Foundations

Part 1 of Credit Products Playground

danieltuzes

Content with changes after version 0.4.0, not yet released
https://danieltuzes.github.io/credit-and-derivatives/

Contents

1 Foundations **1**

1.1 Cash-flow timelines and perspective **4**
1.1.1 What you will be able to do . . . 4
1.1.2 Start with events, not formulas . . . 4
1.1.3 Check your understanding . . . 5
1.1.4 Model boundary and review note . . . 6

1.2 Events, conditional probability, and expectation **7**
1.2.1 What you will be able to do . . . 7
1.2.2 Outcomes, events, and probabilities . . . 7
1.2.3 Conditional probability restricts the outcomes . . . 8
1.2.4 Expectation from disjoint events . . . 8
1.2.5 Check your understanding . . . 10
1.2.6 Model boundary and review note . . . 10

1.3 Rate quotes, compounding, and basis points **12**
1.3.1 What you will be able to do . . . 12
1.3.2 A rate quote needs its compounding convention . . . 12
1.3.3 Common unit boundary . . . 13
1.3.4 Check your understanding . . . 13
1.3.5 Model boundary and review note . . . 14

1.4 Discount factors **16**
1.4.1 What you will be able to do . . . 16
1.4.2 Undo accumulation . . . 16
1.4.3 Meaning before arithmetic . . . 16
1.4.4 Check your understanding . . . 17
1.4.5 Model boundary and review note . . . 18

1.5 Present value of a cash-flow schedule **19**
1.5.1 What you will be able to do . . . 19
1.5.2 Make cash flows comparable before adding . . . 19
1.5.3 Check your understanding . . . 20
1.5.4 Model boundary and review note . . . 21

1.6 No-arbitrage and replication **23**
1.6.1 What you will be able to do . . . 23
1.6.2 Compare cash flows, not labels . . . 23
1.6.3 Value the replicating portfolio . . . 23
1.6.4 Where replication is used next . . . 24
1.6.5 Model boundary and review note . . . 24

1.7 Risk-neutral pricing is not a risk-free probability **26**
1.7.1 What you will be able to do . . . 26
1.7.2 Two measures for two different uses . . . 26
1.7.3 From a partition to a price . . . 27
1.7.4 Worked example . . . 27
1.7.5 Uncertainty remains . . . 28
1.7.6 Advanced boundary: uniqueness . . . 28
1.7.7 Check your understanding . . . 28
1.7.8 Model boundary and review note . . . 29

A Notation **30**

B Abbreviations **34**

List of Figures

1.5.1 One-payment discounting explorer 21

List of Tables

1.3.1 One rate in three representations 13

1.7.1 Scenarios of the worked example 27

Part 1

Foundations

This part builds the valuation tools that the rest of the course reuses. The first three chapters set up the inputs: a timeline of signed, dated cash flows, events and expectations over a finite partition, and rate quotes with their compounding convention and units. The next two chapters turn a rate into a discount factor and add discounted cash flows into a present value. The last two chapters value one payoff from another. A replicating portfolio fixes a price by no-arbitrage, and risk-neutral pricing states the same price as a discounted expectation under pricing weights, which are not forecasts.

The argument

1. A valuation starts from signed cash flows placed on one time axis
2. An expectation over disjoint events weights each value by its probability
3. A rate quote needs its compounding convention to give a periodic rate
4. A discount factor converts one future unit into its value at time zero
5. Present value discounts each dated amount before adding
6. A portfolio that matches every future cash flow has the same price
7. Pricing weights state that price as a discounted expectation

After this part you can

- Read the amount and timing of each cash flow from a stated valuation-time timeline.
- Assign positive and negative signs to cash flows from an explicitly stated perspective.
- Identify events and a disjoint exhaustive partition in a stated probability model.
- Calculate the probability of one event conditional on another positive-probability event.
- Calculate a finite expected value from event probabilities and either event values or conditional event means.
- Interpret a nominal annual rate together with its stated compounding frequency.
- Calculate the per-period rate $j^{(m)}/m$ from a nominal annual rate $j^{(m)}$ compounded m times per year.
- Convert rate changes among decimal, percentage-point, and basis-point units.
- Explain a discount factor as today's value of one unit paid at a future date.
- Calculate $D(0,T) = (1+j^{(m)}/m)^{(-mT)}$ from a nominal annual rate $j^{(m)}$ compounded m times per year.
- Explain present value as the valuation-time sum of discounted cash flows under stated assumptions.
- Discount and add dated deterministic cash flows to calculate present value.

- Explain why two portfolios with identical dated state-contingent cash flows must have the same value under the lesson's no-arbitrage assumptions.
- Calculate a claim's no-arbitrage value from the current values and signed quantities of a portfolio that exactly replicates its future cash flows.
- Interpret risk-neutral probabilities as no-arbitrage pricing weights rather than real-world forecasts.
- Calculate time-zero value from finite-state risk-neutral probabilities, signed terminal payoffs, and a deterministic discount factor.

Leads to

- Bonds
- Rates and curves
- Derivative foundations
- Credit risk
- CDS

Chapter 1.1

Cash-flow timelines and perspective

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/foundations/cash-flow-timelines/

1.1.1 What you will be able to do

After this lesson, you should be able to:

- read an ordered schedule of payment amounts and times;
- state the time unit and the time origin of a timeline;
- apply a consistent sign convention from a named perspective.

1.1.2 Start with events, not formulas

A financial calculation begins with a list of events. The valuation-time, written 0, is the origin of every time in this lesson. The local payment-index k identifies one row. The index is not a time or a date. At row k, write the signed-cash-flow as CF_k and its payment-time as t_k. Row k pairs one amount with one time:

$$CF_k \quad \text{occurs at} \quad t_k \tag{1.1.1}$$

The shared subscript k pairs each amount with its time. The amount, the time, and the row index still have different meanings.

Example 1.1.1: Reading cash-flow timelines

Read a timeline

Suppose, in a simplified primary-issuance example, a bondholder pays an issuer USD 980 now, receives USD 30 after one year, and receives USD 1,030 after two years.

Payment index k	Payment time t_k, years from time 0	Signed cash flow CF_k, USD	Event
0	0	-980	Purchase price paid
1	1	+30	Coupon received
2	2	+1,030	Coupon and principal received

This table does **not** say whether USD 980 is a fair price. It only describes who exchanges what, and when.

Reverse the perspective

From that issuer's perspective, every amount above changes sign:

Time	Bondholder	Issuer
0 years	USD -980	USD +980
1 year	USD +30	USD -30
2 years	USD +1,030	USD -1,030

Changing the perspective does not change the dates or the absolute values of the amounts. Cash flows from opposite perspectives must be converted to one perspective before they are added. Otherwise the sum mixes two different models.

Same amount, different time

USD +100 at $t_1 = 0.5$ years and USD +100 at $t_2 = 2$ years are two distinct events. The amounts match; the times do not. A later lesson introduces the discount factors needed to compare them at time zero.

1.1.3 Check your understanding

These direct and transfer items cover reading a schedule and reversing perspective. Each question stays collapsed until you open it; answers and explanations appear once you check.

Knowledge check 1.1.1: Cash-flow timelines

1. From a holder's perspective, a timeline is measured in years from valuation time 0 and shows USD +25 at t = 1 and USD +100 at t = 2. Which reading is correct?

   (a) The USD 25 receipt occurs after one year and the USD 100 receipt after two years

   (b) The USD 100 receipt occurs after one year and the USD 25 receipt after two years

   (c) Both cash flows occur at valuation time

2. Time is measured in years from valuation time 0. Where does a payment scheduled 18 months later belong?

   (a) t = 0.5

   (b) t = 1

(c) t = 1.5

3. Use the convention that inflows are positive and outflows are negative. From an investor's perspective, how is principal received at maturity signed?

   (a) Positive
   (b) Negative
   (c) It has no sign

4. Use the convention that inflows are positive and outflows are negative. From an issuer's perspective, how is a coupon paid to investors signed?

   (a) Positive
   (b) Negative
   (c) It has no sign

Answers: *end of the chapter.*

1.1.4 Model boundary and review note

The times in this lesson are inputs, measured in years. Real schedules require calendar dates, business-day adjustments, and year-fraction conventions; none is inferred in this lesson.

The dated-schedule and timeline treatment here follows Tuckman & Serrat[1]. Reversing the holder's perspective changes the sign of every cash flow. This is a bookkeeping convention, not a claim from one source page. The lesson stays `draft` pending human confirmation of the printed locators.

References

[1] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). §1.1, the dated coupon schedule in Table 1.1; §1.6 and Figure 1.3, a timeline of dated events and day-count intervals.

Answers to the knowledge checks

Knowledge check 1.1.1: Cash-flow timelines

1. **(a)** The USD 25 receipt occurs after one year and the USD 100 receipt after two years. Each amount is paired with its stated time measured from valuation time 0.

2. **(c)** t = 1.5. Under the stated units, 18 months is 1.5 years after valuation time.

3. **(a)** Positive. Principal received is an inflow to the investor, so it is positive under the stated convention.

4. **(b)** Negative. A coupon payment is an outflow from the issuer, so it is negative under the stated convention.

Chapter 1.2

Events, conditional probability, and expectation

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/foundations/probability-events-and

1.2.1 What you will be able to do

After this lesson, you should be able to:

- identify events that form a disjoint and exhaustive partition;
- calculate a conditional probability;
- calculate an expectation when a random value is constant on each event;
- use conditional event means when the value varies inside an event.

1.2.2 Outcomes, events, and probabilities

An **outcome** is one possible result of a modeled experiment. An **event** is a specified set of outcomes. Use the local probability-event A for one such set and the local probability-measure $\mathbb{P}$ for the rule that assigns probabilities to events. In this lesson, $\mathbb{P}$ is a general probability measure. A later finance lesson uses $\mathbb{P}$ only for real-world probabilities and introduces a different symbol for pricing probabilities. This lesson does not choose between the two.

In this lesson, a finite partition is a finite list of events that are:

- **disjoint:** no outcome belongs to two different partition events; and
- **exhaustive:** every possible outcome belongs to one of them.

Use the local probability-event-index j to select an event and the local probability-partition-size m to count the events. For a disjoint and exhaustive list $A_1, \dots, A_m$, additivity of the probability measure gives the sum of the event probabilities:

$$\sum_{j=1}^{m} \mathbb{P}(A_j) = 1 \tag{1.2.1}$$

If the events overlap, the sum of their probabilities can count some outcomes twice. If the events are not exhaustive, the sum omits some outcomes. These event and additivity definitions follow Shreve's probability-space treatment[1].

1.2.3 Conditional probability restricts the outcomes

Suppose the question asks for the probability of event A **given** another event. Write the conditioning event as B. When $\mathbb{P}(B) > 0$, the conditional probability of A given B is defined as:

$$\mathbb{P}(A \mid B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} \tag{1.2.2}$$

The denominator is the probability of the outcomes that the condition "given B" keeps. For example, suppose 30% of all outcomes are in B, and 12% are in both A and B. Then the conditional probability is:

$$\mathbb{P}(A \mid B) = \frac{0.12}{0.30} = 0.40 \tag{1.2.3}$$

In a default model, consider the probability of "default during this period given survival to its start." Its denominator is the probability of survival to the start of the period. This conditional probability is different from the unconditional probability of default by a horizon. Shreve develops conditioning as averaging within the information or event retained by the condition[2].

1.2.4 Expectation from disjoint events

Use the local probability-random-value X for a numerical value determined by the outcome and the expectation operator $\mathbb{E}$ for its probability-weighted average. In this lesson every random value is constant on each event of a finite partition, so the general integral definition of $\mathbb{E}$ reduces to the weighted sum below. The expectation has the same unit as X; a random USD payment has an expected value in USD, not in percent.

This section gives two formulas. The first formula requires X to be constant on each event. The second formula does not.

1.2.4.1 Exact value on every event

First suppose X is constant on every partition event. Use the local probability-event-value x_j for its value on A_j. Then the expectation is the sum of the event values, each multiplied by the probability of its event:

$$\mathbb{E}[X] = \sum_{j=1}^{m} x_j \mathbb{P}(A_j), \qquad X = x_j \text{ on } A_j \tag{1.2.4}$$

Shreve states the expectation of a finite-valued random variable in the form of (1.2.4)[3].

1.2.4.2 Value varies inside an event

If X is not constant inside A_j, one representative value is in general not exact. For each event with positive probability, replace the value with the conditional mean of X inside that event. The expectation is then:

$$\mathbb{E}[X] = \sum_{j=1}^{m} \mathbb{P}(A_j)\mathbb{E}[X \mid A_j] \tag{1.2.5}$$

An event with zero probability contributes zero to the sum and can be omitted. Its conditional mean need not be defined.

The partition identity (1.2.5) follows by splitting an expectation across disjoint events and using the partial-averaging property of conditional expectation. Shreve gives the disjoint-event integral split and the unconditional-average property of conditional expectation[4].

Example 1.2.1: Expectation over disjoint events

Constant value on each event

Suppose X is a payment in USD. It equals USD 0, USD 50, or USD 100 on three disjoint exhaustive events with probabilities 0.10, 0.20, and 0.70.

Event	Constant payment, USD	Probability	Contribution, USD
A_1	0	0.10	0.00
A_2	50	0.20	10.00
A_3	100	0.70	70.00

The expectation is the sum of the contributions:

$$\mathbb{E}[X] = 0 \times 0.10 + 50 \times 0.20 + 100 \times 0.70 = \text{USD } 80$$

The expectation of USD 80 is a probability-weighted average. An expectation need not be one of the possible payments, and in this example it is not.

Conditional means inside events

Suppose three disjoint exhaustive events have probabilities 0.01, 0.02, and 0.97, and the payment varies inside each event. The conditional means of the payment are USD 40, USD 30, and USD 0. The exact expectation from these conditional means is:

$$\mathbb{E}[X] = 0.01 \times 40 + 0.02 \times 30 + 0.97 \times 0 = \text{USD } 1.00$$

USD 40 is the average payment inside the first event. The example does not claim that every outcome in the first event pays USD 40.

Default inside a period

For a CDS period, one event can be "default occurs in this period." A protection payment discounted at the actual default time varies within that event. Accrued premium also varies within the event, because a default near the end of the period accrues more premium than a default near its start.

So one midpoint value for the whole period is an approximation. The exact contribution of the event is the interval default probability multiplied by the conditional mean of the discounted payment inside the period. If the period is divided into smaller disjoint default-time intervals, the sum has more terms. In a continuous-time model, the limit of this sum is an integral over the default time. The later CDS lesson gives the hazard-rate probabilities and the cash-flow formulas for that integral.

1.2.5 Check your understanding

These items separately test events, conditioning, and both expectation forms. Each question stays collapsed until you open it; answers and explanations appear once you check.

Knowledge check 1.2.1: Events, conditioning, and expectation

1. Which description is an event in a probability model?
   (a) A specified set of possible outcomes
   (b) Only the number assigned as a probability
   (c) Only the average value of a random quantity
2. A one-year credit model allows at most one default. Which list is a disjoint and exhaustive partition of outcomes?
   (a) Default by 6 months; default after 6 months through 1 year; survive through 1 year
   (b) Default by 1 year; default after 6 months through 1 year; survive through 1 year
   (c) Default by 6 months; survive through 1 year
3. Events A and B satisfy P(A and B) = 0.12 and P(B) = 0.30. What is P(A given B)? Enter a decimal probability.
4. In a frequency example, 80 otherwise comparable names are alive at the start of a period and 2 of those names default during it. What conditional default fraction does the sample imply for that period? Enter a decimal.
5. A random payoff is constant at 0, 50, and 100 units on three disjoint exhaustive events with probabilities 0.10, 0.20, and 0.70. What is its expected payoff in units?
6. Three disjoint exhaustive events have probabilities 0.01, 0.02, and 0.97. A payment varies inside the events, but its conditional mean in each event is 40, 30, and 0 USD respectively. What is the unconditional expected payment in USD?

Answers: end of the chapter.

1.2.6 Model boundary and review note

This lesson teaches finite partitions and the idea of the limit of a refinement. It does not construct probability spaces, prove measure-theoretic integration, estimate probabilities from data, calibrate risk-neutral probabilities, or evaluate the default-time integrals used by a CDS model. A later lesson distinguishes real-world and risk-neutral measures.

The lesson paraphrases the cited textbook; it remains `draft` pending human verification of the printed locators, formulas, examples, and notation bindings.

References

[1] Shreve, Stochastic Calculus for Finance II: Continuous-Time Models (2004). Ch. 1 §1.1, printed pp. 1-3, Def. 1.1.2 and eqs. 1.1.2-1.1.5.

[2] Shreve, Stochastic Calculus for Finance II: Continuous-Time Models (2004). Ch. 2 §2.3, printed pp. 66-69, eqs. 2.3.4-2.3.17 and Def. 2.3.1.

[3] Shreve, Stochastic Calculus for Finance II: Continuous-Time Models (2004). Ch. 1 §1.3, printed pp. 13-18, Thms. 1.3.1 and 1.3.4.

[4] Shreve, Stochastic Calculus for Finance II: Continuous-Time Models (2004). Ch. 1 §1.3, printed p. 17, Remark 1.3.2; Ch. 2 §2.3, printed pp. 68-72, eqs. 2.3.16-2.3.17 and 2.3.25.

Answers to the knowledge checks

Knowledge check 1.2.1: Events, conditioning, and expectation

1. **(a)** A specified set of possible outcomes. An event is a set of outcomes. A probability measure assigns a number to that set.

2. **(a)** Default by 6 months; default after 6 months through 1 year; survive through 1 year. The three correct events cannot occur together and cover every outcome in the stated one-default model. The second list overlaps, and the third omits defaults in the second half-year.

3. 0.4 (accepted within $\pm 1e-12$). Divide the intersection probability by the conditioning-event probability: 0.12 / 0.30 = 0.40.

4. 0.025 (accepted within $\pm 1e-12$). Condition on the 80 names alive at period start: 2 / 80 = 0.025, or 2.5%. This sample fraction is not by itself a calibrated pricing probability.

5. 80 (accepted within $\pm 1e-9$). Weight each event value by its probability: $0 \times 0.10 + 50 \times 0.20 + 100 \times 0.70 = 80$.

6. 1 (accepted within $\pm 1e-9$). Use event probability times conditional event mean: $0.01 \times 40 + 0.02 \times 30 + 0.97 \times 0 =$ USD 1.00.

Chapter 1.3

Rate quotes, compounding, and basis points

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/foundations/rates-compounding-and-

1.3.1 What you will be able to do

After this lesson, you should be able to:

- interpret a nominal annual quote together with its frequency;
- calculate the rate applied in one compounding period;
- convert between decimal rates, percentages, and basis points.

1.3.2 A rate quote needs its compounding convention

The phrase "an annual rate of 6%" does not determine the compounding unless its convention is stated. In this lesson, write the nominal-annual-rate as $j^{(m)}$ and pair it with the compounding-frequency m. Write the resulting periodic-rate as r_m.

In this lesson, t denotes a time-from-valuation: a payment-time without a row index. The valuation-time is 0. Use the local number-of-compounding-periods N for the number of times the periodic rate is applied, and write the resulting accumulation-factor as $A(0,t)$.

Under the nominal-rate convention of this lesson, the periodic rate is the nominal annual rate divided by the compounding frequency:

$$r_m = \frac{j^{(m)}}{m} \tag{1.3.1}$$

This division is a convention of the quote. It is not a conversion rule for every kind of annual rate.

Example 1.3.1: Compounding and basis points

Nominal 6%, compounded semiannually

With $j^{(2)} = 0.06$ and $m = 2$, the periodic rate is:

$$r_2 = \frac{0.06}{2} = 0.03$$

The model applies the periodic rate of 3% once every half-year. Over $t = 1.5$ years, the local number-of-compounding-periods is $N = mt = 3$. So the accumulation-factor is:

$$A(0, 1.5) = (1 + 0.03)^3 = 1.092727$$

One unit at valuation-time accumulates to 1.092727 units at the payment-time 1.5 years later.

Nominal is not effective annual

Under the same convention, the effective annual rate, i.e., the growth over one year of compounding, is:

$$(1.03)^2 - 1 = 0.0609 = 6.09\%$$

So the nominal annual rate of 6% and the effective annual rate of 6.09% are different quantities. The name "annual rate" without a convention does not distinguish them.

Basis-point arithmetic

A basis-point (bp) is a unit for rate differences. One basis point is 0.0001, i.e., 0.01%:

$$25\,\text{bp} = 25 \times 0.0001 = 0.0025 = 0.25\%$$

A change from 4.10% to 4.35% is 0.25 percentage points (25 bp). Write the unit explicitly, because the number 25 without a unit is ambiguous.

1.3.3 Common unit boundary

Table 1.3.1 **One rate in three representations.** The same rate as the decimal the domain code takes, the percentage shown to a learner, and the number of basis points above zero.

Representation	Same rate
Decimal used by domain code	0.0435
Percentage shown to a learner	4.35%
Basis points above zero	435 bp

A lab converts a percentage input to a decimal. The domain functions take decimal rates.

1.3.4 Check your understanding

These direct and transfer items cover quote interpretation, periodic-rate calculation, and basis-point conversion. Each question stays collapsed until you open it; answers and explanations appear once you check.

Knowledge check 1.3.1: Rates, compounding, and basis points

1. Under the lesson convention, what does a 6% nominal annual rate compounded semiannually imply?

(a) A 3% rate in each six-month period

(b) A 6% rate in each six-month period

(c) A 6% rate in each month

2. Two rates are both quoted as 6% nominal annually, but one compounds annually and the other monthly. Which statement is correct under the lesson convention?

   (a) Their per-period rates differ, and their one-year accumulation generally differs

   (b) They have identical per-period rates because both are quoted as 6%

   (c) Both quotes must be 6% effective annual rates

3. Enter the per-period rate as a decimal for a 6% nominal annual rate compounded twice per year.

4. Enter the quarterly rate as a decimal for a 4.8% nominal annual rate compounded four times per year.

5. Express 75 basis points as a decimal rate.

6. A rate rises from 3.25% to 3.70%. How many basis points is the increase? Enter a number of basis points.

Answers: *end of the chapter.*

1.3.5 Model boundary and review note

This lesson does not introduce continuous compounding, a general conversion to effective rates, day counts, or a term structure.

The nominal-quote-plus-frequency convention and the basis-point unit follow Tuckman & Serrat[1]. The lesson stays `draft` pending human confirmation of the printed locators.

References

[1] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). §2.1, rates quoted as annual rates over a term of fixed-length periods, simple or compound; §2.4, eq. 2.17, the periodic-compounding form; a basis point is 0.01%, Ch. 0 footnote and Ch. 3.

Answers to the knowledge checks

Knowledge check 1.3.1: Rates, compounding, and basis points

1. **(a)** A 3% rate in each six-month period. For a nominal annual rate $j^{(m)}$ compounded m times per year, the lesson convention uses $r_m = j^{(m)}/m$ in each period: $6\%/2 = 3\%$.

2. **(a)** Their per-period rates differ, and their one-year accumulation generally differs. A nominal quote is incomplete without frequency; the period rate is $r_m = j^{(m)}/m$ and compounding applies it m times.

3. 0.03 (accepted within $\pm 1e-12$). With $j^{(2)} = 0.06$ and $m = 2$, the period rate is $r_2 = 0.06/2 = 0.03$.

4. 0.012 (accepted within $\pm 1e-12$). With $j^{(4)} = 0.048$ and $m = 4$, the quarterly rate is $r_4 = 0.048/4 = 0.012$.

5. 0.0075 (accepted within $\pm 1e-12$). One basis point is 0.0001 in decimal-rate units, so $75 \times 0.0001 = 0.0075$.

6. 45 (accepted within $\pm 1e-9$). The increase is 0.45 percentage point, and each percentage point contains 100 basis points: $0.45 \times 100 = 45$.

Chapter 1.4

Discount factors

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/foundations/discount-factors/

1.4.1 What you will be able to do

After this lesson, you should be able to:

- explain the units and meaning of a discount factor;
- calculate one from a stated nominal rate, frequency, and time;
- distinguish discounting from compounding and from payment probability.

1.4.2 Undo accumulation

In this lesson, t denotes a time-from-valuation: a payment-time without a row index. It is measured in years from valuation-time 0. Write the compounding-frequency as m, the periodic-rate as r_m, the accumulation-factor as $A(0,t)$, and the discount-factor as $D(0,t)$.

The accumulation factor $A(0,t)$ is the amount at time t from one unit invested at time zero. The discount factor is its reciprocal:

$$D(0,t) = \frac{1}{A(0,t)} = \left(1 + r_m\right)^{-mt} \tag{1.4.1}$$

The discount factor, the periodic rate, and the compounding frequency keep their different meanings, also when they appear in one expression.

1.4.3 Meaning before arithmetic

Write the local discounted-unit-value as $V_0(1_t)$: the value at valuation time of one deterministic unit paid at t. If $D(0,2) = 0.94$, the value at valuation time of one unit paid at $t = 2$ years is:

$$V_0(1_2) = 1 \times D(0,2) = 0.94 \tag{1.4.2}$$

The number 0.94 is not an annual interest rate of 94%. It is also not a probability of 94% that the payment occurs.

Example 1.4.1: Discount factors from rates and curves

Annual compounding

With a nominal-annual-rate of 5%, $m = 1$, and $t = 2$, the discount factor is:

$$D(0,2) = \left(1 + \frac{0.05}{1}\right)^{-1\times 2} = \frac{1}{1.05^2} \approx 0.907029$$

Semiannual compounding

With a nominal annual rate of 6%, compounded twice per year, the discount factor for 1.5 years is:

$$D(0,1.5) = \left(1 + \frac{0.06}{2}\right)^{-2(1.5)} = \frac{1}{1.03^3} \approx 0.915142$$

The horizon of 1.5 years contains three compounding periods. The exponent 1.5 would be wrong, because it counts years, not half-year periods.

Read a small curve

With the same nominal annual rate of 6%, compounded semiannually, the discount factors are:

Time t, years	Periods mt	Discount factor $D(0,t)$
0.0	0	1.000000
0.5	1	0.970874
1.0	2	0.942596
1.5	3	0.915142

If the rate is non-negative, the discount factor does not increase with the payment time. If the rate is positive, the discount factor decreases with the payment time.

1.4.4 Check your understanding

These items test meaning, annual calculation, and a semiannual transfer case. Each question stays collapsed until you open it; answers and explanations appear once you check.

Knowledge check 1.4.1: Discount factors

1. If D(0,2) = 0.94, what does 0.94 mean?
   (a) The value today of one unit paid in two years
   (b) A 94% annual interest rate
   (c) The probability that payment occurs
2. For the same positive future cash flow, which discount factor gives the lower present value?
   (a) 0.91
   (b) 0.97
   (c) They give the same value

3. Calculate the dimensionless two-year discount factor from a 5% nominal annual rate compounded annually.

4. Calculate the dimensionless 1.5-year discount factor from a 6% nominal annual rate compounded semiannually.

Answers: *end of the chapter.*

1.4.5 Model boundary and review note

The rate is the same for all times, and the cash flows are deterministic. This lesson does not teach spot curves, interpolation, day counts, default, recovery, liquidity, or funding.

The discount-factor definition (1.4.1) follows Tuckman & Serrat[1]; Hull gives the same interpretation, the present value of one unit paid in the future[2]. The lesson stays `draft` pending human confirmation of the printed locators.

References

[1] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). §1.2, the discount factor as the value today of one unit paid at t; §2.4, eqs. 2.17-2.19, the semiannually compounded form.

[2] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 4, "Bond Pricing".

Answers to the knowledge checks

Knowledge check 1.4.1: Discount factors

1. **(a)** The value today of one unit paid in two years. A discount factor converts a future unit into today's value; it is neither an interest rate nor a survival probability.

2. **(a)** 0.91. Present value is the cash flow multiplied by the discount factor, so the smaller positive factor produces the smaller value.

3. 0.9070294784580498 (accepted within ± 0.000001). $D(0,2) = \frac{1}{1.05^2}$.

4. 0.9151416593531596 (accepted within ± 0.000001). The periodic rate is $0.06/2 = 0.03$ and 1.5 years contains three half-year periods, so $D(0,1.5) = \frac{1}{1.03^3}$.

Chapter 1.5

Present value of a cash-flow schedule

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/foundations/present-value/

1.5.1 What you will be able to do

After this lesson, you should be able to:

- interpret present value from the same perspective as its cash flows;
- discount each dated deterministic amount;
- add discounted amounts without mixing payment dates.

1.5.2 Make cash flows comparable before adding

Use the local present-value-payment-index k to select one row and the local number-of-future-payments n to count the included future rows. At row k, write the signed-cash-flow as CF_k, its payment-time as t_k, and the matching discount-factor as $D(0, t_k)$. Write the schedule's present-value at valuation-time 0 as PV_0.

The present value of the deterministic schedule is the sum of the discounted cash flows:

$$PV_0 = \sum_{k=1}^{n} CF_k D(0, t_k) \tag{1.5.1}$$

The index k and the count n are defined only in this lesson. The other symbols are shared across the course.

Two schedules are compatible when they use the same valuation time, perspective, currency, and discount factors. Let $\mathrm{A} \parallel \mathrm{B}$ be the schedule that contains every row of both schedules A and B. For two compatible schedules, the present value is additive:

$$PV_0(\mathrm{A} \parallel \mathrm{B}) = PV_0(\mathrm{A}) + PV_0(\mathrm{B}) \tag{1.5.2}$$

Example 1.5.1: Present values of dated cash flows

One payment

A deterministic receipt of USD 1,000 is paid after two years, and $D(0, 2) = 0.9070294785$. The

present value of the receipt is:

$$PV_0 = 1000 \times 0.9070294785 = \text{USD } 907.0294785 \approx \text{USD } 907.03$$

The present value is in USD at valuation-time, from the holder's perspective.

Mixed signed cash flows

Suppose the holder receives USD 100 at t_1, pays USD 40 at t_2, and receives USD 300 at t_3. The discount factors are 0.97, 0.93, and 0.88.

Payment index k	Cash flow CF_k, USD	Discount factor $D(0, t_k)$	Discounted cash flow $CF_k D(0, t_k)$, USD
1	+100	0.97	+97.00
2	-40	0.93	-37.20
3	+300	0.88	+264.00

The present value is the sum of the discounted cash flows:

$$PV_0 = 97.00 - 37.20 + 264.00 = \text{USD } 323.80$$

Discounting does not change the sign of a cash flow. Reversing the perspective would change the sign of the present value.

Additivity

Combining two schedules keeps every row, including duplicate payments. For compatible schedules, a property test of the domain code checks (1.5.2). The compatibility conditions are necessary: present values in different currencies, or at different valuation dates, cannot be added without another model.

Figure 1.5.1 shows the one-payment case. The discount-factor for one time-from-valuation t is calculated from one discounting-lab-rate with annual compounding. One discounting-lab-cash-flow multiplied by that discount factor gives the present-value at valuation-time 0.

The tested domain code calculates every value in the figure. The figure states the rate as a percentage. The calculation converts it to a decimal first.

1.5.3 Check your understanding

These items separate interpretation from calculation and include an unfamiliar multi-payment case. Each question stays collapsed until you open it; answers and explanations appear once you check.

Knowledge check 1.5.1: Present value

1. Under stated discount factors and a valuation time of 0, what does the present value of a future cash-flow stream represent?
   (a) The valuation-time sum of the individually discounted cash flows
   (b) The undiscounted total of every future cash flow
   (c) Only the cash flow at the latest payment time
2. Every cash flow in a stream is positive. A second stated discount-factor set is smaller at every payment time. What happens to present value?

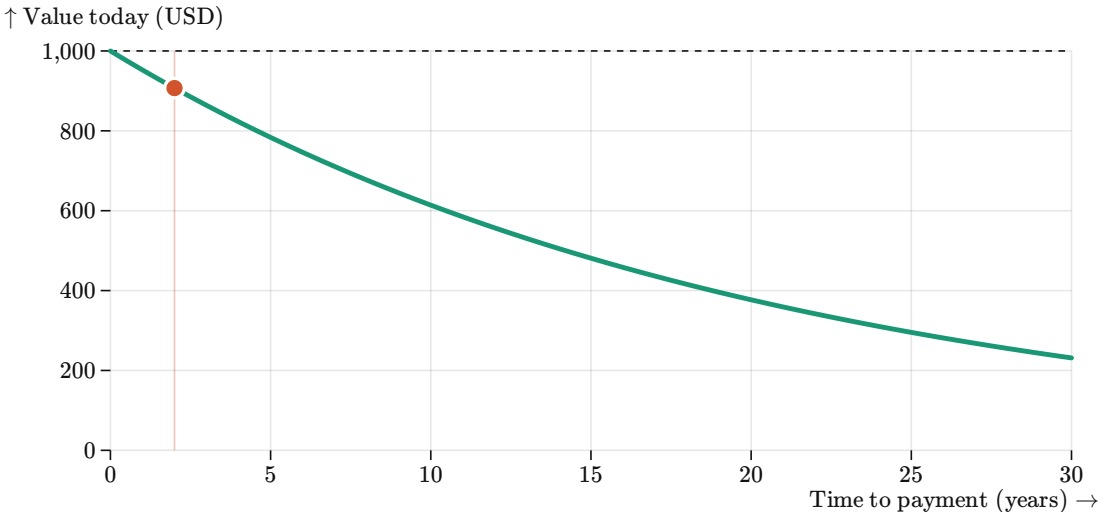


Figure 1.5.1 **One-payment discounting explorer.** The value today, in USD, of one future payment for every time to payment from 0 to 30 years, at one annual rate with annual compounding. The dot and the vertical line mark the time to payment shown; the dashed line marks the undiscounted payment. Shown for a future payment of USD 1,000 in 2 years at an annual rate of 5%: value today USD 907.03.

(a) Present value is lower

(b) Present value is higher

(c) Present value is unchanged

3. From a holder's perspective, a deterministic USD 1,000 receipt at t = 2 years has discount factor D(0,2) = 0.9070294785. What is its time-zero present value in USD?

4. From a holder's perspective, two deterministic USD 100 receipts at different future dates have matching discount factors 0.97 and 0.93. What is their total time-zero present value in USD?

Answers: *end of the chapter.*

1.5.4 Model boundary and review note

This lesson covers deterministic present value only. It omits curves, calendars, uncertainty, credit, recovery, liquidity, tax, and funding.

The rule (1.5.1) follows Tuckman & Serrat[1]; Hull states the same principle[2]. The lesson stays `draft` pending human confirmation of the printed locators.

References

[1] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). §1.2, price as the sum of each cash flow times its discount factor (eqs. 1.1-1.3); full price equals present value (eq. 1.5).

[2] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 4, "Bond Pricing".

Answers to the knowledge checks

Knowledge check 1.5.1: Present value

1. **(a)** The valuation-time sum of the individually discounted cash flows. Present value puts each dated cash flow at the common valuation time using its discount factor, then adds the contributions.

2. **(a)** Present value is lower. Each positive cash flow is multiplied by a smaller non-negative factor, so every contribution and their sum are lower.

3. 907.0294785 (accepted within ± 0.01). Multiply the future cash flow by its discount factor.

4. 190 (accepted within ± 0.01). Discount each dated payment and then add: $100 \times 0.97 + 100 \times 0.93 = 190$.

Chapter 1.6

No-arbitrage and replication

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/foundations/no-arbitrage-and-repli

1.6.1 What you will be able to do

After this lesson, you should be able to:

- explain why identical dated state-contingent cash flows have the same price under the stated assumptions;
- distinguish exact replication from matching only an expected cash flow;
- calculate a claim value from signed positions in a replicating portfolio.

1.6.2 Compare cash flows, not labels

The law of one price applies when two portfolios produce the same signed-cash-flow at every relevant future date and in every modeled state. The names of the portfolios can differ. What matters is the full set of dated cash flows, from the same perspective.

Suppose the two portfolios had different prices at valuation-time 0. In the frictionless model of this lesson, a trader could then buy the cheaper portfolio and sell the more expensive one. The future cash flows of the two positions cancel, so the trader keeps the initial difference in price. This arbitrage argument is the reason for the pricing rule. Tuckman and Serrat state the law of one price and connect deviations from it to an arbitrage trade built from a replicating portfolio[1].

Matching only the expected cash flow is not enough. A payment of USD 100 only in an up state is not the same cash flow as USD 100 only in a down state, even if the two states have equal probability.

1.6.3 Value the replicating portfolio

Use the local replication-security-index i to select one security and the local replication-security-count n to count them. Let the signed replication-security-quantity be a_i and the current replication-security-value be $V_{0,i}$. When this portfolio exactly matches the claim's future cash flows, the replicated-claim-value is the value of the portfolio:

$$V_0^{\text{claim}} = \sum_{i=1}^{n} a_i V_{0,i} \tag{1.6.1}$$

A positive a_i is a long position, and a negative a_i is a short position. The formula uses current values because the portfolio already matches the claim's cash flows at every future date and in every state.

Suppose a claim is replicated by two units of security A worth USD 12 each and one unit of security B worth USD 7. The value of the replicating portfolio, and therefore of the claim, is:

$$V_0^{\text{claim}} = 2(12) + 1(7) = \text{USD } 31 \tag{1.6.2}$$

1.6.4 Where replication is used next

Forward prices use cash-and-carry replication. Put-call parity compares two portfolios with identical expiry cash flows. A binomial option model constructs a state-by-state replicating portfolio. Risk-neutral valuation is another way to calculate that same no-arbitrage value inside a specified model.

Knowledge check 1.6.1: No-arbitrage and replication

1. Portfolio A and claim B have identical signed cash flows at every future date in every modeled state. Under the lesson's frictionless no-arbitrage assumptions, what must be true?
   (a) They must have the same current value
   (b) They need the same contract name but may have unrelated values
   (c) Only their average future cash flow needs the same value
2. Two portfolios have the same expected maturity cash flow, but one pays only in an up state and the other only in a down state. Does the law of one price force their current values to be equal?
   (a) No, matching an expectation does not replicate state-by-state cash flows
   (b) Yes, the same expected cash flow always implies the same value
   (c) Yes, sharing a maturity date is sufficient
3. A claim is exactly replicated by two units of security A worth USD 12 each and one unit of security B worth USD 7. What is the claim's no-arbitrage value in USD?
4. A claim is exactly replicated by buying three units of security A worth USD 18 each and shorting two units of security B worth USD 11 each. What is the claim's no-arbitrage value in USD?

Answers: *end of the chapter.*

1.6.5 Model boundary and review note

Real trades can have bid-ask spreads, different borrowing and lending rates, short-sale constraints, collateral, taxes, liquidity differences, and execution risk. These effects can make a textbook

arbitrage impossible to execute. This lesson sets up a frictionless pricing benchmark. It does not claim that every observed price difference is a riskless profit.

The source locator, notation, examples, and answer keys remain `draft` pending independent human review.

References

[1] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 1 §§1.3-1.4, printed pp. 53-57, law of one price and arbitrage replication.

Answers to the knowledge checks

Knowledge check 1.6.1: No-arbitrage and replication

1. **(a)** They must have the same current value. If identical future cash flows had different current values, buying the cheaper package and selling the dearer one would create an initial receipt with offsetting future obligations under the stated assumptions.

2. **(a)** No, matching an expectation does not replicate state-by-state cash flows. Replication requires the same dated cash flow in every modeled state. Equal averages alone leave different state exposure.

3. 31 (accepted within ± 0.000001). The replicating portfolio costs $2 \times 12 + 1 \times 7 =$ USD 31, so the law of one price gives the claim the same value.

4. 32 (accepted within ± 0.000001). The signed portfolio value is 3×18 - $2 \times 11 =$ USD 32. The short position enters with a negative quantity.

Chapter 1.7

Risk-neutral pricing is not a risk-free probability

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/foundations/risk-neutral-pricing/

1.7.1 What you will be able to do

After this lesson, you should be able to:

- distinguish a real-world probability from a risk-neutral pricing weight;
- explain why changing measure does not remove uncertainty or volatility;
- calculate an expected discounted payoff in a finite-state model;
- state why a simple CDS model can omit a separate volatility input without claiming that volatility never matters.

1.7.2 Two measures for two different uses

Valuation is at valuation-time 0. Let the local risk-neutral-pricing-horizon T be a future time in exact model-years, and let the terminal-random-payoff X_T be the signed USD amount that a claim holder receives or pays at T.

A probability measure assigns mutually consistent weights to events. The real-world-probability-measure $\mathbb{P}$ describes the modeled probabilities of events in the real world. It is the measure for a forecast. The risk-neutral-probability-measure $\mathbb{Q}$ gives the weights of an arbitrage-free pricing model. So the same event can have different probabilities under $\mathbb{P}$ and $\mathbb{Q}$. Shreve makes this distinction between the actual and the risk-neutral measure in the introduction of the change of measure [1]. Hull makes the same distinction between forecasting and valuation [2].

> **There is no "risk-free measure" in this model.** The measure is **risk-neutral**. "Risk-free" describes the cash account and discounting convention used as the pricing benchmark. It does not describe the event probabilities, and it does not make the payoff safe.

The name "risk-neutral" also does not assume that every investor is indifferent to risk. Risk preferences and risk premia are among the reasons why $\mathbb{P}$ and $\mathbb{Q}$ differ. The pricing calculation uses $\mathbb{Q}$ after $\mathbb{Q}$ has been selected or calibrated.

1.7.3 From a partition to a price

In a no-arbitrage model with the cash account as numeraire, $\mathbb{Q}$ is chosen so that discounted traded prices are martingales. In words: after discounting with the cash account, today's traded price equals the conditional $\mathbb{Q}$-expectation of its discounted future value. Shreve states the martingale condition and its discounted-payoff consequence [3]. Tuckman and Serrat give the same expected-discounted-value construction in a finite-state fixed-income setting [4].

Write the expectation under the risk-neutral-probability-measure as $\mathbb{E}^{\mathbb{Q}}$. It is the same averaging operator, taken with the pricing weights of $\mathbb{Q}$. Use the local risk-neutral-scenario-index j and risk-neutral-scenario-count m. Let x_j be the local risk-neutral-scenario-payoff, in USD at T, and let q_j be its dimensionless risk-neutral-scenario-probability. For disjoint, exhaustive scenarios on which the payoff is constant, the risk-neutral expected payoff is:

$$\mathbb{E}^{\mathbb{Q}}[X_T] = \sum_{j=1}^{m} x_j q_j, \qquad q_j \geq 0, \qquad \sum_{j=1}^{m} q_j = 1 \tag{1.7.1}$$

This expected payoff is in USD at T. It is not a present value. Let the input discount-factor be $D(0,T)$, in USD at time zero per USD at T. The claim's present-value is the discounted risk-neutral expected payoff:

$$PV_0 = \mathbb{E}^{\mathbb{Q}}\left[D\left(0,T\right) X_T\right] = D(0,T)\mathbb{E}^{\mathbb{Q}}[X_T] = D(0,T)\sum_{j=1}^{m} x_j q_j \tag{1.7.2}$$

The second equality uses this lesson's assumption that $D(0,T)$ is deterministic. If the discount factor is random, it must in general stay inside the expectation. The expectation of a product equals the product of the expectations only under additional assumptions.

1.7.4 Worked example

Suppose $D(0,T) = 0.95$ and the claim has two disjoint, exhaustive outcomes:

Table 1.7.1 **Scenarios of the worked example.** Each scenario's risk-neutral probability, payoff at the horizon, and contribution to the risk-neutral expected payoff, in USD at the horizon.

Scenario j	Risk-neutral probability q_j	Payoff x_j, USD at T	Contribution $q_j x_j$, USD at T
1	0.25	120	30.00
2	0.75	40	30.00

The risk-neutral expected payoff (1.7.1) and time-zero value (1.7.2) are

$$\mathbb{E}^{\mathbb{Q}}[X_T] = 0.25(120) + 0.75(40) = \text{USD } 60 \text{ at } T \tag{1.7.3}$$

$$PV_0 = 0.95(60) = \text{USD } 57 \text{ at time } 0 \tag{1.7.4}$$

No $\mathbb{P}$ probabilities enter this price, because the example states the $\mathbb{Q}$ pricing weights as inputs. An estimate of the real-world expected payoff would be a different calculation.

1.7.5 Uncertainty remains

A change from $\mathbb{P}$ to $\mathbb{Q}$ changes the weights of the modeled outcomes. It does not remove the randomness. In the standard change-of-measure result for a diffusion, the drift changes and the diffusion coefficient stays the same [5]. So volatility remains a model input for nonlinear claims such as options. A restricted linear model can omit a separate volatility input for reasons specific to its product. The lesson on that product states those reasons.

1.7.6 Advanced boundary: uniqueness

The pricing measure need not be unique. Under the usual technical conditions, the absence of arbitrage is tied to the existence of a suitable martingale measure. Uniqueness of the measure requires a stronger condition: the modeled market must be complete. This section is a boundary note, and the assessment of this lesson does not test it. Shreve develops the existence and uniqueness distinction [6].

1.7.7 Check your understanding

These checks separate interpretation from arithmetic and include unfamiliar scenario weights and a signed negative payoff.

Knowledge check 1.7.1: Risk-neutral pricing

1. Which statement correctly distinguishes a real-world probability measure P from a risk-neutral probability measure Q in this lesson?
   (a) P describes modeled real-world frequencies; Q gives no-arbitrage pricing weights
   (b) P and Q must assign the same probability to every event
   (c) Q assigns certainty to every traded payoff
2. An analyst says, 'After changing from P to Q, uncertainty disappears and volatility is absorbed by the measure.' Which response is correct for the standard diffusion result taught here?
   (a) The change reweights outcomes and changes drift, but it does not remove randomness or the diffusion volatility
   (b) The change makes every future payoff deterministic
   (c) The change sets every volatility parameter to zero
3. At T, a signed payoff is USD 120 with Q-probability 0.70 and USD 20 with Q-probability 0.30. The deterministic discount factor is D(0,T) = 0.96. What is its time-zero risk-neutral value in USD?
4. A contract has signed terminal payoffs of USD 200, USD 50, and USD -40 in three disjoint exhaustive scenarios. Their Q-probabilities are 0.20, 0.50, and 0.30, and D(0,T) = 0.94. What is the time-zero risk-neutral value in USD?

Answers: end of the chapter.

1.7.8 Model boundary and review note

This lesson uses one terminal payoff, a finite scenario partition, deterministic risk-free discounting, and input $\mathbb{Q}$ weights. It does not estimate $\mathbb{P}$, calibrate $\mathbb{Q}$, generate a yield curve, treat a stochastic numeraire, establish market completeness, or value counterparty, liquidity, funding, or option effects.

All source records, notation, formulas, examples, and answer keys remain `draft` pending independent human source, quantitative, and editorial review.

References

[1] Shreve, Stochastic Calculus for Finance II: Continuous-Time Models (2004). Ch. 1 §1.6, printed pp. 32-35.

[2] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 12 §12.2, printed pp. 257-260; Ch. 14 §14.7, printed pp. 311-313.

[3] Shreve, Stochastic Calculus for Finance II: Continuous-Time Models (2004). Ch. 5 §5.2.2, printed pp. 214-217, eqs. 5.2.22-5.2.24; §5.2.4, printed pp. 218-219, eqs. 5.2.29-5.2.31.

[4] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 7 §7.3, printed pp. 182-184, eqs. 7.7-7.8.

[5] Shreve, Stochastic Calculus for Finance II: Continuous-Time Models (2004). Ch. 5 §5.2.2, printed pp. 214-217, especially eqs. 5.2.22-5.2.24.

[6] Shreve, Stochastic Calculus for Finance II: Continuous-Time Models (2004). Ch. 5 §§5.4.3-5.4.4, printed pp. 228-233.

Answers to the knowledge checks

Knowledge check 1.7.1: Risk-neutral pricing

1. **(a)** P describes modeled real-world frequencies; Q gives no-arbitrage pricing weights. P and Q can weight the same outcomes differently. P is the modeled real-world measure; Q is selected so discounted traded prices satisfy the model's no-arbitrage pricing condition.

2. **(a)** The change reweights outcomes and changes drift, but it does not remove randomness or the diffusion volatility. Risk-neutral valuation changes the probability weights used for pricing. In the standard change-of-measure result, the drift changes while the diffusion coefficient remains; uncertainty is still present.

3. 86.4 (accepted within ± 0.01). First compute the Q-expectation: $0.70 \times 120 + 0.30 \times 20 = 90$ future USD. Then discount: $0.96 \times 90 =$ USD 86.40 at time zero.

4. 49.82 (accepted within ± 0.01). The Q-expected terminal payoff is $0.20 \times 200 + 0.50 \times 50 + 0.30 \times (-40) =$ USD 53. Discounting gives $0.94 \times 53 =$ USD 49.82.

Appendix A

Notation

The symbols of this book, in the order they first appear.

k **Payment index**. Labels one row in the ordered cash-flow schedule. The index is an ordering label; the associated payment time is measured in years. *In Chapter 1.1 only.*

$$k \in \{0, 1, 2, ...\}$$

t_k **Payment time**. Time from the valuation date to one scheduled cash-flow event; a time coordinate, not a calendar date or payment index. *Units:* years from the valuation date.

CF_k **Signed cash flow**. Amount received or paid at one event from the stated holder perspective; positive means received and negative means paid by that holder. *Units:* stated currency units.

0 **Valuation time**. Common origin from which later model times and present values are measured. *Units:* years from the valuation date.

A **Probability event**. A specified set of possible outcomes in this lesson's probability model. An indexed form selects one event in a finite disjoint and exhaustive partition. *In Chapter 1.2 only.*

$$A_j$$

j **Probability event index**. Selects one event in the lesson's finite partition. The index is bookkeeping and has no probability or monetary unit. *In Chapter 1.2 only.*

$$j \in \{1, \dots, m\}$$

x_j **Probability event value**. The constant value taken by the random value on one partition event. This quantity can replace a conditional mean only when every outcome in the event gives the same value. *Units:* value units. *In Chapter 1.2 only.*

$\mathbb{P}$ **Probability measure**. Assigns each modeled event a number between zero and one. The symbol is page-local and generic; it does not yet select a real-world or risk-neutral interpretation. *Units:* probability between zero and one. *In Chapter 1.2 only.*

A_j **Probability partition event**. One event of the lesson's finite disjoint and exhaustive partition, selected by the event index. *In Chapter 1.2 only.*

m **Probability partition size**. Counts the events in the lesson's finite partition. The count is a positive integer and does not represent time. *In Chapter 1.2 only.*

$$m \in \mathbb{N}$$

X **Probability random value**. A numerical value determined by which modeled outcome occurs. Its unit is stated by the example, such as USD for a random payment. *Units:* value units. *In Chapter 1.2 only.*

$\mathbb{E}$ **Expectation**. Operator returning the average of a random quantity, each outcome weighted by its probability under a stated measure; a superscript names that measure when more than one is in play. *Units:* value units.

$$\mathbb{E}[X] = \int_{\Omega} X(\omega)\, d\mathbb{P}(\omega)$$

N **Number of compounding periods**. Counts how many times the periodic rate is applied over the stated horizon. Under equal periods, the count equals compounding frequency multiplied by time in years. *In Chapter 1.3 only.*

$$N = mt$$

$A(0,t)$ **Accumulation factor**. Grows one current unit over a stated future horizon under the selected rate model. *Units:* dimensionless currency-units per current currency-unit.

$1\,\mathrm{bp}$ **Basis point**. Rate-change unit equal to one hundredth of one percentage point. *Units:* decimal rate.

m **Compounding frequency**. Number of equal compounding periods per year under the stated rate convention, a positive integer fixed by the model convention. *Units:* compounding periods per year.

$j^{(m)}$ **Nominal annual rate**. Annualized rate quote that must be paired with its compounding frequency. *Units:* decimal rate per year.

r_m **Periodic rate**. Rate applied once in each compounding period under the stated convention, derived from the stated nominal annual quote in this model. *Units:* decimal per compounding period.

t **Time from valuation**. A time measured in years from the valuation time; a time coordinate, not a calendar date or a schedule index. With a row index it is a payment time. *Units:* years from the valuation date.

$V_0(1_t)$ **Discounted unit value**. A local reading aid that makes discount-factor units explicit. It equals one future currency unit multiplied by the discount factor for its payment time. *Units:* valuation currency-units per future currency-unit. *In Chapter 1.4 only.*

$$V_0(1_t) = 1_t D(0,t)$$

$D(0,t)$ **Discount factor**. Converts one deterministic future unit into value at valuation time. *Units:* current currency-units per future currency-unit.

C **Discounting lab cash flow**. The one dated amount the discounting lab discounts. It is page-local to the lab's one-payment model; the lesson's schedule notation writes an indexed amount instead. *Units:* stated currency at payment time. *In Chapter 1.5 only.*

$$PV_0 = C \cdot D(0,t)$$

y **Discounting lab rate**. The one annual rate the discounting lab compounds to build each discount factor. It is page-local to the lab's toy model; elsewhere this lesson takes every discount factor as given rather than deriving it from a rate. *Units:* decimal rate per year. *In Chapter 1.5 only.*

$$D(0,t) = (1+y)^{-t}$$

n **Number of future payments**. The final index in the finite present-value sum. It counts included future rows and does not itself specify maturity in years. *In Chapter 1.5 only.*

$$n \in \mathbb{N}$$

k **Present value payment index**. Selects one dated cash flow in the finite schedule. This lesson-local index keeps sequence bookkeeping separate from the associated time measured in years. *In Chapter 1.5 only.*

$$k \in \{1, \dots, n\}$$

PV_0 **Present value**. Combines dated signed cash flows into one value at valuation time, using the same holder perspective as the signed cash flows. *Units:* stated currency at the valuation time.

V_0^{claim} **Replicated claim value**. Current value assigned to the claim whose future cash flows are exactly reproduced. Under the stated no-arbitrage assumptions it equals the signed current value of the replicating portfolio. *Units:* stated currency at valuation time. *In Chapter 1.6 only.*

n **Replication security count**. Counts the traded securities in the finite replicating portfolio. It is a positive integer rather than a time or cash amount. *In Chapter 1.6 only.*

$$n \in \mathbb{N}$$

i **Replication security index**. Selects one traded security in the finite replicating portfolio. It is a bookkeeping label and has no monetary unit. *In Chapter 1.6 only.*

$$i \in \{1, \dots, n\}$$

a_i **Replication security quantity**. Number of units of security i held in the replicating portfolio. A positive quantity is long and a negative quantity is short from the portfolio holder's perspective. *Units:* security units. *In Chapter 1.6 only.*

$V_{0,i}$ **Replication security value**. Current currency value per unit of one security used in the replicating portfolio. Every value uses the same valuation time and currency. *Units:* stated currency per security at valuation time. *In Chapter 1.6 only.*

T **Risk-neutral pricing horizon**. Future model time at which the one terminal payoff is delivered. It is measured in exact model-years from valuation time zero and is not a calendar date. *Units:* years from valuation time. *In Chapter 1.7 only.*

$$T > 0$$

m **Risk-neutral scenario count**. Counts the disjoint and exhaustive scenarios in the finite pricing partition. It is a positive integer and does not measure time or money. *In Chapter 1.7 only.*

$$m \in \mathbb{N}$$

j **Risk-neutral scenario index**. Selects one scenario in the finite pricing partition. It is a bookkeeping label rather than a probability or time. *In Chapter 1.7 only.*

$$j \in \{1, \dots, m\}$$

x_j **Risk-neutral scenario payoff**. Realized signed terminal amount in one scenario. Positive means received and negative means paid by the claim holder at the horizon. *Units:* stated currency at the horizon. *In Chapter 1.7 only.*

q_j **Risk-neutral scenario probability**. Pricing weight assigned by the risk-neutral measure to one scenario. The weights are dimensionless, non-negative, and sum to one over the finite partition. *Units:* dimensionless probability. *In Chapter 1.7 only.*

$$q_j \geq 0, \quad \sum_{j=1}^{m} q_j = 1$$

$\mathbb{P}$ **Real-world probability measure**. Assigns modeled probabilities intended to describe actual-world event likelihoods; used for forecasting and statistical statements under the stated real-world model. *Units:* dimensionless probability weights between zero and one.

$\mathbb{Q}$ **Risk-neutral probability measure**. Gives model pricing weights under which discounted traded prices satisfy the martingale condition; it prices payoffs relative to a stated numeraire and is not a forecast of actual event frequencies. *Units:* dimensionless probability weights between zero and one.

X_T **Terminal random payoff**. Signed amount delivered by a claim at the stated future horizon, before its outcome is known; positive means received and negative means paid by the claim holder. *Units:* stated currency at the future horizon.

Appendix B

Abbreviations

CDS Credit default swap, the credit derivative the CDS lessons define and value.

CDX The name of a family of standard credit default swap indices, each a standard portfolio of single-name contracts.

DTS Duration times spread, a spread-risk measure for bonds.

FINRA Financial Industry Regulatory Authority.

IMM International Money Market. In the CDS lessons, IMM dates are the standard maturity dates on the twentieth of March, June, September, and December.

ISDA International Swaps and Derivatives Association.

ISO International Organization for Standardization.

JTD Jump to default, the loss on an immediate default of the reference entity.

UTC Coordinated Universal Time, the time standard the date arithmetic counts calendar days in.