Intertemporal Choice and Uncertainty

Simon Naitram[1]

Financial Economics (ECON6043)

April 8, 2021

[1] University of the West Indies, simon.naitram@cavehill.uwi.edu

Tonight's Lecture

- Gamestop
- Discussion on Economics and Finance
- Intertemporal Choice
- Uncertainty

Questions for Summers: Economics and Finance

1. Why would it occur that two fields attempting to answer the same questions simply do not interact?
2. Why does Summers talk about ketchup?
3. Do you side more with the standard finance approach or the standard economics approach?
4. Do you agree with Summers that there's actually a problem here?
5. Explain the debate over excess volatility in the stock market.
    - What are stock prices?
    - What are future dividend streams?
    - What is a required return?
    - What is the difference between ex post and ex ante?

Intertemporal Choice

- One of the main reasons that the financial market is so tricky is the time-based nature of many decisions
- The fundamental tool we use to understand time in economics is the concept of intertemporal choice.
- We begin by examining a simple two-period model.

Intertemporal Budget Constraint

- An individual receives an income in period 1 (m_1) and an income in period 2 (m_2).
- They can choose to consume all of their income in the period they receive it.
- Or they can transfer income from one period to another by **saving**, which gives an interest return of r.
- The intertemporal budget constraint is:

$$(m_1 - c_1)(1 + r) = c_2 - m_2$$
$$(m_1 - c_1)(1 + r) + m_2 = c_2$$

- Note that $m_1 - c_1$ can be negative too.

Intertemporal Preferences

- We can apply the standard theory of consumer choice to time decisions.
- The objects of choice are now the streams of consumption over time (c_1, c_2).
- One popular utility function for intertemporal choice is a utility function that is additive over time:

$$U(c_1, c_2) = u_1(c_1) + u_2(c_2)$$

where $u_t(c_t)$ is the utility of consumption in period t.

Intertemporal Preferences

- We can convert this to a time-stationary utility function:

$$U(c_1, c_2) = u(c_1) + \alpha u(c_2)$$

- The term α is important: it is the **discount factor**
- It answers the question: "how much would you give up today in order to have one dollar of consumption tomorrow?"

The consumer's optimization problem is therefore:

$$\begin{aligned} \max_{c_1, c_2} \quad & U(c_1, c_2) = u(c_1) + \alpha u(c_2) \\ \text{s.t.} \quad & c_2 = (m_1 - c_1)(1 + r) + m_2 \end{aligned}$$

with the first order conditions:

$$\begin{aligned} \frac{u'(c_1)}{(1+r)} &= \lambda \\ u'(c_2) &= \lambda \end{aligned}$$

or more simply:

$$u'(c_1) = \alpha u'(c_2)(1 + r).$$

Uncertainty

- When we begin to talk about time, we need to acknowledge that the future is inherently uncertain.
- We do so by considering that there is a probability of an event happening.
- Probability implies risk.

Lotteries

Consider a simple lottery: you buy a ticket that pays \$100 with probability p or \$0 with probability $(1 - p)$.

- Getting a prize with probability $p = 1$ is the same as getting the prize for certain.
- It doesn't make a difference how we frame the probabilities.
- The consumer only cares about the *net probabilities*—the eventual probability of an event.

Given any two lotteries, the consumer can choose between them.

Expected Utility

- We can use utility functions to describe a consumer's preferences about consumption in different circumstances.
- How an individual values an outcome depends also on the probability that the outcome will occur.
- One very convenient form is a utility function that has the **expected utility property**:

$$u(x, y) = pu(x) + (1 - p)u(y)$$

- This says the utility of a lottery is the expectation of the utility from its prizes.

Expected Utility

- Expected utility is the main microfoundation of financial economics.
- It is a theory that tries to describe how people make financial decisions when they face uncertainty over the outcomes.
- Do we think it is a reasonable representation of people's approach to decision-making?

Expected Utility

Some properties of expected utility:

- If a consumer is indifferent between two outcomes x and y, then:

$$pu(x) + (1-p)y = (1-p)x + py$$

- Additive separability means that when an outcome happens, we don't care about the other outcomes (independence of outcomes).

- If an outcome $x \succ y$ then a lottery $pu(x) + (1-p)y$ is preferred to $qu(x) + (1-q)y$ only if $p > q$.

- For many potential outcomes, we can write expected utility as:

$$\sum_i^n p_i u(x_i)$$

Risk Aversion

- We can describe the consumer's behaviour over financial gambles if we know the representation of their utility that has the expected utility property.
- In particular, we can compare:
  - The utility the consumer would get from getting the **expected value** of the lottery.
  - The **expected utility** the consumer would get from the lottery.

Risk Aversion

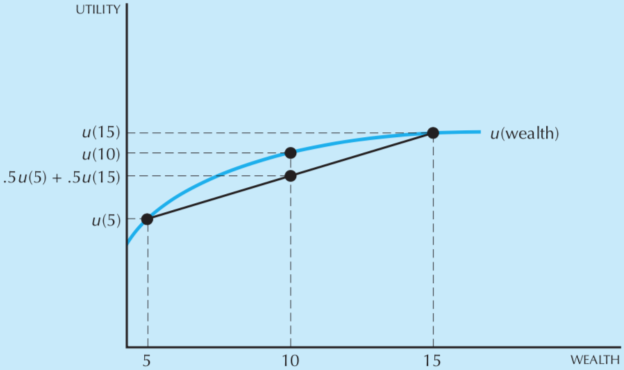

Risk aversion. For a risk-averse consumer the utility of the expected value of wealth, $u(10)$, is greater than the expected utility of wealth, $.5u(5) + .5u(15)$.

Risk Aversion

- In the preceding example, the consumer would prefer to have the expected value of their wealth rather than face the gamble.
- Such behaviour is called **risk aversion**.
- For a consumer to be risk averse, a *chord* drawn between any two points of the graph of the utility function must lie below the function.
- That is—if the consumer's expected utility function over wealth is **concave** then the consumer is risk averse.
- Conversely, if the consumer's expected utility function over wealth is **convex** then the consumer is risk loving.

Risk Aversion

- The more concave the function is, the more risk averse a person is.
- It is convenient for us to have a measure of risk aversion, tied to the concavity of the expected utility function.
- The **Arrow-Pratt measure of absolute risk aversion** is:

$$r(w) = -\frac{u''(w)}{u'(w)}$$

where w is wealth.

Application: Demand for Insurance

- A consumer with wealth W faces a probability p that they will lose an amount L.
- The consumer can purchase insurance that will reimburse them q dollars if the loss occurs at a cost πq.
- How much coverage (q) will the consumer optimally demand?

Application: Demand for Insurance

The consumer's problem is:

$$\max_q \quad p \cdot u(W - L - \pi q + q) + (1 - p) \cdot u(W - \pi q)$$

Taking the derivative wrt q and setting it equal to zero:

$$pu'(W - L - \pi q^* + q^*)(1 - \pi) - (1 - p)u'(W - \pi q^*)(\pi) = 0$$

or, written more succinctly:

$$\frac{u'(W - L - \pi q^* + q^*)}{u'(W - \pi q^*)} = \frac{(1 - p)}{p} \frac{\pi}{(1 - \pi)}$$

Application: Demand for Insurance

- The insurance company's expected profit is:

$$p(\pi q - q) + (1 - p)\pi q$$

- If we assume there's perfect competition in the insurance market and the firm's expected profits are equal to zero then:

$$\pi = p.$$

- The cost of the premium would be exactly its expected value—the *actuarially fair value*.

Application: Demand for Insurance

- Plugging this into the consumer's optimality conditions, we find that the RHS $=$ 1, and therefore:
$$u'(W - L - \pi q^* + q^*) = u'(W - \pi q^*)$$
- If a consumer is strictly risk averse so that $u''(w) < 0$, then the only way for this to happen is if the monetary outcomes are exactly the same in both cases:
$$W - L - \pi q^* + q^* = W - \pi q^*$$
- This implies that at the optimal, $L = q$, so that the consumer completely insures themselves against loss. This holds regardless of the consumer's absolute level of risk aversion, as long as they are risk averse.

Risk Premium

- We can formally express the idea that one individual is more risk averse than another by saying that individual A would be willing to pay more to avoid a risk than another individual B would.
- Consider the case of a payoff that takes the value of a random variable ϵ with expected value zero, so that $E(\epsilon) = 0$.
- Define $\pi_A(\epsilon)$ as the maximum amount that individual A is willing to pay (wealth they're willing to give up) to avoid facing the random variable ϵ.
- The amount $\pi_A(\epsilon)$ that equalises the utility from having a sure payoff versus facing the random payoff is called the **risk premium**.

$$u_A(w - \pi_A(\epsilon)) = E(A(w + \epsilon))$$

Pratt's Theorem

Pratt's theorem tells us that we can compare risk aversion between two individuals in any of three equivalent ways. If:

- $-\frac{u''_A(w)}{u'_A(w)} > -\frac{u''_B(w)}{u'_B(w)}$, or
- $u_A(w) = G(u_B(w))$ where $G(.)$ is some strictly increasing concave function, or
- $\pi_A(\epsilon) > \pi_B(\epsilon)$ for all random variables ϵ with an expected value of zero,

then we can say that individual A is **more risk averse** than individual B.

Relative Risk Aversion

- The main problem with our absolute measure of risk aversion is that it can often change based on the level of wealth that an individual is starting with.
- For example, we might find that individuals might accept more gambles express in absolute dollar terms if they begin with more wealth.
- It is less convincing for us to argue that an individual is willing to risk losing a larger *fraction* of their wealth as they become richer.
- We therefore consider a measure of **relative risk aversion**, which is now dimensionless:

$$\rho = -\frac{u''(w)w}{u'(w)}$$

The Allais Paradox

Choose between the two gambles:

- **Gamble A**: a 100% chance of receiving $1 million
- **Gamble B**: a 10% chance of receiving $5 million, an 89% chance of receiving $1 million, and a 1% chance of receiving nothing.

The Allais Paradox

Choose between the two gambles:

- **Gamble C**: an 11% chance of receiving \$1 million, and an 89% chance of receiving nothing.
- **Gamble D**: a 10% chance of receiving \$5 million, and a 90% chance of receiving nothing.

The Allais Paradox

- Many people prefer A to B and prefer D to C.
- However,

$$u(1) > 0.1u(5) + 0.89u(1) + 0.01u(0)$$

  Rewriting:

$$0.11u(1) > 0.1u(5) + 0.01u(0)$$

  Then adding $0.89u(0)$ to both sides:

$$0.11u(1) + 0.89u(0) > 0.1u(5) + 0.90u(0)$$

- An expected utility maximiser *regardless of their risk aversion* should prefer gamble C to gamble D if the prefer A to B.

Example: Two-Period Model

- Let's return to the two-period model, but now in a portfolio setting.
- We still have c_1 and c_2 as consumption, with an initial endowment now called w_1 that can be invested or consumed in period 1.
- Let's now consider two assets: one that pays a fixed return r_0 as before, and one that pays a *random* return of $\tilde{r}_1$.
- One notation you'll often find useful is the **gross return**:

$$R_0 = (1 + r_0) \quad \text{and} \quad \tilde{R}_1 = 1 + \tilde{r}_1.$$

- Note that the expected return on any risky asset can be written as the risk-free rate plus the risk premium (Varian, Microeconomic Analysis Pg. 188).

Example: Two-Period Model

- The consumer chooses to invest a fraction x of wealth in the risky asset and a fraction $(1 - x)$ in the certain asset.
- Note that wealth is $w_1 - c_1$, so total return on investment is

$$\tilde{R}_1(w_1 - c_1)x + (w_1 - c_1)(1 - x)R_0.$$

- In the second period, the individual consumes all their wealth, so that there is no income in period 2:

$$\tilde{w} = \tilde{c}_2 = (w_1 - c_1)[\tilde{R}_1 x + (1 - x)R_0]$$

- Note that second-period wealth (and therefore consumption) is uncertain since the return on investment is uncertain.

Example: Two-Period Model

- Rewriting the total return on the portfolio as:

$$\tilde{R} = \tilde{R}_1 x + (1 - x) R_0$$

we can simplify second period wealth, and plug it into the consumer's utility function:

$$U(c_1, x) = u(c_1) + \alpha E u(w_1 - c_1)\tilde{R}$$

- Differentiating with respect to c_1 and x, we get the FOCs:

$$u'(c_1) = \alpha E u(c_2)\tilde{R}$$
$$E u'(c_2)(\tilde{R}_1 - R_0) = 0$$

Example: Two-Period Model

- The first condition says that the marginal utility of period 1 consumption needs to be equal to the discounted expected marginal utility of period 2 consumption. This is the standard intertemporal optimization condition.
- The second condition determines the optimal allocation. It says that the consumer should invest in the risky asset until the expected marginal utility of consumption in period 2 is zero—as long as $\tilde{R}_1 > R_0$.

Homework for next week

- Your homework is simple: work through the similar two-period model in Varian, Microeconomic Analysis (which is the graduate text) on Page 184-186.
- We will discuss in class next week how it is solved, what the comparative statics imply, and how we can use this model to create testable predictions that a researcher might wish to examine.

The End.