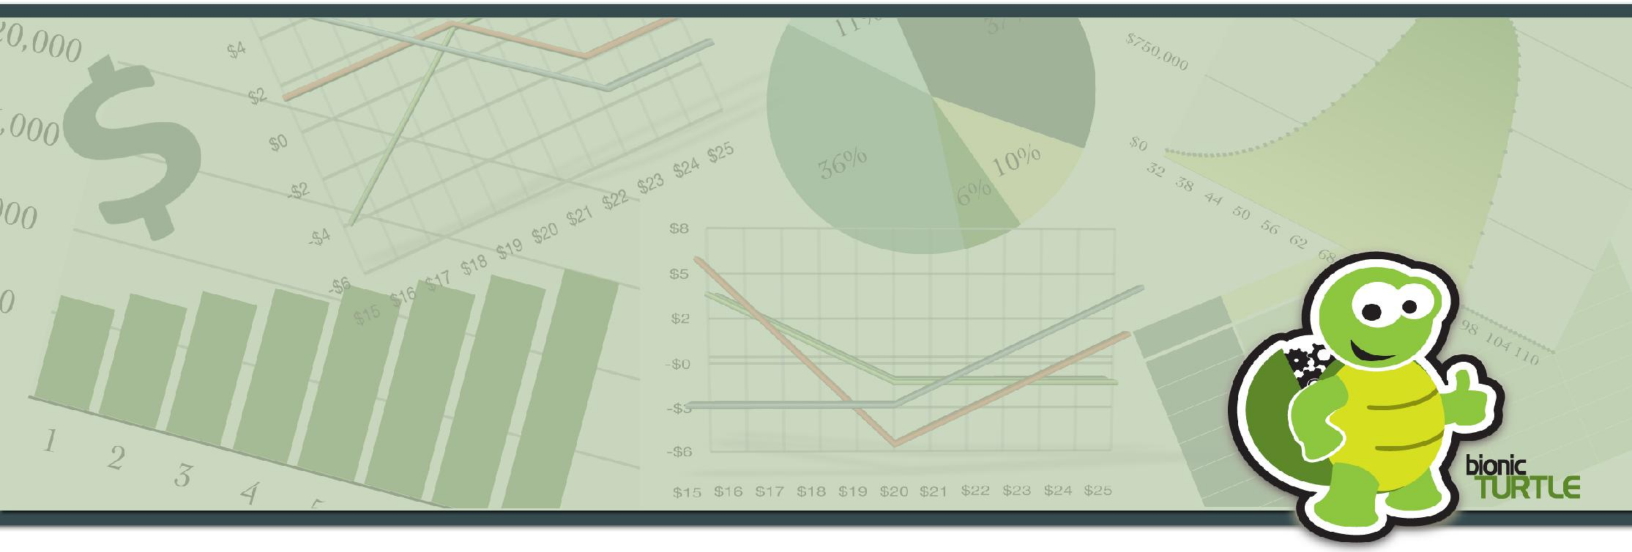

P1.T2. Quantitative Analysis

Chapter 1: Fundamentals of Probability

Bionic Turtle FRM Study Notes

Chapter 1: Fundamentals of Probability

Chapter 1: Fundamentals of Probability

- **Describe an event and an event space.**
- **Describe independent events and mutually exclusive events.**
- **Explain the difference between independent events and conditionally independent events.**
- **Calculate the probability of an event given a discrete probability function.**
- **Define and calculate a conditional probability**
- **Distinguish between conditional and unconditional probabilities.**
- **Explain and apply Bayes' rule.**

Key ideas and/or definitions in this chapter

- **Random (or statistical) experiment:** An observation or measurement process with multiple but uncertain outcomes; and:
  Random variable (or stochastic variable): A stochastic or random variable (r.v.) is a variable whose value is determined by the outcome of an experiment
- **Sample space:** Set of all possible outcomes of an experiment, and:
  Sample point: Each member or outcome of the sample space.
- **Outcome:** The result of a single trial. For example, if we roll two dice, an outcome might be a three (3) and a four (4); a different outcome might be a (5) and a (2); and:
  Event: The result that reflects none, one, or more outcomes in the sample space. Events can be simple or compound. An event is a subset of the sample space. If we roll two-dice, an example of an event might be rolling a seven (7) in total.
- **Discrete random variable:** A random variable (r.v.) that can take a *countable number* of finite number of values (or countably infinite). Examples include a coin (heads or tails); a six-sided die (1, 2, 3, 4, 5, or 6); or a bond default (default or survive); and:
  Continuous random variable: A random variable (r.v.) that can take any *measurable* value within some interval. Examples include asset returns; or time.
- **Mutually exclusive events:** Events which cannot simultaneously occur. If A and B are mutually exclusive, the probability of (A and B) is zero. Put another way, their intersection is the null set; and:
  Collectively exhaustive events (a.k.a., cumulatively exhaustive): Events that cumulatively describe all possible outcomes.
- **Unconditional (aka, marginal) probability, P[A]**: does not depend on prior event;
  Conditional probability, P[A|B]: depends on (given the occurrence of) another event;
  Joint probability, P[A∩B] or P[A,B] or P[AB]: probability of both occurrences.
- **Bayes' rule (aka, Bayes' theorem):** A rule that updates a prior probability, P[A], to a posterior probability, P[A | B], given evidence, P[B], and a likelihood, P[B | A]. Specifically, Bayes tells us the P[A|B] = P[B∩A]/P[B] = P[B|A] × P[A] / P[B] by relying on the fundamental relationship that a joint probability equals an unconditional probability multiplied by a conditional probability: P[B∩A] = P[B|A] × P[A] = P[A|B] × P[B].

Describe an event and an event space.

The sample space, denoted Ω, is the set of *all possible outcomes*. Its contents depend on the situation and might include nominal, ordinal (aka, categorical), interval or ratio data types. If the sample set includes real numbers, it will be denoted by $\mathbb{R}$, as would be probably the case for asset returns.

Events, denoted by ω, are subsets of the sample space. An *event* is a set of outcomes (but an event can contain zero elements). An *elementary event* includes only one outcome. The event space, denoted by $\mathcal{F}$, includes events (aka, outcome combinations) to which we can assign probabilities. For example, imagine two bonds where each bond eithers default or survive. We can observe four events.

{A defaults but B survives}, {B defaults, but A survives},
{Both A and B default}, or {Both A and B survive; aka, Neither A nor B defaults}.

This event space contains four events, which is a finite number and can be thusly be called a *discrete probability space*. Alternatively, we can define the even {only one bond defaults} such that this alternative event space contains only three events: {none default, only one defaults, or both default} but the total probability will still sum to 100.0% as the probability of {only one default} would equal the sum of {A defaults, but B survives} and {B defaults, but A survives}

Describe independent events and mutually exclusive events.

Mutually exclusive events (One random variable and two mutually exclusive events)

If two events are mutually exclusive, the probability of either occurring is the sum of their individual probabilities. Statistically:

$$P[A \cup B] = P[A] + P[B] \text{ if mutally exclusive}$$

where A ∪ B is the union of A and B; i.e., the probability of *either A or B occurring*.[1]

This equality is true only for mutually exclusive events. This property of mutually exclusive events can be extended to any number of events. The probability that any of n mutually exclusive events occurs is the sum of the probabilities of (each of) those n events.

Independent events (More than one random variable)

The random variables X and Y are independent if the conditional distribution of Y given X equals the marginal distribution of Y. Since **independence** implies P (Y=y | X=x) = P(Y=y):

$$\mathrm{P}(Y = y \mid X = x) = \frac{\mathrm{P(X = x, Y = y)}}{\mathrm{P\ (X = x)}}$$

Statistical independence is when the value taken by one variable has no effect on the value taken by the other variable. If the variables are independent, **their joint probability will equal the product of their marginal probabilities**. If they are not independent, they are dependent.

[1] Michael Miller, Mathematics and Statistics for Financial Risk Management, 2nd Edition (Hoboken, NJ: John Wiley & Sons, 2013)

The *most useful test* of statistical independence is given by:

$$P(X = x, Y = y) = P(X = x) \times P(Y = y)$$

That is, random variables X and Y are *independent if their joint distribution is equal to the product of their marginal distributions*. For example, when rolling two dice, the outcome of the second one will be independent of the first. This independence implies that the probability of rolling double-sixes is equal to the product of P(rolling one six) multiplied by P(rolling one six). So, if two die are independent, then:

P (first roll = 6, second roll = 6) = P (rolling a six) × P (rolling a six) = (1/6) × (1/6) = 1/36

Explain the difference between independent events and conditionally independent events.

Independent Events

If the probability of an event does not depend on another event, the events are independent. If two events are independent, then their joint probability is equal to the product of their unconditional probabilities:

$P(A \cap B) = P(A)\,P(B)$

If events are independent, then conditional probability is equal to the unconditional probability:

$$P(B|A) = \frac{P\,(A \cap B)}{P\,(A)} = \frac{P(A)\,P\,(B)}{P(A)} = P(B)$$

By the same logic, if the events are independent, then also true is that $P(A|B) = P(A)$.

Mutually exclusive events cannot both occur such that $P(A \cap B) = 0$. This implies that mutually exclusive events cannot be independent; i.e., mutually exclusive events are dependent.

Conditional Independence

Independence can be conditional on another event. A and B are conditionally independent if:

$P(A \cap B \mid C) = P(A \mid C) \times P(B \mid C)$

$P(A \cap B \mid C)$ denotes the conditional probability that A and B jointly occur conditional on outcome C. Events can be conditionally independent yet unconditionally dependent. Events can be conditionally dependent, yet independent!

Calculate the probability of an event for a discrete probability function.

Probability: Classical or "a priori" definition

The probability of outcome (A) is given by:

$$P(A) = \frac{Number\ of\ outcomes\ favorable\ to\ A}{Total\ n\ umberof\ outcomes}$$

For example, consider a craps roll of two six-sided dice. What is the probability of rolling a seven? i.e., P[X=7]. There are six outcomes that generate a roll of seven: 1+6, 2+5, 3+4, 4+3, 5+2, and 6+1. Further, there are 36 total outcomes. Therefore, the probability is 6/36 =1/6.

In this case, the outcomes need to be mutually exclusive, equally likely, and "cumulatively exhaustive" (i.e., all possible outcomes included in total). A key property of a probability is that the sum of the probabilities for all (discrete) outcomes is 1.0.

Probability: Relative frequency or empirical definition

Relative frequency is based on an **actual number of historical observations** (or Monte Carlo simulations). For example, here is a simulation (produced in Excel) of one hundred (100) rolls of a single six-sided die:

Roll	Expected	Actual Freq	Empirical Distribution
1	16.67%	11	11.0%
2	16.67%	17	17.0%
3	16.67%	18	18.0%
4	16.67%	21	21.0%
5	16.67%	18	18.0%
6	16.67%	15	15.0%
Total	**100.00%**	**100**	**100.0%**

Note the difference between an a priori probability and an empirical probability:

- The a **priori (classical) probability** of rolling a three (3) is 1/6.
- But the **empirical frequency**, based on this sample, is 18%. If we generate another sample, we will produce a different empirical frequency.

This relates also to *sampling variation*. The a priori probability is based on population properties; in this case, the a priori probability of rolling any number is clearly 1/6th. However, a sample of 100 trials will exhibit sampling variation: the number of threes (3s) rolled above varies from the parametric probability of 1/6th. We do not expect the sample to produce 1/6th perfectly for each outcome.