the gym		$-c - p$	b

- Under the fixed fee contract, you go to the gym if $\delta b - c - p \geq 0 \Leftrightarrow c \leq \delta b - p$
- Probability of gym attendance: $\Pr(c \leq \delta b - p)$
- Since c is a random variable, you don't know how happy you'll be going to the gym
- The pay-per-visit contract as a lottery:

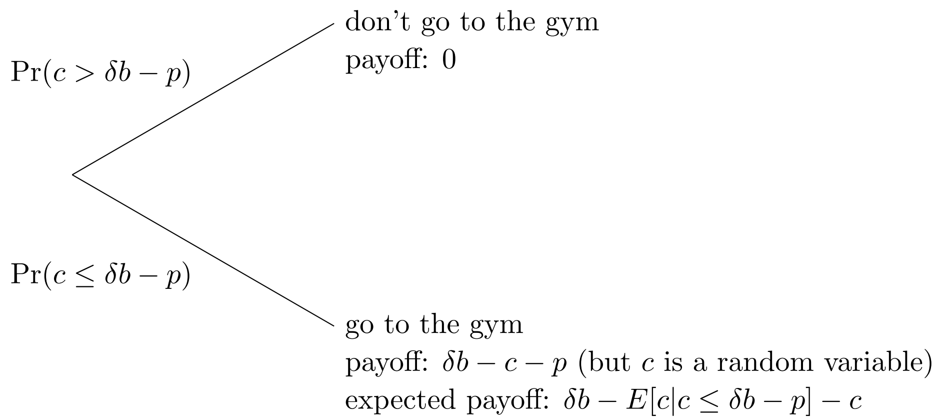

- Expected (discounted) utility of fixed-fee contract (evaluated at time $t = 0$):

$$EU_{PPF} = \delta \left\{ \Pr(c \leq \delta b - p) \left(\delta b - E[c|c \leq \delta b - p] - p \right) \right\}$$

- You (weakly) prefer the fixed-fee contract if $EU_{FF} \geq EU_{PPV}$
- We can partition the possible values of c into three ranges:
  - When $c > \delta b$, you don't go to the gym
  - When $\delta b - p < c \leq \delta b$, you go to the gym if you've chosen the fixed fee contract, but not if you've chosen the pay-per-visit contract
  - When $c \leq \delta b - p$, you go to the gym under either contract
- Comparing payoffs across the two contracts for these three states:

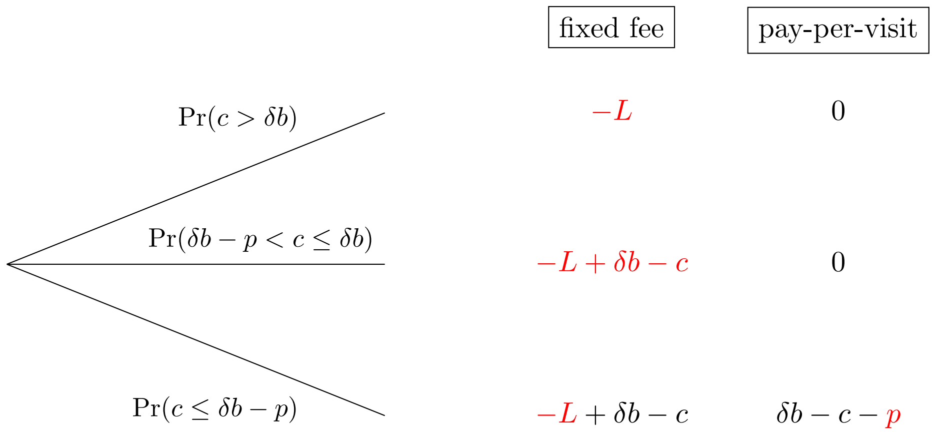

- These are realized payoffs conditional on c, but when the decision-maker evaluates a gym contract, she cares about expected values (so we need to think about the conditional expectation of c)

- The figure above illustrates that both EU_{FF} and EU_{PPV} implicitly contain a term

$$\Pr(c \leq \delta b - p) \left\{ \delta b - E\left[c | c \leq \delta b - p \right] \right\}$$

  which cancels from both when we compare the contracts

- Thus, the individual prefers the fixed fee contract when $EU_{FF} \geq EU_{PPV}$

$$\Leftrightarrow \delta \left\{ -L + \Pr(\delta b - p < c \leq \delta b) \left(\delta b - E[c | \delta b - p < c \leq \delta b] \right) \right\} \geq \delta \left\{ -\Pr(c \leq \delta b - p) p \right\}$$

$$\Leftrightarrow \Pr(c \leq \delta b - p) p + \Pr(\delta b - p < c \leq \delta b) \left(\delta b - E[c | \delta b - p < c \leq \delta b] \right) \geq L$$

- Now, notice that $\delta b - p < c \leq \delta b \Rightarrow \delta b - c < p$, so

$$\Rightarrow \Pr(\delta b - p < c \leq \delta b) \left(\delta b - E[c | \delta b - p < c \leq \delta b] \right) < \Pr(\delta b - p < c \leq \delta b) p$$

- This means that a dynamically-consistent person only prefers the fixed-fee contract when

$$L \leq \Pr(c \leq \delta b) p$$

since (again) $\Pr(c \leq \delta b) = \Pr(c \leq \delta b - p) + \Pr(\delta b - p < c \leq \delta b)$

- $\Pr(c \leq \delta b)$ is the expected number of trips to the gym under the fixed-fee contract, so the inequality above tells us that the fixed fee divided by the expected number of visits should be less than the pay-per-visit price p
- With more periods, this pattern should be even more pronounced since future pay-per-visit payments are discounted substantially when an individual is choosing her contract at $t = 0$ (making the pay-per-visit contract more attractive)
- So what explains the empirical regularity documented by Professor Stefano DellaVigna and Professor Ulrike Malmendier?
- Consider a present-biased decision-maker who discounts all future periods by an additional discount factor, $\beta < 1$, and holds beliefs about their degree of present bias, $\tilde{\beta} \in [\beta, 1]$
- Comparing payoffs across the two contracts for these three states:

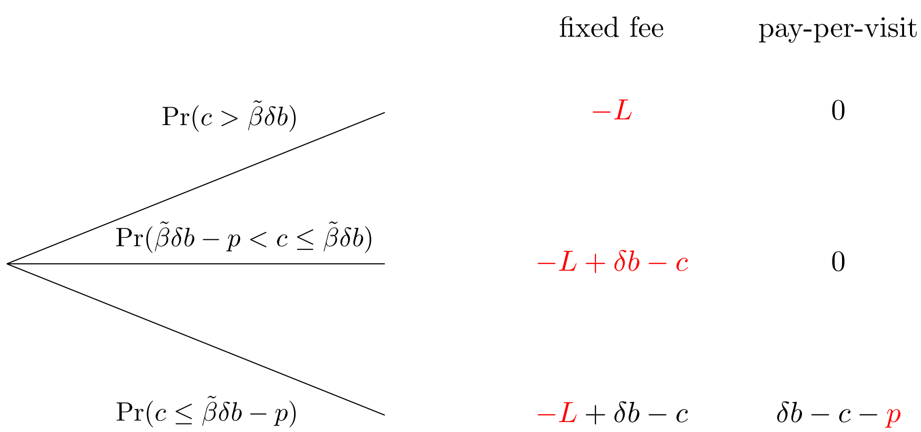

Revised: November 9, 2022
© 2022, Pamela Jakiela