1 The Economic Approach

1.1 What Is Economics?

1.1.1 Canonical Examples of Tradeoffs

1. The allocation of time:
   (a) Work effort vs. leisure time (consumption-leisure tradeoffs)
   (b) Work today vs.work tomorrow (intertemporal tradeoffs)
2. Allocations across people:
   (a) Me vs. you (or me vs. society)
   (b) Kid 1 vs. Kid 2 (intrahousehold allocation)
   (c) Richer vs. poorer households (redistribution)
3. Consumption decisions
   (a) Allocating spending across budget categories (housing vs. food)
   (b) Consumption of specific goods

1.2 Mathematical Models of Individual Choices

1.2.1 Example 1: Cookie Monster

- Suppose Cookie Monster's happiness from cookies is given by:

$$h(c) = 4c - c^2 \tag{1}$$

  where c is cookies and h is happiness measured in some reasonable units

- Why might we suppose this? Just to make up a nonsensical econ problem?

cookies	**happiness**
0	
1	
2	
3	
4	
5	

- *Graph the relationship between cookies and happiness*
- Functional form captures something about Cookie Monster's wellbeing: cookies make him happy, but too many cookies eventually make him worse off than zero cookies
- How many cookies should Cookie Monster eat?
  - What does "should" mean anyway?
- How many cookies will Cookie Monster eat if he wants to be as happy as possible?
  - A: 2
- We could have solved this using calculus to find the value of c that maximizes $h(c)$:
  - To find the critical value(s) of $h(c)$ we first take the derivative:

$$h(c) = 4c - c^2 \Rightarrow \frac{dh}{dc} = 4 - 2c \tag{2}$$

– Setting the derivative equal to zero gives us a **first-order condition** characterizing a (global or local) maximum or minimum:

$$\begin{aligned}\frac{dh}{dc} &= 4 - 2c^* = 0 \\ &\Leftrightarrow 4 = 2c^* \\ &\Leftrightarrow c = 2^*\end{aligned} \tag{3}$$

Which confirms what we saw in the table when we looked at whole numbers of cookies

– A lest step (that we will often skip in practice) is to check the **second-order condition** by looking at the sign of the second derivate:

  * If the function is concave at it's critical value, then the slope is decreasing in c, which means the critical value is a (local or global) maximum
  * If the function is convex, then the slope is increasing and the critical value represents a minimum

– In this case:

$$\frac{d^2h}{dc^2} = -2 < 0 \tag{4}$$

so we have identified a local maximum

– In practice, we will often be working with functions which are known to be concave, so we will not (always) need to check the second-order condition(s)

1.2.2 Example 2: A Lemming Eating a Popsicle

- Calculus is often helpful when we wish to model individual decisions about how much to consume

- A simpler consumption choice is whether or not to consume (or buy) a single unit of an **indivisible good**, a good which it is not possible to consume different amounts of

  – Examples: buying a car, getting vaccinated, having a child, going to college, etc.

  – In many of these cases, it might actually be possible to consume more than one unit, particularly over the long term, but we are often interested in modeling a yes-or-no decision about whether to consume one unit at a particular point in time

- Example (familiar to those who have seen Zootopia): a lemming is deciding whether or not to purchase a popsicle
- There are many lemmings, but we consider a specific individual lemming whom we'll refer to as "lemming i"
- Some lemmings like popsicles more than others: lemming i's valuation of a popsicle (i.e. the happiness they get from eating it) is given by v_i
- The price of a popsicle is p
- Lemming i will buy a popsicle whenever $v_i \geq p$, i.e. whenever the benefits of eating a popsicle (v_i) exceed the costs of eating a popsicle (in this case the price)
- This simple model is surprisingly useful, in part because it can be extended in a number of different ways depending on what aspect of the decision problem or the economic environment we want to study:
  - We can add random "noise" to the benefit an individual gets from eating a popsicle if we want to look at how often a lemming chooses a popsicle (perhaps as a function of the temperature on a given day, or the price of the popsicle, or how many other lemmings nearby are eating popsicles)
  - We can also model heterogeneity across individuals if we want to think about demand for popsicles in a given market (of lemmings)

1.2.3 Example 3: Many Lemmings, Some of Whom Are Eating Popsicles

- Suppose we have a large population of lemmings, and each individual lemming i is characterized by their valuation of popsicles, v_i, which represents the satisfaction lemming i gets from eating a popsicle
- v_i is a **random variable**: if we chose one lemming at random, we would not know how much that particular lemming liked popsicles

  Begin digression about random variables:

- Informally, a **random variable** is a variable that can take on different numeric values, where the realized outcome depends on chance or randomness or some other unobserved, plausibly random-ish process
- Some examples:
    1. The outcome of a roll of a die
    2. The number of times a coin lands on heads after ten coin flips
    3. Tomorrow's weather (1 = sunny, 0 = cloudy)
- A key thing about a random variable is that no one know in advance which of the possible outcomes will be realized, but we do know what outcomes are possible; then at some point the event takes place and we know the outcome (or at least it is possible for someone to know the outcome)
    1. If we roll a six-sided die, the possible realizations of the random variable are 1, 2, 3, 4, 5, and 6
    2. If we flip a "fair" coin twice, the possible outcomes are HH, HT, TH, and TT; and the possible realizations of the random variable "how many heads?" are 0, 1, and 2
    3. With tomorrow's weather, we said: 1 = sunny, 0 = cloudy
- Before the outcome of a random variable is realized, we can talk about the probabilities of different outcomes
    1. If we roll a six-sided die, what is the probability of each of the possible outcomes
    2. With two flips, HH, HT, TH, and TT are all equally likely (so what is the probability of each outcome?); so we can calculate the probability of each of the possible realizations of the random variable "how many heads?"
    3. With tomorrow's weather, we we can let $p \in [0, 1]$ be the probability that tomorrow is sunny
- The probabilities of all possible outcomes of a random variable must add up to 1

- If we observed a number of realizations of the random variable, we could calculate the average across however many trials we ran
- The **expected value** of a random variable is (again, informally) the average over infinitely many trials; it is the mathematical expectation of the variable
  - For the examples we've talked about – all of which have a finite number of possible outcomes – we can calculate the expected value by multiplying each possible outcome by its probability, and then summing over possible outcomes (do this for the examples)
  - These are examples of **discrete random variables** which can take on a finite number of possible values
  - The **probability distribution** of a discrete random variable is a list of its possible realizations and their associated probabilities
- We are often interested in **continuous random variables**, random variables that can take on an infinite number of possible values
  1. Example: a lemming's valuation of a possible falls between a and b, and can take any value in that range (so there are infinitely many possible valuations)
- Examples of normal and uniform random variables
- For continuous random variables, we can think about the probability of a specific outcome x (pdf) or the probability that the outcome is below some specific x (cdf)
- The area under the pdf is always 1: $\int_{-\infty}^{\infty} f(x)dx = 1$
- Discuss pdf and cdf of a uniform random variable

  End digression about random variables
- Back to lemmings: suppose there are a whole bunch of lemmings who differ in terms of how much they like popsicles, v_i
  - Given a price, p, the lemmings with a higher v_i (that is above p) will choose to buy a popsicle, while those with a relatively low v_i will not choose to buy a popsicle

- We don't know v_i for an individual lemming drawn at random from the population, but suppose we know something about the distribution of v_i
- Specifically, we'll often assume that we know the minimum and maximum possible values of v_i, which we will call a and b, and that all values between a and b are equally likely:
  - $a =$ the minimum possible value of v_i
  - $b =$ the maximum possible value of v_i
  - If all values of v_i between a and b are equally likely, we can say that "the random variable V is uniformly distributed on the interval $[a, b]$", or in mathematical notation, "$V \sim U[a, b]$"
  - *Draw the uniform pdf on the board if it is not already there.*
  - This is just saying that any individual lemming i's valuation of a popsicle, v_i is in this interval, and all the possible values between a and b are equally likely
  - Given this, which lemmings buy a popsicle at price p?
    * What proportion of the population buys a popsicle at $p < a$? Why?
    * What proportion of the population buys a popsicle at $p > b$? Why?
    * What if $a < p < b$? *Draw on the pdf.*
  - An individual lemming buys a popsicle $\Leftrightarrow v_i \geq p$
  - What is $\Pr[v_i > p]$ for an individual lemming drawn at random from the population, or equivalently what proportion of the lemmings will buy a popsicle at price $a < p < b$?
    * *Reason through it using the graph*
    * We can also do this with statistics by using two facts
      1. $\Pr[v_i > p] = 1 - \Pr[v_i \leq p]$
      2. The cumulative distribution function for the uniform distribution tells us the probability that an individual draw of v_i is less than or equal to p: if $X \sim U[a, b]$, the cdf of X is

$$F(x) = \frac{x-a}{b-a} \tag{5}$$

for any $x \in [a, b]$

* Given this, we have:

$$\begin{aligned}
\Pr[v_i > p] &= 1 - \Pr[v_i \leq p] \\
&= 1 - F(p) \\
&= 1 - \frac{p-a}{b-a} \\
&= \frac{b-a}{b-a} - \frac{p-a}{b-a} \\
&= \frac{b-p}{b-a}
\end{aligned} \tag{6}$$

- So, we started with a simple model of (lemming) behavior, and we were able to aggregate up to market behavior by imposing some (fairly reasonable) structure on the distribution of individual tastes;
- We'll see that we can use models like this to think about how equilibrium would change in response to different types of events: the introduction of a tax, more popsicle-sellers entering the market, etc.
- The model is simple – too simple – but there may be settings where it captures individual heterogeneity in a tractable and useful way; and we can also think about how the predictions of the model might change if we used a different distribution of v_i

1.2.4 Concluding Remarks

- We've seen two and a half models of individual choice behavior
- We choose the mathematical model to suit the aspect of human (or in this case monster or lemming) behavior that we want to study
- Economists like to say "All models are wrong, but some are useful."
  - Is it wrong about the aspect of behavior we want to study, and in a way that matters (in terms of the model's predictions)?
  - Is it useful in generating non-obvious predictions, or ruling out some behaviors or equilibrium outcomes that seem intuitively plausible?