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

2 Bonds **30**

2.1 Fixed-rate bond contract and cash flows **33**
2.1.1 What you will be able to do . . . 33
2.1.2 What a bond is . . . 33
2.1.3 Read the contract before valuing it . . . 33
2.1.4 Calculate the coupon . . . 34
2.1.5 Construct every promised payment . . . 34
2.1.6 Check your understanding . . . 35
2.1.7 Model boundary and review note . . . 36

2.2 Bond price from discount factors **37**
2.2.1 What you will be able to do . . . 37
2.2.2 Bond pricing is a present-value application . . . 37
2.2.3 Error pattern: discount after adding . . . 38
2.2.4 Check your understanding . . . 38
2.2.5 Model boundary and review note . . . 39

2.3 Yield to maturity as a single-rate summary **40**
2.3.1 What you will be able to do . . . 40
2.3.2 What yield means in this toy model . . . 40
2.3.3 Coupon rate and yield have different roles . . . 40
2.3.4 Check your understanding . . . 41
2.3.5 Model boundary and review note . . . 42

2.4 The bond price-yield relationship **43**
2.4.1 What you will be able to do . . . 43
2.4.2 The price decreases because every denominator increases . . . 43
2.4.3 Check your understanding . . . 44
2.4.4 Model boundary and review note . . . 45

2.5 Settlement, accrued interest, and clean versus dirty price 47
2.5.1 The price paid is not always the price quoted . . . 47
2.5.2 Actual/actual within one coupon period . . . 47
2.5.3 Why this matters for derivatives . . . 48
2.5.4 Sources . . . 48

2.6 Bond forwards 49
2.6.1 The forward contract . . . 49
2.6.2 Cash-and-carry price . . . 49
2.6.3 Basis consistency . . . 49
2.6.4 Sources . . . 50

3 Rates and curves 51

3.1 Discount curve and forward discounting 53
3.1.1 What you will be able to do . . . 53
3.1.2 A curve is a dated collection, not one yield . . . 53
3.1.3 Discount between future times . . . 53
3.1.4 Scope boundary . . . 54
3.1.5 Sources . . . 54

4 Derivative foundations 56

4.1 Forward contracts, delivery price, and value 59
4.1.1 What you will be able to do . . . 59
4.1.2 An obligation, not an option . . . 59
4.1.3 Fair delivery price at inception . . . 59
4.1.4 An existing forward can gain or lose value . . . 60
4.1.5 Model boundary and review note . . . 61

4.2 European option contracts and payoffs 62
4.2.1 What you will be able to do . . . 62
4.2.2 Rights and obligations . . . 62
4.2.3 Expiry payoff . . . 62
4.2.4 Value, premium, payoff, and profit are different . . . 63
4.2.5 Model boundary and review note . . . 64

4.3 Put-call parity and synthetic forwards 65
4.3.1 What you will be able to do . . . 65
4.3.2 Match the expiry cash flows . . . 65
4.3.3 Present-value parity . . . 65
4.3.4 When put-call parity does not apply . . . 66
4.3.5 Model boundary and review note . . . 66

4.4 One-period binomial option valuation **68**
4.4.1 What you will be able to do . . . 68
4.4.2 Two states, two traded positions . . . 68
4.4.3 The same value with risk-neutral weights . . . 69
4.4.4 Worked call example . . . 69
4.4.5 From one period to many periods . . . 69
4.4.6 Model boundary and review note . . . 70

4.5 Multi-period lattice valuation **71**
4.5.1 Repeat the one-period argument backward . . . 71
4.5.2 Timing is part of the model . . . 71
4.5.3 Worked example . . . 72
4.5.4 Sources . . . 72

5 Bond options **73**

5.1 European bond options **75**
5.1.1 The standalone contract . . . 75
5.1.2 Put the bond value and the strike on one price basis . . . 75
5.1.3 Worked example . . . 76
5.1.4 Coupon and exercise ordering . . . 76
5.1.5 Sources . . . 77

5.2 Issuer-default knockout bond options **78**
5.2.1 Contract definition . . . 78
5.2.2 Value the underlying bond first . . . 78
5.2.3 Apply the knockout to option continuation . . . 78
5.2.4 Worked example . . . 79
5.2.5 Put-call parity with a default knockout . . . 79
5.2.6 Model and source boundary . . . 79
5.2.7 Sources . . . 80

6 Credit risk **82**

6.1 Default, hazard, and survival **85**
6.1.1 What you will be able to do . . . 85
6.1.2 Survival probability: no default by time t . . . 85
6.1.3 Interval default probability: a decrease in survival . . . 86
6.1.4 Constant hazard rate: an intensity, not a cumulative probability . . . 86
6.1.5 Check your understanding . . . 87
6.1.6 Model boundary and review note . . . 88

6.2 Recovery and one-period risky present value **90**
6.2.1 What you will be able to do . . . 90
6.2.2 Name the recovery convention before calculating . . . 90
6.2.3 Recovery and loss given default are complements here . . . 91
6.2.4 Calculate the expected maturity amount, then discount it . . . 91

6.2.5 Check your understanding . . . 92
6.2.6 Model boundary and review note . . . 93

7 CDS 94

7.1 CDS premium and protection legs 97
7.1.1 What a CDS is . . . 97
7.1.2 Start from disjoint events . . . 97
7.1.3 The hazard rate gives the probabilities in the formulas . . . 98
7.1.4 Default in the middle of a period . . . 98
7.1.5 Exact premium-leg expectation in this model . . . 99
7.1.6 Exact protection-leg expectation in this model . . . 99
7.1.7 Closed forms and the zero-upfront par spread . . . 100
7.1.8 Generalizing the premium leg: an extinguishing swap . . . 101
7.1.9 Why there is no volatility input here . . . 102
7.1.10 Both legs for one contract . . . 102
7.1.11 Check your understanding . . . 102
7.1.12 Model boundary and review note . . . 104

7.2 CDS market-standard quote and upfront 106
7.2.1 What you will be able to do . . . 106
7.2.2 Three quantities with three different roles . . . 106
7.2.3 First decide which quantities are inputs and which are solved . . . 107
7.2.4 Why one upfront amount is enough at inception . . . 107
7.2.5 Risk-neutral does not mean risk-free . . . 108
7.2.6 Convert MSQ to signed upfront . . . 108
7.2.7 Convert signed upfront back to MSQ . . . 109
7.2.8 When MSQ and par spread are equal . . . 109
7.2.9 What a real standard converter adds . . . 110
7.2.10 Check your understanding . . . 111
7.2.11 Model boundary and review note . . . 112

7.3 Credit curve and market observables 114
7.3.1 What you will be able to do . . . 114
7.3.2 A traded price is an expectation . . . 114
7.3.3 What the market shows, the model assumes, and the model solves . . . 115
7.3.4 The credit curve is two objects . . . 115
7.3.5 Transforming marks is not fitting the curve . . . 117
7.3.6 A note on jargon . . . 119
7.3.7 Check your understanding . . . 119
7.3.8 Model boundary and review note . . . 121

7.4 Equivalent notional and quote risk 125
7.4.1 What you will be able to do . . . 125
7.4.2 One market level moves everything . . . 125
7.4.3 Risky DV01 by bump and reprice . . . 126
7.4.4 The liquid-tenor equivalent . . . 126

7.4.5 What an equivalent-ratio hedge does not do 128
7.4.6 Check your understanding 129
7.4.7 Model boundary and review note 130

A Notation **132**

B Abbreviations **146**

List of Figures

1.5.1 One-payment discounting explorer 21

2.4.1 Bond price–yield explorer 45

7.1.1 CDS leg explorer 103

List of Tables

1.3.1 One rate in three representations 13

1.7.1 Scenarios of the worked example 27

2.4.1 Signs of a yield shock and the price change 44

7.2.1 Inputs and solved quantities of each calculation 107

7.3.1 Observed, contractual, convention, and solved quantities 116
7.3.2 CDS desk jargon 120

List of Diagrams

4.5.1 Backward-induction order 71
5.1.1 European bond-option valuation flow 75
5.2.1 Issuer-default knockout valuation flow 79

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?

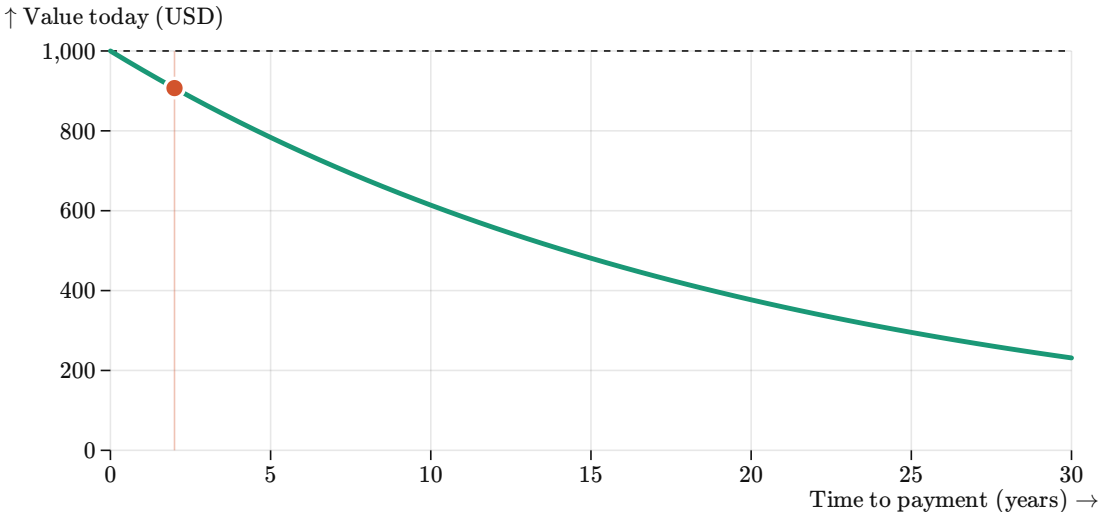


Figure 1.5.1 **One-payment discounting explorer.** The value today, in USD, of one future payment for every time to payment from 0 to 30 years, at one annual rate with annual compounding. The dot and the vertical line mark the time to payment shown; the dashed line marks the undiscounted payment. Shown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

Part 2

Bonds

Part 1 gave a timeline of signed cash flows, rate quotes with their compounding, discount factors, and present value as a sum of discounted dated amounts. A fixed-rate bond is the first contract the course values with these tools: its contract terms fix the promised payments, and its price is their present value. The last chapter, on bond forwards, also uses the forward contract of Part 4.

This part values a simplified fixed-rate bond with deterministic promised payments. The first chapter turns face value, coupon rate, payment frequency, and maturity into the payment schedule. The next two chapters price the bond, first from one discount factor per payment and then from a single yield to maturity. The fourth chapter explains why the price decreases when the yield increases and why the price-yield curve is not linear. The last two chapters add trading details: accrued interest and the clean and dirty price at settlement, and the delivery price of a bond forward.

The argument

1. The contract terms fix every promised payment
2. The price is the present value of those payments
3. A yield to maturity summarizes the discounting in one rate
4. A higher yield gives a lower price, along a curve
5. Accrued interest separates the quoted clean price from the cash dirty price
6. Cash-and-carry gives a bond forward delivery price on the dirty-price basis

After this part you can

- Identify face value, annual coupon rate, coupon frequency, maturity, and principal redemption in a simplified fixed-rate bond description.
- Calculate each level coupon as face value times annual coupon rate divided by payments per year.
- Construct the coupon and principal cash-flow schedule of a simplified fixed-rate bond.
- Price a simplified fixed-rate bond by multiplying each promised cash flow by its stated discount factor and summing the results.
- Interpret yield to maturity as the single nominal annual rate, compounded at coupon frequency, that reproduces price from promised cash flows in the simplified model.
- Price a simplified fixed-coupon bond from a nominal annual yield compounded at the coupon frequency.
- Explain why the price of positive fixed cash flows moves inversely to yield while discounting denominators remain positive.
- Recognize that equal yield changes generally produce unequal price changes in the simplified fixed-cash-flow model.
- Calculate actual/actual coupon-period accrued interest and convert consistently between clean quoted price and dirty invoice price.

- Calculate a no-arbitrage bond forward delivery price from the current dirty bond price, the present value of pre-delivery coupons, and the delivery-date discount factor.

Builds on

- Foundations
- Derivative foundations

Leads to

- Bond options

Chapter 2.1

Fixed-rate bond contract and cash flows

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/bonds/fixed-rate-contract-and-cash

2.1.1 What you will be able to do

After this lesson, you should be able to:

- identify the terms that determine a toy fixed-rate bond's payments;
- calculate the level coupon amount;
- construct the complete promised cash-flow schedule.

2.1.2 What a bond is

A bond is a debt claim. The issuer is the borrower; the bondholder is the lender who owns the claim. In the fixed-rate example used here, the issuer promises a sequence of coupon payments and repayment of principal at maturity. The price paid for the bond is a different quantity from the promised cash flows.

The word "promises" is important. A contract can specify an amount without a guarantee that the issuer pays it. This first bond lesson constructs the contractual schedule. The credit-risk lessons introduce default probability and recovery later. Tuckman and Serrat define a coupon bond by its coupon rate, its maturity, and its face amount (also called the par amount or the principal amount). They show the coupons and the principal on the cash-flow timeline[1].

2.1.3 Read the contract before valuing it

Write the simplified contract's face-value as F, its annual-coupon-rate as c, its bond-payment-frequency as m_{B}, and its maturity-time as T. Write the number-of-bond-payments as n and index those dates with the bond-payment-index k. The matching payment-time is t_k, measured from valuation-time 0.

Write the level coupon-payment as C and the promised bond-cash-flow on date k as CF_k^{bond}.

The number of payment dates is the payment frequency multiplied by the maturity:

$$n = m_{\mathrm{B}} T \tag{2.1.1}$$

The bond payment frequency is a different concept from the compounding-frequency m of the earlier lessons. So it has a different symbol, with the subscript B.

2.1.4 Calculate the coupon

The level coupon payment is the annual coupon rate multiplied by the face value, divided by the payment frequency:

$$C = \frac{cF}{m_{\mathrm{B}}} \tag{2.1.2}$$

The coupon rate is applied to the face value, not to the market price of the bond.

2.1.5 Construct every promised payment

For the bondholder, the promised bond cash flow on date k is the coupon, plus the face value on the final date:

$$CF_k^{\mathrm{bond}} = C + \mathbf{1}_{\{k=n\}} F \tag{2.1.3}$$

The indicator $\mathbf{1}_{\{k=n\}}$ equals one on the final date and zero on every other date. So the face value is paid only on the final date.

On the regular schedule, the payment time of date k, measured from valuation time, is:

$$t_k = \frac{k}{m_{\mathrm{B}}}, \qquad 0 \text{ is the valuation time} \tag{2.1.4}$$

Example 2.1.1: Coupon schedules by payment frequency

Annual coupons

A two-year bond has a face value of USD 1,000, an annual coupon rate of 5%, and one payment per year. The coupon is:

$$C = \frac{0.05 \times 1000}{1} = \text{USD } 50$$

Payment index k	Payment time t_k, years	Coupon, USD	Principal, USD	Total, USD
1	1.0	50	0	50
2	2.0	50	1,000	1,050

Semiannual coupons

An 18-month bond has a face value of USD 1,000, an annual coupon rate of 4%, and two payments per year. The coupon and the number of payments are:

$$C = \frac{0.04 \times 1000}{2} = \text{USD } 20, \qquad n = 2 \times 1.5 = 3$$

Its payments at 0.5, 1.0, and 1.5 years are USD 20, USD 20, and USD 1,020.

Quarterly coupon amount

For a face value of USD 500, an annual coupon rate of 4.8%, and quarterly payments, the coupon is:

$$C = \frac{0.048 \times 500}{4} = \text{USD } 6$$

The maturity determines the number of payments. Under the level-coupon assumption, the maturity does not change the coupon amount.

2.1.6 Check your understanding

These items cover term interpretation, coupon arithmetic, and schedule construction. Each question stays collapsed until you open it; answers and explanations appear once you check.

Knowledge check 2.1.1: Fixed-rate bond contract and cash flows

1. A simplified bond has USD 1,000 face value, a 5% annual coupon rate, two coupon payments per year, and two years to maturity. Which field determines the principal redeemed at maturity under the lesson convention?

   (a) Face value
   (b) Annual coupon rate
   (c) Coupon frequency

2. A simplified fixed-rate bond's market price changes while its contract is unchanged. Which listed field remains a contractual input to promised coupon cash flows?

   (a) The annual coupon rate
   (b) The new market price
   (c) The size of the market-price change

3. Under the lesson convention, a bond has USD 1,000 face value, a 6% annual coupon rate, and two equal coupon payments per year. What is each coupon amount in USD?

4. Under the lesson convention, a bond has USD 2,000 face value, a 4.5% annual coupon rate, and two equal coupon payments per year. What is each coupon amount in USD?

5. From the holder's perspective, what are the promised cash flows of a two-year bond with USD 1,000 face value, a 5% annual coupon rate, and one coupon payment per year?

   (a) USD +50 after year one and USD +1,050 after year two
   (b) USD +50 after each year and no principal
   (c) USD +1,100 after year two only

6. From the holder's perspective, what are the promised cash flows of an 18-month bond with USD 1,000 face value, a 4% annual coupon rate, and semiannual coupons?

(a) USD +20, USD +20, and USD +1,020 at six-month intervals

(b) USD +40, USD +40, and USD +1,040 at six-month intervals

(c) USD +1,020 at every coupon date

Answers: *end of the chapter.*

2.1.7 Model boundary and review note

This lesson describes promised cash flows only. It does not assign a price, yield, payment probability, or expected recovery.

The contract terms and the level-coupon schedule follow Tuckman & Serrat[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). Ch. 1 §1.1, printed p. 50 and Table 1.1.

[2] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). §1.1, "Government Coupon Bonds": a bond is set by coupon rate, maturity date, and face (par, principal); each semiannual coupon is half the annual rate times face, with principal repaid at maturity (Table 1.1).

Answers to the knowledge checks

Knowledge check 2.1.1: Fixed-rate bond contract and cash flows

1. **(a)** Face value. The simplified contract redeems principal at face value; the coupon rate and frequency determine periodic coupons.

2. **(a)** The annual coupon rate. In the simplified model, promised coupons follow the contractual face value, annual coupon rate, and frequency rather than the current market price.

3. 30 (accepted within $\pm 1e-9$). Coupon = face value $\times$ annual coupon rate / payments per year = 1,000 $\times$ 0.06 / 2 = 30.

4. 45 (accepted within $\pm 1e-9$). Coupon = 2,000 $\times$ 0.045 / 2 = 45.

5. **(a)** USD +50 after year one and USD +1,050 after year two. The coupon is 5% of face value each year; principal is added to the final coupon.

6. **(a)** USD +20, USD +20, and USD +1,020 at six-month intervals. Each semiannual coupon is 1,000 $\times$ 4% / 2 = 20, with principal paid only at maturity.

Chapter 2.2

Bond price from discount factors

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/bonds/price-from-discount-factors/

2.2.1 What you will be able to do

After this lesson, you should be able to calculate a toy bond's price from its promised cash-flow schedule and input discount factors.

2.2.2 Bond pricing is a present-value application

Use the bond-payment-index k to select one of the number-of-bond-payments n. Write the promised bond-cash-flow as CF_k^{bond}, its payment-time as t_k, and its matching discount-factor as $D(0, t_k)$. Write the bond-price as P_0. The bond price is the present-value at valuation-time 0 of the promised payments. So the bond price is the schedule sum (1.5.1), applied to the promised payments of a bond:

$$P_0 = \sum_{k=1}^{n} CF_k^{\text{bond}} D(0, t_k) \tag{2.2.1}$$

In this lesson, "price" means this present value. It does not mean a clean price, a quoted price, or a probability. Those concepts are outside this lesson.

Example 2.2.1: Bond prices from discount factors

Two discount factors

The annual-coupon bond from the preceding lesson pays USD 50 after one year and USD 1,050 after two years. Suppose the matching discount factors are 0.96 and 0.91.

Payment index k	Cash flow, USD	Discount factor	Present value, USD
1	50	0.96	48.00
2	1,050	0.91	955.50

The bond price is the sum of the present values:

$$P_0 = 50(0.96) + 1050(0.91) = \text{USD } 1{,}003.50$$

The price is above USD 1,000 because, with these two discount factors, the present value of the promised payments is larger than the face value. This lesson does not use a yield.

Unfamiliar schedule

Consider payments of USD 20 at 0.5 years, USD 20 at 1.0 year, and USD 1,020 at 1.5 years, with discount factors 0.98, 0.95, and 0.92. The bond price is:

$$P_0 = 20(0.98) + 20(0.95) + 1020(0.92) = 19.60 + 19.00 + 938.40 = \text{USD } 977.00$$

Zero-coupon boundary

If every coupon is zero, the only payment is the face-value F at maturity-time T. The bond price is then:

$$P_0 = FD(0, T)$$

This is (2.2.1) with one payment. It is not a separate rule.

2.2.3 Error pattern: discount after adding

For the two-payment example, compare a wrong calculation, which applies one discount factor to the sum of the payments, with the calculation in (2.2.1):

$$\begin{aligned} (50 + 1050)(0.91) &= 1001.00 \\ 50(0.96) + 1050(0.91) &= 1003.50 \\ 1001.00 &\neq 1003.50 \end{aligned} \tag{2.2.2}$$

One common discount factor is correct only when the model assigns that discount factor to every payment.

2.2.4 Check your understanding

This set includes one coupon-bond calculation and one transfer schedule. Each question stays collapsed until you open it; answers and explanations appear once you check.

Knowledge check 2.2.1: Bond price from discount factors

1. From the holder's perspective, a simplified bond pays USD +50 at t = 1 year and USD +1,050 at t = 2 years. The matching discount factors are 0.95 and 0.90. What is its time-zero price in USD?

2. From the holder's perspective, a simplified bond pays USD +20, USD +20, and USD +1,020 at three successive dates. Their matching discount factors are 0.98, 0.95, and 0.91. What is its time-zero price in USD?

Answers: end of the chapter.

2.2.5 Model boundary and review note

The discount factors are inputs; they are not calibrated. This lesson has no credit risk, although the course is about credit products.

The pricing rule (2.2.1) follows Tuckman & Serrat[1]; Hull discounts each cash flow at its own zero rate the same way[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, each bond price as its cash flows times the matching discount factors (eqs. 1.1-1.3), checked against market prices in Table 1.4.

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

Answers to the knowledge checks

Knowledge check 2.2.1: Bond price from discount factors

1. 992.5 (accepted within $\pm 1e-9$). Price $= 50 \times 0.95 + 1{,}050 \times 0.90 = 47.5 + 945 = 992.5$.
2. 966.8 (accepted within $\pm 1e-9$). Price $= 20 \times 0.98 + 20 \times 0.95 + 1{,}020 \times 0.91 = 19.6 + 19 + 928.2 = 966.8$.

Chapter 2.3

Yield to maturity as a single-rate summary

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/bonds/yield-to-maturity/

2.3.1 What you will be able to do

After this lesson, you should be able to:

- interpret yield to maturity as a model-dependent single-rate summary;
- state its annualization and compounding convention;
- calculate the price from an input yield without changing the promised cash flows.

2.3.2 What yield means in this toy model

Use the bond-payment-index k to select one of the number-of-bond-payments n, and write its promised bond-cash-flow as CF_k^{bond}. Write the bond-price as P_0 and the bond-payment-frequency as m_{B}.

Write yield-to-maturity as $y^{(m_{\text{B}})}$. It is the one nominal annual rate, compounded at the bond payment frequency, that makes the discounted promised payments equal the bond price:

$$P_0 = \sum_{k=1}^{n} \frac{CF_k^{\text{bond}}}{\left(1 + y^{(m_{\text{B}})}/m_{\text{B}}\right)^k} \tag{2.3.1}$$

The symbol $y^{(m_{\text{B}})}$ could be mistaken for a spot rate, an effective annual rate, an expected return, or another quote. In this course it means only the bond yield defined above.

2.3.3 Coupon rate and yield have different roles

- The coupon rate determines the promised coupon-payment C in the contract.
- The input yield determines the denominators in the price formula.

A change of the input yield does not change the coupons or the face-value F of the bond.

Example 2.3.1: Pricing from a yield to maturity

Annual payments

A two-year bond with an annual coupon rate of 5% and a face value of USD 1,000 pays USD 50 and USD 1,050. At an annual yield of 6% with annual compounding, the price is:

$$P_0 = \frac{50}{1.06} + \frac{1050}{1.06^2} \approx \text{USD } 981.67$$

Semiannual compounding

For the same face value and coupon rate over two years, but with semiannual payments, $C = 25$, $m_{\mathrm{B}} = 2$, and $n = 4$. A nominal annual yield of 6% gives a periodic rate of 3%, and the price is:

$$P_0 = \frac{25}{1.03} + \frac{25}{1.03^2} + \frac{25}{1.03^3} + \frac{1025}{1.03^4} \approx \text{USD } 981.41$$

The annual and the semiannual example use the same yield number, 6%, but different payment times and different compounding.

Zero-coupon special case

With no coupons and maturity-time T, the price is:

$$P_0 = \frac{F}{\left(1 + y^{(m_{\mathrm{B}})}/m_{\mathrm{B}}\right)^{m_{\mathrm{B}}T}}$$

For USD 1,000 due in three years at 4% with annual compounding, the price is approximately USD 889.00.

2.3.4 Check your understanding

These items separate yield interpretation from annual and periodic-compounding calculations. Each question stays collapsed until you open it; answers and explanations appear once you check.

Knowledge check 2.3.1: Yield to maturity

1. In the simplified lesson model, what is yield to maturity?
   (a) The single nominal annual rate, compounded at coupon frequency, that reproduces price from promised cash flows
   (b) An additional cash flow paid by the issuer at maturity
   (c) The currency amount of each coupon
2. Two bonds have different promised cash-flow amounts and maturities, but the same coupon frequency and the same quoted yield to maturity under the simplified nominal convention. Which statement follows?
   (a) The same summary rate is used, but each bond's own cash flows still determine its price
   (b) The two bonds must have the same price
   (c) The two bonds must have identical coupon cash flows

3. Price in USD a two-year bond with USD 1,000 face value, a 5% annual coupon rate, and one coupon payment per year at a 6% nominal annual yield compounded annually. Assume valuation on a coupon date and no accrued interest.

4. Price in USD a two-year bond with USD 1,000 face value, a 5% annual coupon rate, and two coupon payments per year at a 6% nominal annual yield compounded semiannually. Assume valuation on a coupon date and no accrued interest.

Answers: *end of the chapter.*

2.3.5 Model boundary and review note

This lesson calculates the price from an input yield. It does not solve numerically for the yield from a price. A real bond implementation also needs schedules, day counts, accrued interest, clean and dirty prices, and possibly credit and option models.

The single-rate definition and the $2T$-period pricing formula follow Tuckman & Serrat[1]; Hull gives the same definition, solved iteratively[2]. FINRA defines it as the return to an investor who buys at the market price and holds to maturity[3]. 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). §3.2, "Yield to Maturity": the single rate that discounts a bond's cash flows to its market price (eq. 3.5 and the general formula), also its internal rate of return.

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

[3] FINRA, Understanding Bond Yield and Return. "Yield to maturity" section. https://www.finra.org/investors/insights/bond-yield-return

Answers to the knowledge checks

Knowledge check 2.3.1: Yield to maturity

1. **(a)** The single nominal annual rate, compounded at coupon frequency, that reproduces price from promised cash flows. The toy model defines yield to maturity as one rate that discounts all promised cash flows to the stated price under the stated compounding convention.

2. **(a)** The same summary rate is used, but each bond's own cash flows still determine its price. A common yield does not erase differences in amount or timing; the model discounts each bond's own promised cash flows.

3. 981.6660733357066 (accepted within ± 0.01). $P_0 = \frac{50}{1.06} + \frac{1{,}050}{1.06^2}$.

4. 981.4145079859479 (accepted within ± 0.01). Each coupon is USD 25 and the periodic yield is $0.06/2 = 0.03$, so $P_0 = \frac{25}{1.03} + \frac{25}{1.03^2} + \frac{25}{1.03^3} + \frac{1{,}025}{1.03^4}$.

Chapter 2.4

The bond price-yield relationship

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/bonds/price-yield-relationship/

2.4.1 What you will be able to do

After this lesson, you should be able to:

- explain why the price of a bond with fixed positive cash flows decreases when the yield increases;
- hold contract cash flows fixed while changing only the discounting input;
- recognize that the price-yield relationship is curved rather than linear.

2.4.2 The price decreases because every denominator increases

Write the bond-payment-frequency as m_{B}. In this lesson, y denotes the yield-to-maturity $y^{(m_{\mathrm{B}})}$ without its superscript, as the argument of a function. Write the price-yield-curve as $P(y)$.

Use the bond-payment-index k to select one of the number-of-bond-payments n, and write each promised positive bond-cash-flow as CF_k^{bond}. The price at yield y is the sum of the promised cash flows, each discounted at the yield:

$$P(y) = \sum_{k=1}^{n} \frac{CF_k^{\mathrm{bond}}}{\left(1 + y/m_{\mathrm{B}}\right)^k} \tag{2.4.1}$$

Assume that the yield stays non-negative and that every cash flow stays fixed. If the yield increases, every denominator increases. So every discounted cash flow decreases, and the bond-price decreases. Table 2.4.1 also assumes that the new yield is non-negative.

Define the local yield-shock Δy as new yield minus old yield, and define the local price-change ΔP as new price minus old price.

The table holds only if the promised cash flows stay fixed.

Example 2.4.1: Price-yield comparisons

Three points form a curve

Consider a five-year bond with an annual coupon rate of 5%, semiannual payments, and a face value of 100. Its prices at three yields are:

Table 2.4.1 **Signs of a yield shock and the price change.** For a rise and a fall in the yield, the direction of every discount denominator and the sign of the price change, with the promised cash flows held fixed.

Yield shock	Denominators	Price change
$\Delta y > 0$	Increase	$\Delta P < 0$
$\Delta y < 0$	Decrease	$\Delta P > 0$

Nominal annual yield	Price per 100 face
5%	100.0000
6%	95.7349
7%	91.6834

The price decreases as the yield increases. Each yield increase is one percentage point, but the two price changes are different: 4.2651 and 4.0515. So the three points are not on a straight line.

Equal shocks, unequal price changes

Start from a yield of 5%:

- a yield decrease to 4% increases the price from 100.0000 to 104.4913, a change of +4.4913;
- a yield increase to 6% decreases the price from 100.0000 to 95.7349, a change of −4.2651.

The two yield shocks have equal size, but the absolute price changes are different. This difference is a result of the curvature of the price-yield curve. This lesson does not calculate duration, DV01, or convexity.

Hold other inputs fixed

Suppose the issuer changes a coupon, a default occurs, or the settlement date changes. Then more than the yield has changed, and the comparison no longer isolates the effect of the yield.

One tested domain function calculates every price in Figure 2.4.1. The notation panel only explains the symbols; it does not calculate prices.

2.4.3 Check your understanding

These direct and transfer items cover direction and nonlinearity. Each question stays collapsed until you open it; answers and explanations appear once you check.

Knowledge check 2.4.1: Bond price-yield shape

1. Within the lesson's non-negative-yield domain, all else equal, what happens to a fixed bond's price when its yield rises?

   (a) The price falls

   (b) The price rises

   (c) The price is unchanged

2. A bond's yield falls from 6% to 5% while promised cash flows stay fixed. Which explanation

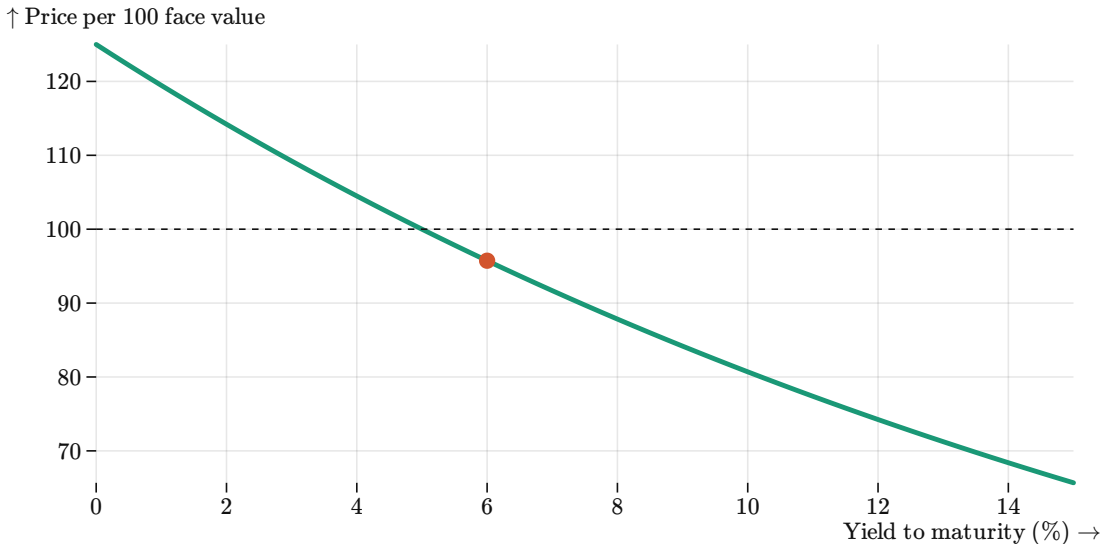


Figure 2.4.1 **Bond price–yield explorer.** The price per 100 face value of one fixed-coupon bond at every yield to maturity from 0% to 15%. The dot marks the price at the yield shown; the dashed line marks par, a price of 100. Shown for a 5% annual coupon rate, 2 coupon payments a year and a 5-year maturity: at a 6% yield to maturity, the price is 95.73 per 100 face value.

is correct?

(a) Discounting denominators shrink, so price rises

(b) The contractual coupons grow, so price rises

(c) The principal shrinks, so price falls

3. A fixed bond's prices at yields of 4%, 5%, and 6% are 103.0, 100.0, and 97.2. What do the unequal changes for equal one-percentage-point yield steps show?

(a) The price-yield relationship is not a straight line

(b) Price is independent of yield

(c) The promised cash flows must change with yield

4. For three equally spaced yields, the middle bond price differs from the average of the two endpoint prices. What is the best interpretation?

(a) The sampled price-yield relationship has curvature

(b) The sampled price-yield relationship is exactly linear

(c) Yield has no relationship to price

Answers: *end of the chapter.*

2.4.4 Model boundary and review note

The plotted curve summarizes a toy model with one yield. It is not a credit curve and not a recommendation. It omits accrued interest, term-structure effects, default, recovery, liquidity, tax,

funding, and embedded options.

The inverse direction and the curvature both follow Tuckman & Serrat[1]; Hull notes the same negative price-yield relationship[2]. FINRA states that price and yield are inversely related[3]. 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). §4.1, "Price-Rate Curves", Figure 4.2, prices decrease as rates rise; §4.5, "Convexity", price-rate curves of coupon bonds are convex.

[2] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). §4.8, Duration.

[3] FINRA, Understanding Bond Yield and Return. "Price and yield" section. https://www.finra.org/investors/insights/bond-yield-return

Answers to the knowledge checks

Knowledge check 2.4.1: Bond price-yield shape

1. **(a)** The price falls. A higher yield produces smaller discount factors for every positive future cash flow.

2. **(a)** Discounting denominators shrink, so price rises. The promised coupons do not change; the change comes from discounting them at a lower yield.

3. **(a)** The price-yield relationship is not a straight line. The first equal yield step changes price by 3.0 while the second changes it by 2.8, so the slope is not constant.

4. **(a)** The sampled price-yield relationship has curvature. For a straight line, the value at the midpoint equals the average of endpoint values; a difference is evidence of curvature.

Chapter 2.5

Settlement, accrued interest, and clean versus dirty price

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/bonds/settlement-clean-and-dirty-p

2.5.1 The price paid is not always the price quoted

A bond price quote commonly shows the clean-bond-price, also called the flat or quoted price. The buyer pays the dirty-bond-price, also called the full, cash, or invoice price. The dirty price is the clean price plus the accrued-interest:

$$P^{\text{dirty}} = P^{\text{clean}} + AI \tag{2.5.1}$$

On a coupon date, the accrued interest is zero, so the two prices are equal. For this reason, the earlier bond lessons, which settle on a coupon date, use one bond-price symbol.

2.5.2 Actual/actual within one coupon period

Under the convention of this lesson, let the local coupon-period-elapsed-days be n_{elapsed} and the local coupon-period-total-days be n_{period}. For one coupon-payment C, the accrued interest is the coupon multiplied by the fraction of the coupon period that has elapsed:

$$AI = C\frac{n_{\text{elapsed}}}{n_{\text{period}}} \tag{2.5.2}$$

Consider the coupon period from February 15, 2021 to August 15, 2021, which has 181 days. Settlement on May 17, 2021 is 91 actual days after the start of the period. With a coupon of USD 31.25, (2.5.2) gives the accrued interest:

$$AI = 31.25\left(\frac{91}{181}\right) = 15.71132597 \text{ USD per bond} \tag{2.5.3}$$

The domain code checks that the dates are valid ISO dates and counts UTC calendar days. It is not a production schedule engine: business-day rules, ex-coupon periods, stubs, and other day counts require separate conventions.

2.5.3 Why this matters for derivatives

The cash amount paid at the settlement of a bond forward or a bond option depends on whether its price is clean or dirty. If a strike is quoted as a clean price and the modeled bond price is a dirty price, add the accrued interest at the exercise date to the strike before the payoff is calculated. Also, each coupon paid before delivery must be assigned to the party that owns the bond on the coupon date.

Knowledge check 2.5.1: Bond settlement

1. A coupon payment is USD 31.25. Settlement is 91 actual days into an actual 181-day coupon period. Calculate accrued interest in USD.

2. A bond has a clean cash price of USD 9,178.125 and accrued interest of USD 15.711. Calculate its dirty invoice price in USD.

Answers: end of the chapter.

2.5.4 Sources

The identity between the clean and the dirty price (2.5.1) and the actual/actual calculation (2.5.2) follow Tuckman and Serrat, Chapter 1 §1.6 [1]. Hull's European bond option discussion emphasizes matching cash versus quoted strike and bond prices [2].

References

[1] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022).

[2] Hull, Options, Futures, and Other Derivatives (8th ed., 2012).

Answers to the knowledge checks

Knowledge check 2.5.1: Bond settlement

1. 15.71132596685083 (accepted within ± 0.000001). Accrued interest is $31.25 \times 91 / 181 =$ USD 15.71132597.

2. 9193.836 (accepted within ± 0.000001). Dirty price = clean price + accrued interest = USD 9,193.836.

Chapter 2.6

Bond forwards

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/bonds/bond-forwards/

2.6.1 The forward contract

The holder of a long bond forward must pay the forward-delivery-price and receive the named bond at forward-delivery-time. The holder of the forward does not own the bond before delivery. So coupons paid before delivery belong to the current owner of the bond.

2.6.2 Cash-and-carry price

At valuation-time, start with the current dirty-bond-price P^{dirty}. Subtract the local predelivery-coupon-present-value I_0^{pre}. The difference is the prepaid forward value: the value today of receiving the bond at delivery. Dividing the prepaid forward value by the discount factor to delivery converts it to an amount at delivery. This amount is the fair bond-forward-price:

$$F^B_{0,T_{\text{fwd}}} = \frac{P^{\text{dirty}} - I_0^{\text{pre}}}{D(0, T_{\text{fwd}})} \tag{2.6.1}$$

For a dirty price of USD 960, a present value of the predelivery coupons of USD 95.45, and a discount factor to delivery of 0.92, (2.6.1) gives the bond forward price:

$$F^B_{0,T_{\text{fwd}}} = \frac{960 - 95.45}{0.92} = 939.72826087 \text{ USD per bond} \tag{2.6.2}$$

An existing contract can have a different delivery price. The earlier lesson on forward value calculates the value of such a contract. This lesson calculates the fair delivery price of a new contract.

2.6.3 Basis consistency

A clean price quote must first be converted to a dirty price. In the same way, the accrued interest at the delivery date must be added to a clean delivery price or a clean strike. If clean and dirty prices are mixed, the result can look like a pricing discrepancy, but it is only a mismatch of conventions.

Knowledge check 2.6.1: Bond forward

1. A bond's current dirty price is USD 960. The present value of coupons paid before delivery is USD 95.45, and D(0,T) = 0.92. Calculate the fair dirty bond forward delivery price in USD.

2. A bond's dirty spot price is USD 1,025, pre-delivery coupons have present value USD 24, and D(0,T) = 0.97. Calculate the fair dirty delivery price in USD.

Answers: *end of the chapter.*

2.6.4 Sources

The known-income cash-and-carry relation (2.6.1) and its coupon-bond application follow Hull, Chapter 5 §§5.3–5.5; Hull also treats the dirty-price basis of the bond forward and the strike in Chapter 28 §28.1 [1].

References

[1] Hull, Options, Futures, and Other Derivatives (8th ed., 2012).

Answers to the knowledge checks

Knowledge check 2.6.1: Bond forward

1. 939.7282608695652 (accepted within ±0.000001). The prepaid forward is 960 - 95.45 = 864.55. Divide by 0.92 to obtain USD 939.72826087.

2. 1031.958762886598 (accepted within ±0.000001). The fair delivery price is (1,025 - 24) / 0.97 = USD 1,031.95876289.

Part 3

Rates and curves

Part 1 defined a discount factor for one future date from a rate quote, and Part 2 applied one discount factor to each promised bond payment. This part collects the discount factors of many dates into one deterministic discount curve. The credit-curve chapter of Part 7 uses that curve as the risk-free discount input of a CDS.

This part has one chapter. It reads a deterministic discount curve that is input at a set of dates, interpolates log-linearly between those dates, and calculates the forward discount factor between two future dates as the ratio of their discount factors from valuation time. The curve is an input: the chapter does not calibrate it to market prices, and it does not extrapolate beyond the last input date.

The argument

1. A discount curve holds one discount factor for each input date
2. Log-linear interpolation gives the dates between the inputs
3. The ratio of two discount factors discounts from one future date to another

After this part you can

- Interpret a dated set of positive valuation-time discount factors as a deterministic curve, including its interpolation and no-extrapolation conventions.
- Calculate the deterministic discount factor between two future model times from their valuation-time discount factors.

Builds on

- Foundations

Leads to

- CDS

Chapter 3.1

Discount curve and forward discounting

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/rates/discount-curve-and-forward-d

3.1.1 What you will be able to do

After this lesson, you should be able to read a deterministic discount curve, state its interpolation boundary, and calculate a discount factor between two future dates.

3.1.2 A curve is a dated collection, not one yield

In the earlier lessons, the discount-factor $D(0,t)$ gave the value at valuation-time of a deterministic payment at one payment-time. A discount curve gives the discount factor for many payment times. Each node of the curve is one payment time with its discount factor. Every node must have the same currency, the same valuation date, and the same collateral and funding basis. Nodes with different bases do not form one consistent curve.

The curve in this lesson accepts only positive discount factors at its nodes, and it sets $D(0,0) = 1$. Between two nodes, the logarithm of the discount factor is interpolated linearly. The curve is not extrapolated beyond its final node. Log-linear interpolation is a convention of this lesson. The lesson does not claim that it is the only correct interpolation.

A discount factor can be larger than one. For example, $D(0,1) = 1.01$ implies a negative continuously compounded one-year zero rate. The curve requires positive discount factors, but it does not require discount factors below one.

3.1.3 Discount between future times

Let the local forward-discount-start-time t be earlier than the local forward-discount-end-time T. The forward-discount-factor is defined as the ratio of the two discount factors from valuation time:

$$Z(t,T) = \frac{D(0,T)}{D(0,t)} \tag{3.1.1}$$

Its unit is currency at time t per unit of currency at time T. If $D(0,1) = 0.98$ and $D(0,3) = 0.94$, the forward discount factor from year 1 to year 3 is:

$$Z(1,3) = \frac{0.94}{0.98} = 0.9591836735 \tag{3.1.2}$$

The forward discount factor has its own symbol, Z. The symbol D keeps its meaning: a discount factor from valuation time.

The ratio in (3.1.1) is needed for multi-period valuation. Each lattice node needs the discount factor over its next period. A discount factor from time zero is not the correct discount factor for that period.

3.1.4 Scope boundary

The curve is an input. It is deterministic and not calibrated. Later lessons with state-dependent discounting give a one-period discount factor at each node. They do not assume that today's forward discount factor is realized in every future state.

Knowledge check 3.1.1: Discount curve and forward discounting

1. What does a deterministic curve node D(0,3) = 0.94 mean?
   (a) At time zero, 0.94 currency units buy one currency unit paid at model-year 3
   (b) At model-year 3, 0.94 currency units are paid at time zero
   (c) The annual rate is 94% under every compounding convention
2. Can a valid input curve contain D(0,1) = 1.01 in this lesson?
   (a) Yes; a positive factor above one corresponds to a negative continuously compounded zero rate
   (b) No; every discount factor is mathematically restricted below one
   (c) Only if issuer default is certain
3. Given D(0,1) = 0.98 and D(0,3) = 0.94, calculate the forward discount factor Z(1,3).
4. Given D(0,2) = 0.95 and D(0,5) = 0.88, calculate the forward discount factor Z(2,5).

Answers: end of the chapter.

3.1.5 Sources

Discount-factor pricing and the law-of-one-price interpretation follow Tuckman and Serrat, Chapter 1 §§1.2–1.4 and Chapter 2 §§2.1 and 2.4 [1]. The later derivatives use of dated factors is consistent with Hull's interest-rate and forward-pricing chapters [2].

References

[1] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022).

[2] Hull, Options, Futures, and Other Derivatives (8th ed., 2012).

Answers to the knowledge checks

Knowledge check 3.1.1: Discount curve and forward discounting

1. **(a)** At time zero, 0.94 currency units buy one currency unit paid at model-year 3. A discount factor is a dated price ratio. It is not itself an annualized rate quotation.

2. **(a)** Yes; a positive factor above one corresponds to a negative continuously compounded zero rate. The domain requires positive factors, not factors below one. Negative rates can produce factors above one.

3. 0.9591836734693877 (accepted within $\pm 1e-9$). Z(1,3) = D(0,3) / D(0,1) = 0.94 / 0.98.

4. 0.9263157894736842 (accepted within $\pm 1e-9$). Z(2,5) = 0.88 / 0.95. The result is time-2 currency per one unit paid at time 5.

Part 4

Derivative foundations

Part 1 gave present value, no-arbitrage replication, and risk-neutral pricing as a discounted expectation over a finite set of states. This part applies those results to the first derivative contracts. Replication fixes the forward delivery price and put-call parity, and the risk-neutral expectation values an option, first on a one-period tree and then node by node on a multi-period lattice.

This part covers forwards and European options on one underlying. The first chapter separates a forward's obligation, its fair delivery price, and the value of an existing contract. The second chapter separates an option's right and expiry payoff from its current value, premium, and profit. The third chapter matches call and put portfolios to derive put-call parity. The last two chapters value an option: on a one-period binomial tree by replication and by risk-neutral weights, and on a multi-period lattice by backward induction with local discount factors.

The argument

1. A forward is an obligation, and its delivery price makes a new contract worth zero
2. An option is a right, and its payoff at expiry is not its value today
3. Matched call and put portfolios give put-call parity
4. A two-state payoff is replicated, or weighted by risk-neutral probabilities
5. Backward induction repeats the one-period step node by node

After this part you can

- Interpret the long and short obligations, delivery time, delivery price, and signed expiry payoff of a forward contract.
- Calculate the fair delivery price of a newly struck forward from spot value, present value of deterministic income before delivery, and the delivery-date discount factor.
- Calculate the signed value to the long of an existing forward from the current forward price, its fixed delivery price, quantity, and the delivery-date discount factor.
- Identify the underlying, call or put right, holder and writer, strike, expiry, and settlement perspective of a European option.
- Calculate the expiry payoff of a long or short European call or put from the underlying expiry value, strike, and quantity.
- Distinguish an option's current value or premium from its expiry payoff and from profit after accounting for the premium under a stated financing convention.
- Explain European put-call parity by matching the expiry cash flows of a call-minus-put portfolio and a forward-style position under stated income and discounting assumptions.
- Calculate a missing European call or put value from the other option value, prepaid forward value, strike, and expiry discount factor.
- Calculate the underlying units and risk-free cash position that replicate a claim's two one-period state payoffs.

- Calculate a European option's one-period binomial value using model-implied risk-neutral weights and deterministic discounting.
- Calculate node values on an input multi-period recombining pricing lattice by risk-neutral expectation and one-period node discounting.

Builds on

- Foundations

Leads to

- Bonds
- Bond options

Chapter 4.1

Forward contracts, delivery price, and value

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/derivatives/forward-contracts-and-

4.1.1 What you will be able to do

After this lesson, you should be able to:

- read the long and short obligations of a forward;
- distinguish delivery price, current forward price, and contract value;
- calculate a fair delivery price with deterministic income;
- calculate the signed value of an existing forward.

4.1.2 An obligation, not an option

The exchange occurs at the forward-delivery-time T_{fwd}. Under the physical-settlement convention, the long pays the forward-delivery-price K_{fwd} and receives one unit of the asset. If the derivative-underlying-value at delivery is $S_{T_{\text{fwd}}}$, the local long-forward-expiry-payoff per unit is the asset value minus the delivery price:

$$\Pi^{\text{long}}_{T_{\text{fwd}}} = S_{T_{\text{fwd}}} - K_{\text{fwd}} \tag{4.1.1}$$

The payoff of the short is the negative of the payoff of the long. Hull distinguishes this bilateral obligation from an option holder's exercise right and gives the corresponding long and short payoff directions[1].

4.1.3 Fair delivery price at inception

The forward-price $F_{0,T_{\text{fwd}}}$ is the delivery price that would make a new contract worth zero at valuation-time 0. Suppose the underlying is worth S_0 and the underlying-income-present-value paid before delivery is I_0. The long does not own the asset before delivery and does not receive the income. So the prepaid value of the asset delivered is $S_0 - I_0$.

Under the stated cash-and-carry assumptions, with the delivery-time discount-factor $D(0, T_{\text{fwd}})$, the law of one price gives the forward price:

$$F_{0,T_{\text{fwd}}} = \frac{S_0 - I_0}{D(0, T_{\text{fwd}})} \tag{4.1.2}$$

Hull derives the no-income and known-income forward prices by matching a spot purchase and financing strategy to the forward exchange [2].

For $S_0 = 100$, $I_0 = 4$, and $D(0, T_{\text{fwd}}) = 0.95$, the forward price is:

$$F_{0,T_{\text{fwd}}} = \frac{100 - 4}{0.95} = 101.0526316 \tag{4.1.3}$$

The delivery price is not a premium paid at inception. A new contract with a fair delivery price has signed value zero, because its delivery price equals the current forward price.

4.1.4 An existing forward can gain or lose value

After inception, K_{fwd} stays fixed, but the current forward price can change. For the local forward-contract-quantity N, the forward-contract-value to the long at time zero is the discounted difference between the current forward price and the delivery price:

$$V_0^{\text{fwd}} = N D(0, T_{\text{fwd}}) \left(F_{0,T_{\text{fwd}}} - K_{\text{fwd}} \right) \tag{4.1.4}$$

If $N = 10$, the current forward price is 98, the contract delivery price is 102, and the discount factor is 0.94, the value to the long is:

$$V_0^{\text{fwd}} = 10(0.94)(98 - 102) = -37.60 \tag{4.1.5}$$

The short's value is USD +37.60 under the same inputs. Hull derives this discounted forward-price difference for an existing contract [3].

Knowledge check 4.1.1: Forward contracts and value

1. What does the long side of a physically settled forward contract agree to do at delivery?
   (a) Buy the underlying for the fixed delivery price
   (b) Choose whether to buy, with no obligation
   (c) Receive the delivery price as an inception premium
2. A long forward has delivery price USD 80 and the underlying is worth USD 74 at delivery. What is the long's payoff per unit?
   (a) USD -6
   (b) USD +6
   (c) USD 0 because the long can decline delivery
3. An investment asset is worth USD 100, pays no income before delivery, and D(0,T) = 0.95. What fair delivery price gives a new forward zero value, in USD per unit?
4. A bond's current full price is USD 900. A known coupon before delivery has present value USD 40, and D(0,T) = 0.96. What is the fair full forward delivery price in USD?

5. The current forward price is USD 110, an existing long forward's delivery price is USD 105, and D(0,T) = 0.97. What is the contract value to the long per unit, in USD?

6. A long forward covers 10 units. The current forward price is USD 98 per unit, its fixed delivery price is USD 102, and D(0,T) = 0.94. What is its signed value to the long in USD?

Answers: *end of the chapter.*

4.1.5 Model boundary and review note

This lesson uses a toy cash-and-carry model for an investment asset with known income and deterministic discounting. It does not model futures-style daily settlement, repo specialness, collateral, counterparty credit, taxes, borrowing asymmetry, delivery options, or an asset that cannot be financed or shorted on the assumed terms.

The formulas are implemented in a pure tested domain module. Sources, notation, examples, code, and answer keys remain `draft` pending independent human review.

References

[1] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 1 §1.3, printed pp. 5-7, forward contracts and payoffs.

[2] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 5 §§5.3-5.5, printed pp. 103-108, cash-and-carry forward pricing and known income.

[3] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 5 §5.7, printed pp. 109-111, valuation of existing forward contracts.

Answers to the knowledge checks

Knowledge check 4.1.1: Forward contracts and value

1. **(a)** Buy the underlying for the fixed delivery price. A forward is a bilateral obligation. At delivery the long pays the fixed delivery price and receives the underlying under the stated physical-settlement convention.

2. **(a)** USD -6. The long payoff is underlying value minus delivery price: 74 - 80 = USD -6. Unlike an option holder, the forward long cannot simply decline delivery in this model.

3. 105.2631578947 (accepted within ±0.000001). The prepaid forward is USD 100. Divide by the delivery-date discount factor: 100 / 0.95 = USD 105.2631578947.

4. 895.8333333333 (accepted within ±0.000001). Subtract income received before delivery, then carry the prepaid value to delivery: (900 - 40) / 0.96 = USD 895.8333333333.

5. 4.85 (accepted within ±0.000001). Discount the delivery-date advantage: 0.97 × (110 - 105) = USD 4.85.

6. −37.6 (accepted within ±0.000001). The signed value is 10 × 0.94 × (98 - 102) = USD -37.60.

Chapter 4.2

European option contracts and payoffs

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/derivatives/european-option-contra

4.2.1 What you will be able to do

After this lesson, you should be able to:

- identify the holder, writer, underlying, strike, and expiry of a European option;
- calculate long and short call and put payoffs;
- distinguish current value, premium, expiry payoff, and profit.

4.2.2 Rights and obligations

A European call gives its holder the right, but not the obligation, to buy one unit of the underlying for the option-strike-price K at the option-expiry-time T_{opt}. A European put gives the holder the corresponding right to sell. If the holder exercises, the writer must carry out the other side of the exchange. Hull introduces call and put rights and European exercise [1].

4.2.3 Expiry payoff

Let the derivative-underlying-value at expiry be $S_{T_{\text{opt}}}$. The local long-call-expiry-payoff and long-put-expiry-payoff, per unit, are

$$\Pi^{\text{call}}_{T_{\text{opt}}} = \max\Big(S_{T_{\text{opt}}} - K, 0\Big) \tag{4.2.1}$$

$$\Pi^{\text{put}}_{T_{\text{opt}}} = \max\Big(K - S_{T_{\text{opt}}}, 0\Big) \tag{4.2.2}$$

The maximum with zero expresses the holder's right not to exercise when exercise would give a negative payoff. A writer's signed payoff is the negative of the matching holder payoff.

For $S_{T_{\text{opt}}} = 115$ and $K = 100$, one call pays 15 and one put pays zero. For $S_{T_{\text{opt}}} = 70$ and $K = 85$, one put pays 15. A writer of three such puts has a signed payoff of -45.

4.2.4 Value, premium, payoff, and profit are different

The European call value c_t or European put value p_t is a current value. At inception, the holder pays the local option-premium-paid π_0, under the stated settlement convention. The payoff is determined later, by the state at expiry.

If financing is ignored, the local option-holder-expiry-profit of a call is the payoff minus the premium:

$$G_{T_{\text{opt}}} = \Pi^{\text{call}}_{T_{\text{opt}}} - \pi_0 \tag{4.2.3}$$

For example, a call bought for 6 that later pays 9 has a payoff of 9 and a profit of 3, if financing is ignored. A profit calculation that includes financing first moves the premium and all other cash flows to one common date. Hull's option examples distinguish premium, exercise, payoff, and profit [2].

Knowledge check 4.2.1: European option contracts and payoffs

1. What right does the holder of a European call have in this lesson?
   (a) The right to buy the underlying at the strike at expiry
   (b) The right to sell the underlying at any time before expiry
   (c) An obligation to buy the underlying at expiry
2. Who has the exercise right and who must perform if a European put is exercised?
   (a) The holder chooses; the writer must perform
   (b) The writer chooses; the holder must perform
   (c) Both sides independently choose whether to perform
3. Two European calls each have strike USD 100. The underlying is worth USD 115 at expiry. What is the total holder payoff in USD?
4. A trader writes three European puts with strike USD 85. The underlying is worth USD 70 at expiry. What is the writer's signed payoff in USD?
5. A call is purchased today for a USD 6 premium and later pays USD 9 at expiry. Ignoring financing, which quantities are correctly labeled?
   (a) Expiry payoff USD 9; profit USD 3
   (b) Expiry payoff USD 3; profit USD 9
   (c) Both payoff and profit are USD 9
6. An out-of-the-money option expires with zero payoff. Does that imply its value was zero on every earlier date?
   (a) No, earlier value reflected the remaining possible future payoffs
   (b) Yes, a zero realized payoff makes every earlier value zero
   (c) Yes, option value always equals expiry intrinsic value

Answers: end of the chapter.

4.2.5 Model boundary and review note

This lesson covers European calls and puts only. It does not value them or cover American or Bermudan exercise, automatic exercise, assignment, settlement delay, corporate actions, collateral, counterparty default, or barrier and other path-dependent terms.

The payoff functions are implemented in a pure tested domain module. Sources, notation, examples, code, and answer keys remain `draft` pending independent human review.

References

[1] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 1 §1.5, printed pp. 7-9, call and put option rights and European exercise.

[2] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 9 practice questions 9.1-9.2 and 9.9-9.10, printed pp. 211-212, payoff versus profit.

Answers to the knowledge checks

Knowledge check 4.2.1: European option contracts and payoffs

1. **(a)** The right to buy the underlying at the strike at expiry. A European call holder has a right, not an obligation, to buy at the strike at the single expiry time.

2. **(a)** The holder chooses; the writer must perform. The option premium buys the holder the exercise right; the writer bears the corresponding obligation under the stated contract.

3. 30 (accepted within ± 0.000001). Each call pays max(115 - 100, 0) = USD 15. Two calls pay USD 30.

4. -45 (accepted within ± 0.000001). Each holder receives max(85 - 70, 0) = USD 15, so the writer pays 3×15 = USD 45 and has payoff USD -45.

5. **(a)** Expiry payoff USD 9; profit USD 3. The contract payoff is USD 9. Ignoring financing, profit subtracts the USD 6 premium and is USD 3.

6. **(a)** No, earlier value reflected the remaining possible future payoffs. Before expiry, more than one future state may remain possible. A realized zero payoff does not retrospectively erase the value of the contingent right.

Chapter 4.3

Put-call parity and synthetic forwards

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/derivatives/put-call-parity/

4.3.1 What you will be able to do

After this lesson, you should be able to:

- show that long call plus short put has a forward-style expiry payoff;
- state the exact contract terms that must match;
- solve a missing European call or put value using discounted strike and known income.

4.3.2 Match the expiry cash flows

For the same derivative-underlying-value, option-strike-price K, and option-expiry-time T_{opt}, the holder payoffs satisfy this identity in every state at expiry:

$$\max\Big(S_{T_{\text{opt}}} - K, 0\Big) - \max\Big(K - S_{T_{\text{opt}}}, 0\Big) = S_{T_{\text{opt}}} - K \tag{4.3.1}$$

The left side is long call plus short put. The right side is the payoff of a long forward-style exchange at strike K. This is an identity in every expiry state, not a forecast or an average.

4.3.3 Present-value parity

Let the underlying-income-present-value be I_0. The local prepaid-forward-value is $\widetilde{F}_{0,T_{\text{opt}}} = S_0 - I_0$. Applying the law of one price to the matched portfolios, with European call value c_0, European put value p_0, and the strike-payment discount-factor $D(0, T_{\text{opt}})$, gives put-call parity in present-value form:

$$c_0 - p_0 = \widetilde{F}_{0,T_{\text{opt}}} - KD(0, T_{\text{opt}}) \tag{4.3.2}$$

The prepaid forward value is the discounted forward-price $F_{0,T_{\text{opt}}}$ for delivery at option expiry:

$$\widetilde{F}_{0,T_{\text{opt}}} = D(0, T_{\text{opt}})F_{0,T_{\text{opt}}} \tag{4.3.3}$$

Here $F_{0,T_{\text{opt}}}$ is the current forward price, not the fixed delivery price of an existing forward. Substituting it into the present-value form gives put-call parity in forward form:

$$c_0 - p_0 = D(0, T_{\text{opt}}) \left(F_{0,T_{\text{opt}}} - K \right) \tag{4.3.4}$$

Hull derives put-call parity by comparing the expiry values of two portfolios and invoking no arbitrage[1].

For $c_0 = 8$, $\widetilde{F}_{0,T_{\text{opt}}} = 97$, $K = 100$, and $D(0, T_{\text{opt}}) = 0.95$, the present-value form gives the put value:

$$p_0 = 8 - 97 + 100(0.95) = 6 \tag{4.3.5}$$

4.3.4 When put-call parity does not apply

The equality needs the same underlying, strike, expiry, exercise style, quantity, income treatment, and settlement basis. American exercise adds a choice of exercise time. A knockout triggered by default changes the cash flows in the default state. If both options terminate on issuer default, the call payoff minus the put payoff is zero in that state. So the matching forward-style claim must also terminate on issuer default. The formula for options without a knockout does not apply unchanged.

Knowledge check 4.3.1: Put-call parity

1. Which matching conditions are required for the European call and put in the lesson's put-call parity?
   (a) Same underlying, strike, expiry, exercise style, and settlement basis
   (b) Only the same current premium
   (c) Only the same call-or-put direction
2. Does vanilla European put-call parity automatically apply to an option that is extinguished by issuer default?
   (a) No, the default state changes the matched cash flows and needs an explicit survival-contingent parity
   (b) Yes, the same formula applies to every contract called an option
   (c) Yes, because default can be ignored after a strike is fixed
3. A European call is worth USD 8, the prepaid forward value is USD 97, the strike is USD 100, and D(0,T) = 0.95. What put value satisfies parity, in USD?
4. A European put is worth USD 7, the prepaid forward value after known income is USD 92, the strike is USD 100, and D(0,T) = 0.96. What call value satisfies parity, in USD?

Answers: end of the chapter.

4.3.5 Model boundary and review note

Put-call parity does not depend on an option pricing model. It still relies on the stated assumptions about trading, funding, income, and matching contract terms. Put-call parity does not give an

option price by itself, and it contains no volatility model.

The parity functions are implemented in a pure tested domain module. Sources, notation, examples, code, and answer keys remain `draft` pending independent human review.

References

[1] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 10 §10.4, printed pp. 221-225, European put-call parity and arbitrage portfolios.

Answers to the knowledge checks

Knowledge check 4.3.1: Put-call parity

1. **(a)** Same underlying, strike, expiry, exercise style, and settlement basis. Parity compares portfolios whose state-by-state expiry cash flows match. Different underlying, strike, expiry, style, or settlement basis breaks that match.

2. **(a)** No, the default state changes the matched cash flows and needs an explicit survival-contingent parity. If both options terminate on default, call-minus-put is also zero in that state. The corresponding synthetic forward must carry the same default contingency.

3. 6 (accepted within ± 0.000001). From c - p = prepaid forward - K D(0,T), p = 8 - 97 + 100 $\times$ 0.95 = USD 6.

4. 3 (accepted within ± 0.000001). Rearrange parity: c = p + prepaid forward - K D(0,T) = 7 + 92 - 96 = USD 3.

Chapter 4.4

One-period binomial option valuation

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/derivatives/one-period-binomial-op

4.4.1 What you will be able to do

After this lesson, you should be able to:

- construct the two-position portfolio that matches both state payoffs;
- calculate the tree's risk-neutral up-state weight;
- show that replication and discounted risk-neutral expectation give the same value.

4.4.2 Two states, two traded positions

Let the underlying be worth S_0 at valuation-time and take exactly the local values binomial-up-underlying-value S^{up} or binomial-down-underlying-value S^{down} at the option-expiry-time T_{opt}. A claim pays local binomial-up-claim-payoff X^{up} or binomial-down-claim-payoff X^{down}.

The local binomial-hedge-units Δ is the number of units of the underlying that matches the difference between the two claim payoffs:

$$\Delta = \frac{X^{\text{up}} - X^{\text{down}}}{S^{\text{up}} - S^{\text{down}}} \tag{4.4.1}$$

The local binomial-risk-free-expiry-cash is then chosen so that the portfolio matches the claim payoff in the down state:

$$M_{T_{\text{opt}}} = X^{\text{down}} - \Delta\, S^{\text{down}} \tag{4.4.2}$$

So the local binomial-claim-value is the value of the replicating portfolio at time zero:

$$V_0 = \Delta\, S_0 + D(0, T_{\text{opt}}) M_{T_{\text{opt}}} \tag{4.4.3}$$

4.4.3 The same value with risk-neutral weights

The local binomial-risk-neutral-up-weight q is chosen so that the model reproduces the current value of the underlying:

$$q = \frac{S_0/D(0,T_{\text{opt}}) - S^{\text{down}}}{S^{\text{up}} - S^{\text{down}}} \tag{4.4.4}$$

Under the risk-neutral-probability-measure $\mathbb{Q}$, the claim value is the discounted risk-neutral expected payoff:

$$V_0 = D(0,T_{\text{opt}}) \left(q\, X^{\text{up}} + (1-q)\, X^{\text{down}} \right) \tag{4.4.5}$$

Hull derives the one-step hedge and then shows that its no-arbitrage value equals the discounted risk-neutral expected payoff [1].

4.4.4 Worked call example

Let $S_0 = 100$, $S^{\text{up}} = 120$, $S^{\text{down}} = 80$, $D(0,T_{\text{opt}}) = 1/1.05$, and let the option-strike-price be $K = 100$. The call payoffs are $X^{\text{up}} = 20$ and $X^{\text{down}} = 0$.

Then $\Delta = 0.5$, $M_{T_{\text{opt}}} = -40$, and $q = 0.625$. Both methods give the same claim value:

$$V_0 = 0.5(100) + \frac{-40}{1.05} = \frac{0.625(20) + 0.375(0)}{1.05} = 11.9047619 \tag{4.4.6}$$

The negative cash position is borrowing. It is not a negative option value.

4.4.5 From one period to many periods

Repeating this valuation one period at a time, from expiry backward, is backward induction on a multi-period tree. The later bond-option lessons need a discount factor at each node and an additional state for survival and default. An option with a knockout triggered by default cannot be valued from the distribution at expiry alone, because the option can terminate before expiry.

Knowledge check 4.4.1: One-period binomial option valuation

1. A no-income underlying is worth USD 100 now and will be USD 120 or USD 80 at expiry. D(0,T) = 1/1.05. A call with strike USD 100 pays USD 20 or USD 0. The model-implied Q up-weight is 0.625. What is the call value in USD?

2. In the same tree, a put with strike USD 100 pays USD 0 in the up state and USD 20 in the down state. The Q up-weight is 0.625 and D(0,T) = 1/1.05. What is the put value in USD?

3. A claim pays USD 20 when the underlying is USD 120 and USD 0 when it is USD 80. How many underlying units are required in its one-period replicating portfolio?

4. A claim pays USD 10 when the underlying is USD 130 and USD -2 when it is USD 90. Its hedge quantity is 0.3 underlying units. What risk-free cash amount is required at expiry to reproduce the down-state payoff, in USD?

Answers: *end of the chapter.*

4.4.6 Model boundary and review note

This complete two-state model has one non-income-paying traded underlying, one risk-free account, and deterministic discounting. It does not estimate real-world probabilities or volatility. It does not model multiple time steps, stochastic rates, default, or transaction and funding effects.

The two valuation methods are implemented together in a pure tested domain module. The hand-calculated answer and implementation both remain subject to independent human quantitative review.

References

[1] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 12 §§12.1-12.2, printed pp. 253-259, one-step binomial replication and risk-neutral valuation.

Answers to the knowledge checks

Knowledge check 4.4.1: One-period binomial option valuation

1. 11.9047619048 (accepted within ±0.000001). Discount the Q-weighted payoff: $(1/1.05) \times [0.625 \times 20 + 0.375 \times 0]$ = USD 11.9047619048.

2. 7.1428571429 (accepted within ±0.000001). The Q down-weight is 0.375. Discounting gives $(1/1.05) \times [0.625 \times 0 + 0.375 \times 20]$ = USD 7.1428571429.

3. 0.5 (accepted within ±0.000001). The hedge quantity is the payoff difference divided by the underlying-value difference: (20 - 0) / (120 - 80) = 0.5.

4. −29 (accepted within ±0.000001). Cash at expiry is $M_{T_{\text{opt}}} = X^{\text{down}} - \Delta\, S^{\text{down}} = -2 - 0.3 \times 90 =$ USD −29. A negative amount is borrowing.

Chapter 4.5

Multi-period lattice valuation

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/derivatives/multiperiod-lattice-va

4.5.1 Repeat the one-period argument backward

The one-period binomial lesson valued two terminal payoffs. A multi-period lattice applies the same pricing step at every node, beginning with the last row and moving backward to valuation-time.

At the node with lattice-time-index i and lattice-state-index j, let the lattice-node-up-weight be $q_{i,j}$ and the lattice-node-discount-factor be $d_{i,j}$. Let the local next-date cash flow $CF_{i+1,j}$ be paid at the successor node $(i+1, j)$; the node value $V_{i+1,j}$ excludes it. The lattice-node-value at node (i, j) is the discounted risk-neutral expectation of the successor cash flows plus the successor node values:

$$V_{i,j} = d_{i,j}\left[\left(1 - q_{i,j}\right)\left(CF_{i+1,j} + V_{i+1,j}\right) + q_{i,j}\left(CF_{i+1,j+1} + V_{i+1,j+1}\right)\right] \tag{4.5.1}$$

The two successors of node (i, j) are the nodes $(i+1, j)$ and $(i+1, j+1)$. The labels "up" and "down" are conventions. The input up weights and values must use the same order.

4.5.2 Timing is part of the model

In (4.5.1), each next-date cash flow is added to the value of its successor node before discounting. So the value of a coupon bond at a node is ex-coupon: it excludes the coupon paid at that node. A change of the order of exercise and coupon payment changes the contract, and it can change the option payoff.

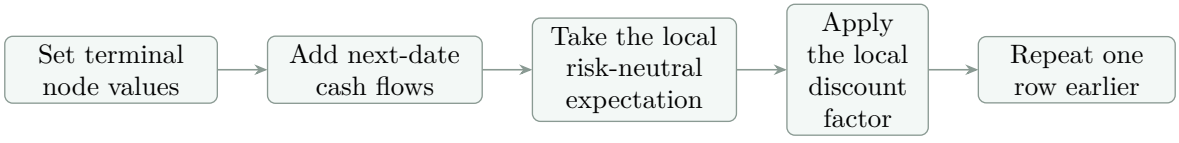


Diagram 4.5.1 **Backward-induction order.** Start from terminal node values. At each earlier row, add any next-date cash flow to each successor value, form the local risk-neutral weighted average, multiply by the local discount factor, and repeat.

4.5.3 Worked example

The terminal values are USD 0, USD 20, and USD 40. At time index 1, both nodes have discount factor 0.90, and their up weights are 0.25 and 0.75. At time index 0, the discount factor is 0.95 and the up weight is 0.50. Applying (4.5.1) row by row gives the node values:

$$\begin{aligned} V_{1,0} &= 0.90\left[(1-0.25)\,(0)+0.25\,(20)\right] = 4.50 \\ V_{1,1} &= 0.90\left[(1-0.75)\,(20)+0.75\,(40)\right] = 31.50 \\ V_{0,0} &= 0.95\left[(1-0.50)\,(4.50)+0.50\,(31.50)\right] \\ &= 17.10 \text{ USD} \end{aligned} \tag{4.5.2}$$

`backwardInductionValue` reproduces the three node rows `[[17.1], [4.5, 31.5], [0, 20, 40]]` in its independent reference test.

Knowledge check 4.5.1: Multi-period lattice

1. At a node, next-period claim values are 10 in the down state and 30 in the up state. The risk-neutral up weight is 0.4 and the one-period discount factor is 0.95. Calculate the current node value.

2. At a bond-tree node, next-period ex-coupon values are 90 and 110, both states also pay a coupon of 5, the risk-neutral up weight is 0.25, and the node discount factor is 0.96. Calculate the current node value.

Answers: end of the chapter.

4.5.4 Sources

The replication, risk-neutral weighting, and backward-induction construction follow Hull, Chapter 12 §§12.1–12.3 [1]. The node discount factors are model inputs. Calibration and short-rate dynamics are outside this lesson.

References

[1] Hull, Options, Futures, and Other Derivatives (8th ed., 2012).

Answers to the knowledge checks

Knowledge check 4.5.1: Multi-period lattice

1. 17.1 (accepted within ± 0.000001). The node value is $0.95 \times (0.6 \times 10 + 0.4 \times 30) = 17.1$.

2. 96 (accepted within ± 0.000001). Add the coupon before averaging: $0.96 \times [0.75 \times 95 + 0.25 \times 115] = 96$.

Part 5

Bond options

Part 2 put a bond's quoted and cash prices on one basis, and Part 4 valued a European option by backward induction on a lattice. This part combines them: the underlying of the option is a bond, and the lattice gives the bond values at expiry. The second chapter also uses the survival probability and the recovery of Part 6, because issuer default ends the option.

This part values European options whose underlying is a bond. The first chapter defines a standalone call or put on a bond, puts the strike and the bond value on one cash-price basis, and values the expiry payoffs on an input lattice. The second chapter adds issuer default. Default before exercise extinguishes the option with no rebate, the bond's recovery stays a separate cash flow of the bond, and each backward step discounts only the survival-weighted option values.

The argument

1. A bond option's strike and bond value share one cash-price basis
2. The lattice gives the bond value at each expiry node
3. Backward induction values the expiry payoffs
4. Issuer default before exercise extinguishes the option
5. Each backward step weights the option value by survival

After this part you can

- Interpret a standalone European call or put on a named bond, including exercise time, dirty-versus-clean strike basis, coupon timing, and holder rights.
- Calculate a European bond option payoff at expiry and value it by backward induction on an input default-free bond-value lattice.
- Interpret an issuer-default knockout bond option as a contract extinguished with zero option rebate if the bond issuer defaults before exercise, separately from any recovery on the bond.
- Calculate an issuer-default knockout bond option value on an input joint rate-credit lattice by applying survival to option continuation and zero option payoff on pre-expiry default.

Builds on

- Bonds
- Derivative foundations
- Credit risk

Chapter 5.1

European bond options

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/bond-options/european-bond-options

5.1.1 The standalone contract

A European bond call gives the holder the right, but not the obligation, to buy a named bond at option-expiry-time. A put gives the right to sell. This option is not the issuer's call feature inside a callable bond. The two rights are held by different parties and belong to different instruments.

5.1.2 Put the bond value and the strike on one price basis

The local bond-option-expiry-value $B^{\text{dirty}}_{T_{\text{opt}}}$ and option-strike-price K must both be dirty cash prices. If a contract states a clean strike, add the accrued-interest at the exercise date to the strike before the comparison.

In this lesson, exercise occurs immediately after any coupon paid on the expiry date. The call payoff and the put payoff are:

$$X^{B,\text{call}}_{T_{\text{opt}}} = \max\left(B^{\text{dirty}}_{T_{\text{opt}}} - K, 0\right) \tag{5.1.1}$$

$$X^{B,\text{put}}_{T_{\text{opt}}} = \max\left(K - B^{\text{dirty}}_{T_{\text{opt}}}, 0\right) \tag{5.1.2}$$

The lattice first gives the bond values at expiry. Apply (5.1.1) or (5.1.2) at the expiry nodes. Then apply backward induction to calculate the lattice-node-value at time zero.

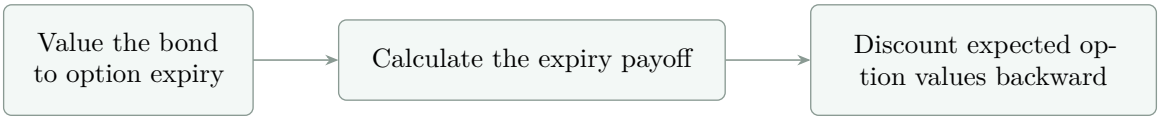


Diagram 5.1.1 **European bond-option valuation flow.** First value the bond at every option-expiry node, then apply the call or put payoff, then discount the risk-neutral expected payoffs backward through the input lattice.

5.1.3 Worked example

Suppose a one-period input lattice gives dirty bond values at expiry of USD 90 and USD 110. The strike is USD 100, the risk-neutral up weight is 0.60, and the discount factor is 0.95. Applying (5.1.1) and one backward step gives the call value:

$$\begin{aligned} X_{T_{\text{opt}}}^{B,\text{call,down}} &= \max(90-100,0) = 0 \\ X_{T_{\text{opt}}}^{B,\text{call,up}} &= \max(110-100,0) = 10 \\ V_0 &= 0.95\left[(1-0.60)\,(0)+0.60\,(10)\right] \\ &= 5.70 \text{ USD per option} \end{aligned} \tag{5.1.3}$$

The tested domain code does the same two steps: `europeanOptionPayoff` calculates the expiry payoffs, and `backwardInductionValue` discounts them.

5.1.4 Coupon and exercise ordering

In this lesson, a coupon on the expiry date is paid before exercise. So the bond value at expiry is ex-coupon. A contract that is exercised before the coupon, or that transfers the right to the coupon, requires different terminal inputs.

Knowledge check 5.1.1: European bond option

1. A standalone European call on a bond gives its holder which right?
   (a) The right, but not obligation, to buy the named bond at the stated strike and exercise time
   (b) The issuer's right to redeem every outstanding bond
   (c) Automatic ownership of every coupon paid before option exercise
2. An option strike is quoted as a clean bond price and exercise occurs between coupon dates. What is needed before comparing it with a dirty bond value?
   (a) Add the exercise-date accrued interest under the stated convention
   (b) Subtract every coupon paid over the bond's entire life
   (c) No adjustment; clean and dirty prices are identical between coupon dates
3. At expiry a bond has dirty ex-coupon values of USD 78, USD 84, and USD 89 across three nodes. Calculate call payoffs for strike USD 82 and report their sum in USD.
4. One period before expiry, a default-free bond call is worth USD 0 in the down node and USD 6 in the up node. The risk-neutral up weight is 0.4 and the node discount factor is 0.95. Calculate the current option value in USD.

Answers: *end of the chapter.*

5.1.5 Sources

The bond-option definition and the need to distinguish a cash (dirty) strike from a quoted (clean) strike follow Hull, Chapter 28 §28.1 [1]. This lesson uses an input lattice, not Black's market model, so that the next lesson can model the issuer-default state explicitly.

References

[1] Hull, Options, Futures, and Other Derivatives (8th ed., 2012).

Answers to the knowledge checks

Knowledge check 5.1.1: European bond option

1. **(a)** The right, but not obligation, to buy the named bond at the stated strike and exercise time. This lesson models a standalone holder option, not an embedded issuer call feature.

2. **(a)** Add the exercise-date accrued interest under the stated convention. Payoff inputs must share one price basis. A clean strike becomes a dirty cash strike after adding applicable accrued interest.

3. 9 (accepted within ± 0.000001). The node payoffs are 0, 2, and 7, whose sum is USD 9. The sum is requested only as a compact payoff check, not as a probability-weighted value.

4. 2.28 (accepted within ± 0.000001). Backward induction gives $0.95 \times (0.6 \times 0 + 0.4 \times 6) =$ USD 2.28.

Chapter 5.2

Issuer-default knockout bond options

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/bond-options/issuer-default-knocko

5.2.1 Contract definition

In this lesson, "knockout" has one trigger: if the issuer defaults before the exercise time, the option is extinguished and pays no rebate. This knockout is not a barrier on the market price of the bond.

The underlying bond is a different claim. The bond can still pay recovery after default. So the recovery is part of the bond value, and it is never part of the option value in the default state.

5.2.2 Value the underlying bond first

An alive node is a node at which the issuer has not defaulted. At an alive node with lattice-time-index i and lattice-state-index j, use the local node discount factor $d_{i,j}$, the local risk-neutral up weight $q_{i,j}$, the conditional-node-survival-probability $s_{i,j}$, the next surviving cash flow CF^B_{i+1}, and the local bond recovery $R^B_{i,j}$. The local surviving-bond-continuation-value is the bond value at the next time, conditional on survival. It is the next surviving cash flow plus the expected alive bond value:

$$H^B_{i,j} = CF^B_{i+1} + (1 - q_{i,j})B^{\text{alive}}_{i+1,j} + q_{i,j}B^{\text{alive}}_{i+1,j+1} \tag{5.2.1}$$

The Alive-state bond value is the discounted expectation over survival and default. Survival gives the continuation value from (5.2.1), and default gives the recovery:

$$B^{\text{alive}}_{i,j} = d_{i,j}\left[s_{i,j}H^B_{i,j} + \left(1 - s_{i,j}\right)R^B_{i,j}\right] \tag{5.2.2}$$

The values at each node are ex-cash-flow. The model assumes that recovery is paid at the next lattice date. This timing is a simplification.

5.2.3 Apply the knockout to option continuation

At call expiry, if the issuer has not defaulted, the option value is the call payoff:

$$O^{\text{KO}}_{T,j} = \max\left(B^{\text{alive}}_{T,j} - K, 0\right) \tag{5.2.3}$$

Before expiry, the option value in the default state is zero. So the option value at a node is the discounted, survival-weighted expectation of the option values at the two successor nodes:

$$O^{\mathrm{KO}}_{i,j} = d_{i,j}s_{i,j}\left[\left(1-q_{i,j}\right)O^{\mathrm{KO}}_{i+1,j} + q_{i,j}O^{\mathrm{KO}}_{i+1,j+1}\right] \tag{5.2.4}$$

(5.2.4) has no recovery term, unlike (5.2.2). If the conditional survival probability at the current node is zero, the option value is zero, also when the bond recovery is positive.

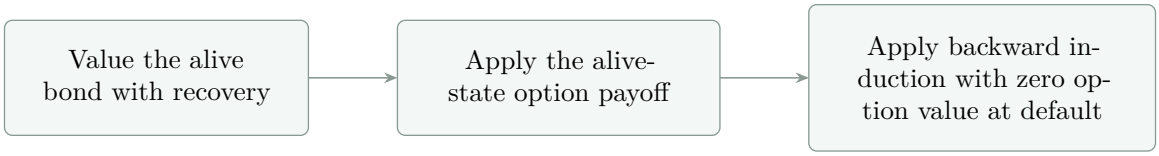


Diagram 5.2.1 **Issuer-default knockout valuation flow.** Value the underlying bond first, including its separate input recovery. At option expiry, apply the payoff only in alive states. Before expiry, discount only the survival-weighted option values, because the option value at default is zero.

5.2.4 Worked example

Consider a lattice with two annual steps. At option expiry, the input values of the final period give alive bond values of USD 82.80 and USD 88.65. With a call strike of USD 85, (5.2.3) gives terminal option values of USD 0 and USD 3.65. At valuation time, the discount factor is 0.95, the conditional survival probability is 0.90, and the up weight is 0.50. The option values are:

$$\begin{aligned} O^{\mathrm{KO}}_{1,0} &= \max(82.80-85,0) = 0 \\ O^{\mathrm{KO}}_{1,1} &= \max(88.65-85,0) = 3.65 \\ O^{\mathrm{KO}}_{0,0} &= 0.95(0.90)\left[(1-0.50)\,(0)+0.50\,(3.65)\right] \\ &= 1.560375 \text{ USD per option} \end{aligned} \tag{5.2.5}$$

The domain function `valueDefaultKnockoutBondOption` performs the same bond recursion and option recursion, in the same order.

5.2.5 Put-call parity with a default knockout

Put-call parity for options without a knockout does not hold unchanged for this contract. In a default state before expiry, the call value minus the put value is zero. So the matching synthetic forward must also be extinguished by the same default event. A default-free forward does not match the options in every state.

5.2.6 Model and source boundary

Hull describes European bond options. Separately, Hull describes CDS forwards and CDS options that cease to exist if the reference entity defaults. This lesson combines the two descriptions for teaching: an option on the issuer's bond that is knocked out by issuer default. The lesson implements the stated contract. It does not claim that all bond options use this convention. The joint lattice for rates and credit is an input; it is not calibrated.

Knowledge check 5.2.1: Issuer-default knockout bond option

1. In this lesson's contract, what happens if the issuer defaults before option exercise?
   (a) The option is extinguished with zero option rebate
   (b) The option holder automatically receives the bond's recovery amount
   (c) The option survives unless the bond price also crosses a barrier
2. The bond pays USD 40 recovery after default, but an otherwise identical call is extinguished on that default. What enters the option's default branch?
   (a) USD 0; the USD 40 belongs only in the bond valuation
   (b) USD 40 because every bond cash flow passes through to the option
   (c) The option strike because exercise is forced
3. At an alive node, the risk-neutral expected next option value conditional on survival is USD 10, conditional survival over the period is 0.90, and the node discount factor is 0.95. Default gives zero option rebate. Calculate the current option value in USD.
4. At an alive node, next option values conditional on survival are USD 2 down and USD 8 up. The risk-neutral up weight is 0.25, conditional survival is 0.80, and the discount factor is 0.96. Calculate the current knockout option value in USD.

Answers: end of the chapter.

5.2.7 Sources

The bond-option cash-price conventions follow Hull, Chapter 28 §28.1. Hull describes the analogous extinguishment at default for CDS forwards and CDS options in Chapter 24 §24.5. Backward induction follows Chapter 12 [1]. A recovery convention must state its base and its timing, as the fixed-income treatment emphasizes [2].

References

[1] Hull, Options, Futures, and Other Derivatives (8th ed., 2012).

[2] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022).

Answers to the knowledge checks

Knowledge check 5.2.1: Issuer-default knockout bond option

1. **(a)** The option is extinguished with zero option rebate. Issuer default is the knockout event. Recovery on the underlying bond is modeled separately and is not an option rebate.
2. **(a)** USD 0; the USD 40 belongs only in the bond valuation. The contract separates underlying recovery from option termination.
3. 8.55 (accepted within ± 0.000001). The value is $0.95 \times [0.90 \times 10 + 0.10 \times 0] =$ USD 8.55.

4. 2.688 (accepted within ± 0.000001). The survival-state expectation is $0.75 \times 2 + 0.25 \times 8 = 3.5$. Multiply by survival and discounting: $0.96 \times 0.80 \times 3.5 =$ USD 2.688.

Part 6

Credit risk

Part 1 gave events and expectations over a partition, present value, and risk-neutral pricing weights. This part applies them to a new event, the default of a reference entity. Default and survival partition the outcomes at each future date, so a survival curve gives default probabilities, and a risky payment is valued as a discounted expectation over the survival and default states.

This part introduces the default time of a reference entity in the course's toy models. The first chapter reads an input survival curve, calculates the default probability of each interval, and uses a constant hazard rate. The second chapter values one claim under recovery of par: it pays its full amount on survival and a recovered fraction of par on default, both at maturity, and the claim's value is the discounted expectation of those two payments.

The argument

1. The default time is random, and survival to a date is an event
2. A survival curve gives the default probability of each interval
3. A constant hazard rate gives an exponential survival curve
4. Recovery of par sets the payment in the default state
5. A risky payment is the discounted expectation over survival and default

After this part you can

- Interpret survival probability as the modeled probability, conditional on survival at valuation time, of no default through a stated future time.
- Calculate time-zero cumulative or interval default probability from survival probabilities at stated times.
- Interpret a hazard rate as an instantaneous conditional default intensity, measured per unit of time, among names that have survived to that time.
- Calculate survival probability as S(0,t) = exp(-lambda t) under a stated non-negative constant hazard rate and consistent time units.
- Interpret the recovery rate as the fraction of par paid at maturity after default in the lesson's one-period recovery-of-par model.
- Calculate loss given default as one minus the recovery rate in the lesson's fractional recovery-of-par model.
- Calculate the time-zero present value of a one-period recovery-of-par claim by probability-weighting its survival and default payoffs at maturity and applying the stated discount factor.

Builds on

- Foundations

Leads to

- Bond options
- CDS

Chapter 6.1

Default, hazard, and survival

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/credit/default-hazard-and-survival

6.1.1 What you will be able to do

After this lesson, you should be able to:

- interpret a survival probability without confusing it with a discount factor;
- calculate the default probability of one interval from the survival curve;
- explain what a constant hazard rate means in this toy model;
- calculate the survival probability from a constant hazard rate.

6.1.2 Survival probability: no default by time t

The survival-probability $S^{\mathbb{Q}}(0,t)$ is the probability, under the risk-neutral-probability-measure $\mathbb{Q}$, that the reference entity has not defaulted between valuation-time 0 and a future time-from-valuation t. With the default-time τ, the survival probability is defined as:

$$S^{\mathbb{Q}}(0,t) = \mathbb{Q}(\tau > t) \tag{6.1.1}$$

In this lesson, the survival probabilities are inputs: the lesson states their values and does not calculate them. In practice, a survival curve is fitted to market prices. The CDS lesson on the credit curve fits one to CDS quotes.

A survival probability is not a discount factor. A survival probability is the probability of a modeled event. A discount factor converts an amount paid at a future time to its value at time zero. A survival probability and a discount factor can have equal values and still mean different things.

We assume that the survival curve starts at one:

$$S(0,0) = 1 \tag{6.1.2}$$

We also assume that the survival curve does not increase with time. This lesson does not estimate survival probabilities from historical data, and it does not fit them to market quotes. Tuckman and Serrat introduce survival and default probabilities for CDS valuation[1]. Hull distinguishes risk-neutral default probabilities, which are used in valuation, from real-world default probabilities [2].

6.1.3 Interval default probability: a decrease in survival

The local credit-period-index i selects one interval of the survival schedule, and the number-of-credit-periods n counts the intervals. Interval i starts at t_{i-1} and ends at the payment-time t_i, with $t_{i-1} < t_i$. The interval-default-probability of interval i is the decrease in survival probability over the interval:

$$\Delta q_i^{\mathbb{Q}} = S^{\mathbb{Q}}(0, t_{i-1}) - S^{\mathbb{Q}}(0, t_i) \tag{6.1.3}$$

The survival probability at the start of the interval comes first in the subtraction. Because the survival curve does not increase, this order gives $\Delta q_i \geq 0$. The sum of the interval default probabilities from time zero to t_n is the cumulative default probability by t_n:

$$\sum_{i=1}^{n} \Delta q_i = 1 - S(0, t_n) \tag{6.1.4}$$

This identity holds because the sum telescopes: each intermediate survival probability appears once with a plus sign and once with a minus sign. The identity uses two assumptions: $S(0,0) = 1$, and at most one modeled default.

6.1.4 Constant hazard rate: an intensity, not a cumulative probability

In this lesson, the Constant hazard rate λ is constant. It is a conditional default intensity, measured per model-year. The hazard rate is not the probability of default by the end of one year. It is also not a discount rate. With the local credit-small-time-increment Δt, the hazard rate is defined as the limit of a conditional default probability per unit of time:

$$\lambda = \lim_{\Delta t \downarrow 0} \frac{\mathbb{Q}(t < \tau \leq t + \Delta t \mid \tau > t)}{\Delta t} \tag{6.1.5}$$

The probability in this definition is conditional on the event $\tau > t$. So the hazard rate at time t describes only the paths on which the reference entity has survived to t.

With a constant hazard rate, the survival probability is:

$$S^{\mathbb{Q}}(0, t) = \exp(-\lambda t) \tag{6.1.6}$$

So the cumulative default probability by t is $1 - S(0, t)$. This lesson does not introduce piecewise hazard rates, curve calibration, or a statistical estimator. Tuckman and Serrat derive the conditional intensity and the constant-hazard survival probability[3], and so does Hull [4].

Example 6.1.1: Survival and default probabilities

One interval drop

Suppose the one-year survival probability is $S(0,1) = 0.97$. With $S(0,0) = 1$, the interval default probability of the first year is:

$$\Delta q_1 = 1 - 0.97 = 0.03 = 3\%$$

The 3% is the probability of default during the first year. It is not discounted.

Three years at constant hazard

Let $\lambda = 0.02$ per model-year and $t = 3$ model-years. Substituting these values into (6.1.6) gives the three-year survival probability:

$$S(0,3) = \exp(-0.02 \times 3) = \exp(-0.06) \approx 0.941765$$

The cumulative default probability by three years is:

$$1 - S(0,3) \approx 0.058235 = 5.8235\%$$

The cumulative default probability is not $\lambda t = 6\%$. The product λt is the exponent in the survival probability. It is not the exact cumulative default probability.

Two unequal interval drops

Consider this survival curve:

Interval index i	Time t_i, model-years	Survival $S(0,t_i)$	Interval default Δq_i
0	0	1.00	—
1	1	0.97	0.03
2	2	0.92	0.05

The interval default probability of the second interval, 0.05, is larger than that of the first interval, 0.03, although both intervals are one model-year long. The cumulative two-year default probability is $0.03 + 0.05 = 0.08 = 1 - 0.92$.

6.1.5 Check your understanding

The assessment separates survival meaning, interval subtraction, hazard interpretation, and constant-hazard arithmetic. Each competency has a direct item and a transfer item.

Knowledge check 6.1.1: Default, hazard, and survival

1. A model gives S(0,3) = 0.92 for a name that has survived to valuation time 0. Which interpretation is correct?
   (a) The modeled probability of no default through year 3 is 92%
   (b) The modeled probability of default by year 3 is 92%
   (c) The three-year discount factor is necessarily 0.92
2. Under the same model and valuation time, name A has five-year survival probability 0.84 and name B has five-year survival probability 0.93. Which conclusion is supported?
   (a) Name A has the larger modeled probability of default by year 5
   (b) Name B has the larger modeled probability of default by year 5
   (c) The two names have the same modeled probability of default by year 5
3. A model gives two-year survival probability S(0,2) = 0.93. Enter the cumulative probability of default by year 2 as a decimal.

4. Viewed from valuation time 0, survival probabilities are S(0,1) = 0.97 and S(0,2) = 0.91. Enter the unconditional probability of default during the interval from year 1 to year 2 as a decimal.

5. In the lesson model, what does a hazard rate at time t describe?

   (a) Instantaneous default intensity conditional on survival to time t

   (b) Cumulative default probability from time 0 through time t

   (c) The interest rate used to discount every cash flow

6. Two names have survived to the same instant. Name A's current hazard rate is 2% per year and name B's is 5% per year. Over the next very short interval, which statement matches the hazard interpretation?

   (a) Name B has the higher modeled conditional default likelihood

   (b) Name A has the higher modeled conditional default likelihood

   (c) The rates alone prove both names have already defaulted

7. Under the constant-hazard exponential model, calculate the dimensionless three-year survival probability for an annual hazard rate of 2% (0.02 per year).

8. Under the constant-hazard exponential model, calculate the dimensionless survival probability over 18 months for an annual hazard rate of 3.2% (0.032 per year).

Answers: *end of the chapter.*

6.1.6 Model boundary and review note

This lesson uses a toy model with at most one default, a constant hazard rate, exact model-year times, and survival probabilities that are inputs. The model omits:

- calendar dates, day counts, and business-day adjustments;
- stochastic interest rates;
- dependence between interest rates and default;
- estimation and calibration.

The lesson makes no claim that pricing-model probabilities are forecasts of realized default frequencies.

All materially changed source locators, formulas, examples, notation, 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. 14 §14.6, printed pp. 367–370, eqs. 14.4–14.7.

[2] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 23 §23.5, printed pp. 528–530.

[3] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.6, printed pp. 367–368; Appendix A14.1, printed p. 505, eqs. A14.1–A14.2.

[4] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 23 §§23.2 and 23.5, printed pp. 522–523 and 528–530.

Answers to the knowledge checks

Knowledge check 6.1.1: Default, hazard, and survival

1. **(a)** The modeled probability of no default through year 3 is 92%. S(0,3) is the model's probability, conditional on survival at valuation time, that no default occurs through year 3.

2. **(a)** Name A has the larger modeled probability of default by year 5. Lower survival to the same horizon implies higher cumulative default probability to that horizon: 1 - 0.84 exceeds 1 - 0.93.

3. 0.07 (accepted within $\pm 1e-12$). Cumulative default probability is 1 - S(0,2) = 1 - 0.93 = 0.07.

4. 0.06 (accepted within $\pm 1e-12$). The interval default probability is S(0,1) - S(0,2) = 0.97 - 0.91 = 0.06.

5. **(a)** Instantaneous default intensity conditional on survival to time t. Hazard is an instantaneous conditional intensity with inverse-time units, not a cumulative probability or discount rate.

6. **(a)** Name B has the higher modeled conditional default likelihood. Conditional on both names having survived to the instant, the higher current hazard corresponds to greater modeled default likelihood over a sufficiently short next interval.

7. 0.941764533584249 (accepted within ±0.000001). S(0,3) = exp(-0.02 × 3) = exp(-0.06) = 0.9417645336.

8. 0.953133787077505 (accepted within ±0.000001). Eighteen months is 1.5 years, so survival is exp(-0.032 × 1.5) = exp(-0.048) = 0.9531337871.

Chapter 6.2

Recovery and one-period risky present value

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/credit/recovery-and-risky-present-

6.2.1 What you will be able to do

After this lesson, you should be able to:

- interpret recovery as a fraction of a named base paid at a named time;
- calculate loss given default as the complement of recovery;
- calculate present value in one explicitly labeled recovery-of-par-paid-at-maturity model.

6.2.2 Name the recovery convention before calculating

This lesson uses one recovery convention. A claim has a positive local recovery-par-amount F and one local recovery-payment-time T, measured from valuation-time 0.

- If the reference entity survives through T, the holder receives F at T.
- If the reference entity defaults at or before T, the holder receives RF at T, where the recovery-rate R is deterministic.

The survival payment and the default payment are both paid at T. This timing is a model assumption, not a description of an actual instrument. This lesson does not teach recovery paid at default, recovery of market value, settlement delay, or a CDS settlement mechanism.

Tuckman and Serrat define a bond recovery rate relative to the face amount. They also identify the complementary loss fraction in their credit-loss discussion [1]. Payment of the recovery at maturity is a convention of this lesson. The lesson does not claim that it is the only recovery convention.

6.2.3 Recovery and loss given default are complements here

The loss-given-default fraction is defined as:

$$\text{LGD} = 1 - R, \qquad 0 \leq R \leq 1 \tag{6.2.1}$$

For example, $R = 0.40$ gives $\text{LGD} = 0.60$ in this toy model. The value 0.40 is chosen only for the arithmetic. The lesson makes no claim that 40% is a market convention, a forecast, or an appropriate assumption for any real entity.

6.2.4 Calculate the expected maturity amount, then discount it

The input survival-probability is $S^{\mathbb{Q}}(0,T)$ under the risk-neutral-probability-measure $\mathbb{Q}$. In the one-default model, the probability of default by T is $1 - S^{\mathbb{Q}}(0,T)$.

For the single interval from 0 through T, $1 - S(0,T)$ is the interval-default-probability. The expected amount paid at T under $\mathbb{Q}$ is the survival payment and the default payment, each multiplied by its probability:

$$F\left[S\left(0,T\right) + R\left(1 - S\left(0,T\right)\right)\right] \tag{6.2.2}$$

This is the two-event case of the expectation from disjoint events in the probability lesson: survival pays F, and default pays RF. Shreve gives the expectation of a finite-valued random variable[2]. Tuckman and Serrat weight credit cash flows by survival and default probabilities and then discount them[3].

The survival payment and the default payment are both paid at T. So one input discount-factor $D(0,T)$ applies to both. The recovery-of-par-present-value is the discounted expected amount:

$$PV_0^{\text{RoP}} = F\,D(0,T)\left[S^{\mathbb{Q}}\left(0,T\right) + R\left(1 - S^{\mathbb{Q}}\left(0,T\right)\right)\right] \tag{6.2.3}$$

This present-value is calculated from the perspective of the holder of the claim: the modeled future receipts and their value are positive. A purchase price paid by the holder would have the opposite cash-flow sign.

Example 6.2.1: Recovery-of-par present values

Partial recovery

Let $F =$ USD 1,000, $D(0,T) = 0.96$, $S(0,T) = 0.92$, and $R = 0.40$. The expected fraction of par paid at maturity is:

$$0.92 + 0.40(1 - 0.92) = 0.952$$

Multiplying this fraction by the par amount and the discount factor gives the present value:

$$PV_0^{\text{RoP}} = 1000(0.96)(0.952) = \text{USD } 913.92$$

The expected amount at maturity, before discounting, is USD 952.00. The discount factor converts this amount to its value at valuation time.

Zero-recovery boundary

Let $F =$ USD 1,000, $D(0,T) = 0.94$, $S(0,T) = 0.90$, and $R = 0$. The default payment is zero, so the present value is:

$$PV_0^{\text{RoP}} = 1000(0.94)\left[0.90 + 0\,(0.10)\right] = \text{USD } 846.00$$

With zero recovery, the present value is the discounted expected survival payment. This example makes no statement about real recovery outcomes.

Full-recovery boundary

Keep $F = \text{USD } 1{,}000$, $D(0,T) = 0.94$, and $S(0,T) = 0.90$, but set $R = 1$. The survival payment and the default payment are now both the full par amount, paid at T:

$$PV_0^{\text{RoP}} = 1000(0.94)\left[0.90 + 1\,(0.10)\right] = \text{USD } 940.00$$

With full recovery, the present value does not depend on the survival probability. The reason is that both modeled events pay the same amount on the same date. This result holds only under the timing convention of this lesson.

6.2.5 Check your understanding

The assessment tests the meaning of recovery, loss-given-default arithmetic, and direct and transfer calculations under the same one-period timing convention.

Knowledge check 6.2.1: Recovery and risky present value

1. In the lesson's one-period recovery-of-par model, what does a recovery rate of 40% mean if default occurs by maturity?

   (a) A recovery cash flow equal to 40% of par is paid at maturity
   (b) A recovery cash flow equal to 40% of the current market price is paid at maturity
   (c) A recovery cash flow equal to 60% of par is paid at maturity

2. A one-period claim has USD 2,000 par, a current price of USD 1,850, and a 25% recovery-of-par rate. Under the lesson model, which positive recovery cash flow is used after default?

   (a) USD 500 at maturity
   (b) USD 462.50 at maturity
   (c) USD 1,500 at maturity

3. Enter loss given default as a decimal when the recovery rate is 35% (0.35).

4. A recovery-of-par model pays USD 900 after default on a claim with USD 1,500 par. Enter loss given default as a decimal.

5. A one-period defaultable zero-coupon claim has USD 1,000 par, survival probability 0.90 to maturity, recovery-of-par rate 0.40 after default, and maturity discount factor 0.95. Both the survival payoff and any recovery are paid at maturity. What is its time-zero present value in USD?

6. A one-period defaultable zero-coupon claim has USD 2,500 par, cumulative default probability 0.08, loss given default 0.55, and maturity discount factor 0.92. The recovery-of-par amount after default is paid at maturity. What is its time-zero present value in USD?

Answers: end of the chapter.

6.2.6 Model boundary and review note

The model has one maturity date, one deterministic recovery rate, one input survival probability, and one input discount factor. Recovery is a fraction of par paid at maturity, also when default occurs earlier. The lesson omits:

- settlement at the default time;
- coupons and multiple intervals;
- stochastic recovery and recovery of market value;
- dependence between interest rates and default;
- calibration;
- liquidity, funding, taxes, and counterparty risk.

All materially changed content and source locators remain ‘draft‘ pending independent human review.

References

[1] Tuckman & Serrat, Fixed Income Securities: Tools for Today’s Markets (4th ed., 2022). Ch. 14 §14.2, printed p. 353.

[2] Shreve, Stochastic Calculus for Finance II: Continuous-Time Models (2004). Ch. 1 §1.3, printed pp. 13–18.

[3] Tuckman & Serrat, Fixed Income Securities: Tools for Today’s Markets (4th ed., 2022). Ch. 14 §14.7, printed pp. 371–372, eq. 14.8 and Table 14.11.

Answers to the knowledge checks

Knowledge check 6.2.1: Recovery and risky present value

1. **(a)** A recovery cash flow equal to 40% of par is paid at maturity. Under this narrow model, recovery uses par as its base and is paid at the common maturity date.
2. **(a)** USD 500 at maturity. Recovery is 0.25 × USD 2,000 par = USD 500; the current price is not the recovery base in this model.
3. 0.65 (accepted within $\pm 1e-12$). Loss given default is 1 - recovery rate = 1 - 0.35 = 0.65.
4. 0.4 (accepted within $\pm 1e-12$). The recovery rate is 900 / 1,500 = 0.60, so loss given default is 1 - 0.60 = 0.40.
5. 893 (accepted within ±0.01). The probability-weighted maturity payoff is 1,000 × [0.90 + (1 - 0.90) × 0.40] = 940. Discounting gives 0.95 × 940 = 893.
6. 2198.8 (accepted within ±0.01). The probability-weighted payoff factor is 1 - 0.08 × 0.55 = 0.956. Present value is 2,500 × 0.92 × 0.956 = 2,198.80.

Part 7

CDS

Part 6 gave survival probabilities from a hazard rate, recovery of par, and the present value of a payment that depends on default. A credit default swap exchanges payments that depend on the same default time, so each of its legs is a discounted expectation over the default intervals. The credit-curve chapter also uses the discount curve of Part 3.

This part values a credit default swap in a toy model and connects the value to the way the market quotes it. The first chapter builds the premium and protection legs from exact default-time expectations and solves the par spread. The second chapter converts a market-standard quote into a signed upfront amount against a fixed running coupon, and back. The third chapter separates what the market shows from what the model assumes or solves, and fits a credit curve to several tenors. The last chapter measures a position's sensitivity to the liquid contract's quote and sizes the equivalent notional that offsets it.

The argument

1. Each leg is a discounted expectation over the default intervals
2. The par spread makes the two legs equal with zero upfront
3. A market quote converts to an upfront amount against a fixed coupon
4. Several tenor marks fit a piecewise-constant hazard curve
5. Bump and reprice gives the quote risk, and the equivalent notional offsets it

After this part you can

- Interpret a CDS as running-premium payments exchanged for contingent protection, partition default into disjoint time events, and read the protection buyer's signed value as protection minus premium.
- Calculate the positive premium-leg present-value magnitude from scheduled survival-contingent premiums and exact accrued premium integrated over modeled default time under the stated constant-hazard model.
- Calculate the positive protection-leg present-value magnitude by integrating discounted loss-given-default over exact modeled default time under the stated constant risk-neutral hazard.
- Calculate the zero-upfront CDS running spread that equates the exact default-time premium- and protection-leg present-value magnitudes under one internally consistent model.
- Distinguish the lesson's market-standard quote from the fixed running coupon and explain the quote's role in a simplified CDS conversion.
- Calculate the signed upfront amount that makes protection-buyer inception value zero for an input quote-implied premium annuity and fixed running coupon.
- Calculate the market-standard quote implied by a signed upfront amount under the lesson's simplified flat-hazard conversion.
- Classify each quantity in a standard CDS valuation as market-observed, contractual, conventional, or model-solved, and explain why the traded upfront alone identifies the pricing-model hazard rate.

- Interpret a credit curve as the term structure of standard-contract marks at quoted tenors together with the survival curve fitted to them, and distinguish a tenor-by-tenor change of quotation from a fit of one hazard term structure.
- Calculate the survival probability implied by fitted piecewise-constant hazard segments and verify that a tenor's standard contract reprices to its mark under the fitted curve.
- Calculate a CDS position's signed value change for a one-basis-point rise in the liquid tenor's market-standard quote by re-solving the flat hazard rate and revaluing the position for its own side.
- Calculate the notional and side of the most liquid standard CDS that offsets a position's sensitivity to the liquid tenor's quote, and express it as an equivalent ratio to the position notional.

Builds on

- Foundations
- Rates and curves
- Credit risk

Chapter 7.1

CDS premium and protection legs

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/cds/premium-protection-legs-and-pa

7.1.1 What a CDS is

A credit default swap (CDS) is a bilateral derivative on a named reference entity. The protection buyer pays a running premium to the protection seller. If a covered credit event occurs before maturity, the seller makes a loss-related protection payment. Unlike buying the entity's bond, entering a CDS does not itself lend principal to the reference entity. Tuckman and Serrat describe the fee leg and the contingent leg (in this course, the premium leg and the protection leg) in Chapter 14[1]; Hull describes the same cash-flow directions for buyer and seller, and the quarterly premium convention[2].

This lesson reports both leg present values as positive magnitudes. From the protection buyer's perspective, the net present value is the protection-leg present value minus the premium-leg present value:

$$PV_0^{\text{buyer}} = PV_0^{\text{prot}} - PV_0^{\text{prem}} \tag{7.1.1}$$

The seller has the opposite signed value in this two-party model.

7.1.2 Start from disjoint events

Let the random default-time be τ. Use the local cds-period-index i across the number-of-cds-periods n. The schedule has dates $0 = t_0 < t_1 < \cdots < t_n = T$. Define the disjoint interval events

$$G_i = \{t_{i-1} < \tau \le t_i\}, \qquad i = 1, \dots, n \tag{7.1.2}$$

and add $G_0 = \{\tau > T\}$ for survival through maturity. These events are disjoint and exhaustive under the one-default assumption.

For a local cds-event-payoff Y that is constant at y_i on each event, the risk-neutral expectation is:

$$\mathbb{E}^{\mathbb{Q}}[Y] = \sum_{i=0}^{n} y_i \, \mathbb{Q}(G_i) \tag{7.1.3}$$

This is the expectation from disjoint events of the probability lesson, with risk-neutral pricing weights. If Y is not constant within an interval, the expectation uses the conditional means:

$$\mathbb{E}^{\mathbb{Q}}[Y] = \sum_{i=0}^{n} \mathbb{Q}(G_i)\, \mathbb{E}^{\mathbb{Q}}[Y \mid G_i] \tag{7.1.4}$$

For this reason, one payment at the end of the period is not an exact treatment of a default in the middle of the period. If every interval is divided into smaller default-time intervals, the limit of the sum is an integral. Shreve gives the finite disjoint-event expectation rule used here[3]. Tuckman and Serrat's CDS table also weights cash flows by survival and default events[4].

7.1.3 The hazard rate gives the probabilities in the formulas

Let $\lambda(u)$ be the risk-neutral default intensity at time u, given survival to u. It is a deterministic function of time, measured per model-year under the risk-neutral-probability-measure $\mathbb{Q}$. It generalizes the constant Constant hazard rate λ of this lesson. Its meaning is the conditional intensity of the credit lesson, but it need not be constant: $\lambda(u)$ is the proportional rate at which the survival probability decreases at u, given survival to u. As a differential equation, this meaning is:

$$-\,dS^{\mathbb{Q}}(0,u) = \lambda(u)\, S^{\mathbb{Q}}(0,u)\, du, \qquad S^{\mathbb{Q}}(0,0) = 1 \tag{7.1.5}$$

Dividing by the (positive) survival probability turns the left side into a derivative of its logarithm, $-d\ln\left(S^{\mathbb{Q}}(0,u)\right) = \lambda(u)\, du$. Integrating from 0 to t and using $\ln\left(S^{\mathbb{Q}}(0,0)\right) = \ln(1) = 0$ gives the survival probability:

$$S^{\mathbb{Q}}(0,t) = \exp\left(-\int_0^t \lambda(u)\; du\right) \tag{7.1.6}$$

Integrating the same differential equation over one period, not from 0, gives the interval default probability for a general hazard rate:

$$\Delta q_i^{\mathbb{Q}} = \int_{t_{i-1}}^{t_i} \lambda(u)\, S^{\mathbb{Q}}(0,u)\, du = S^{\mathbb{Q}}(0,t_{i-1}) - S^{\mathbb{Q}}(0,t_i) \tag{7.1.7}$$

The constant Constant hazard rate λ of this lesson is the special case $\lambda(u) \equiv \lambda$. Then $\int_0^t \lambda(u)\, du = \lambda t$, and the two results above become:

$$S^{\mathbb{Q}}(0,t) = \exp(-\lambda t), \qquad \Delta q_i^{\mathbb{Q}} = \exp(-\lambda t_{i-1}) - \exp(-\lambda t_i) \tag{7.1.8}$$

The next two sections, on the premium leg and on the protection leg, use the general $\lambda(u)$. The section on closed forms after them uses the constant λ.

These probabilities are pricing probabilities, not forecasts, for a constant and for a general hazard rate. A hazard rate implied from CDS prices can be a risk-neutral quantity rather than a prediction of realized defaults[5]. The constant-hazard survival relation is given in §14.6 and Appendix A14.1[6].

7.1.4 Default in the middle of a period

Suppose default occurs at u in $(t_{i-1}, t_i]$. The input recovery-rate is R.

- The scheduled premium at t_i is not paid, because the reference entity did not survive to that date.

- The buyer pays premium accrued from t_{i-1} to u. In this model its year fraction is $a_i(u) = u - t_{i-1}$.
- The seller pays protection of $\text{LGD}N = (1 - R)N$.
- Both amounts triggered by default are discounted from the modeled default time u, not from t_i.

Hull notes that a final accrued premium is paid when default occurs between scheduled dates[2]. The half-period and midpoint calculations in Hull and in Tuckman and Serrat are discrete approximations[7][8]. This lesson integrates over the modeled default time exactly, for a general hazard rate and a general discount curve. The closed forms later in the lesson use a constant hazard rate and a constant interest rate.

7.1.5 Exact premium-leg expectation in this model

Let $D(0, u)$ be the general dated discount curve of the rates lesson; this section does not assume a flat rate. The cds-scheduled-premium-annuity is:

$$A_0^{\text{sched}} = \sum_{i=1}^{n} \alpha_i D(0, t_i) S^{\mathbb{Q}}(0, t_i) \tag{7.1.9}$$

The exact cds-accrued-premium-annuity is:

$$A_0^{\text{accr}} = \sum_{i=1}^{n} \int_{t_{i-1}}^{t_i} a_i(u) D(0, u) \lambda(u) S^{\mathbb{Q}}(0, u) \, \mathrm{d}u \tag{7.1.10}$$

The cds-premium-annuity is the sum of the two annuities. The premium-leg present value is the annual spread, as a decimal, multiplied by the notional and by the premium annuity:

$$A_0^{\text{prem}} = A_0^{\text{sched}} + A_0^{\text{accr}}, \qquad PV_0^{\text{prem}} = sNA_0^{\text{prem}} \tag{7.1.11}$$

A spread quoted in basis-point per year is divided by 10,000 before it is used in (7.1.11).

7.1.6 Exact protection-leg expectation in this model

The protection-leg present value, as a positive magnitude, is:

$$PV_0^{\text{prot}} = N\text{LGD} \sum_{i=1}^{n} \int_{t_{i-1}}^{t_i} D(0, u) \lambda(u) S^{\mathbb{Q}}(0, u) \, \mathrm{d}u \tag{7.1.12}$$

Equation (7.1.12) is an expectation. Each possible payoff at a default time is discounted and multiplied by the risk-neutral probability density of that default time. The integral then adds over all default times.

7.1.7 Closed forms and the zero-upfront par spread

This section uses a constant Constant hazard rate, $\lambda(u) \equiv \lambda$, which is a special case of the cds-general-hazard-rate. It also uses a constant continuously-compounded-risk-free-rate r. So the general discount curve of the last two sections is $D(0,u) = \exp(-ru)$.

For the regular model-year schedule, $\alpha_i = t_i - t_{i-1}$ and the maturity-time is $T = t_n$. Define the abbreviation $\kappa = r + \lambda$. Then the scheduled-premium annuity and the accrued-premium annuity are:

$$A_0^{\text{sched}} = \sum_{i=1}^{n} \alpha_i \exp(-\kappa t_i) \tag{7.1.13}$$

$$A_0^{\text{accr}} = \sum_{i=1}^{n} \lambda \exp(-\kappa t_{i-1}) \frac{1 - \exp(-\kappa\alpha_i)(1 + \kappa\alpha_i)}{\kappa^2} \tag{7.1.14}$$

The cds-protection-factor, per unit of notional, is:

$$B_0^{\text{prot}} = \text{LGD}\, \lambda \frac{1 - \exp(-\kappa T)}{\kappa} \tag{7.1.15}$$

At $\kappa = 0$, the continuous limits replace the two fractions by ${\alpha_i}^2/2$ and T, respectively. The implementation uses numerically stable forms of those limits.

The cds-par-spread $s^\star$ is the running spread that makes the legs equal with no upfront amount:

$$s^\star = \frac{B_0^{\text{prot}}}{A_0^{\text{prem}}}, \qquad PV_0^{\text{prem}}(s^\star) = PV_0^{\text{prot}} \tag{7.1.16}$$

Tuckman and Serrat state the condition that the fair spread makes the two legs equal, and the relationship between spread and upfront[9]. The next lesson distinguishes this model par spread from a standard running coupon and a signed upfront amount.

Example 7.1.1: Exact default-time CDS legs

A finite event partition

Suppose a two-period protection payoff, simplified to one value per event, is USD 6,000 on G_1, USD 5,800 on G_2, and zero on the survival event G_0. The risk-neutral probabilities of G_1, G_2, and G_0 are 0.02, 0.03, and 0.95. The expected payoff before discounting is:

$$0(0.95) + 6{,}000(0.02) + 5{,}800(0.03) = \text{USD } 294$$

This finite sum is exact only because the payoff is constant on each event. In an exact CDS valuation, the discounted payoff varies with the default time inside an event. So the exact valuation uses an integral over all default times.

One-year exact legs

Take notional $N = \text{USD } 1{,}000{,}000$, running spread $s = 0.01$ per year (100 bp/year), recovery $R = 0.40$, constant risk-neutral hazard $\lambda = 0.02$ per model-year, continuously compounded risk-free rate $r = 0.04$ per model-year, two equal periods, and maturity $T = 1$ model-year.
The tested domain code gives:

Factor	Value
Scheduled premium annuity	0.956105034 model-years
Accrued-premium annuity	0.004828691 model-years
Premium annuity	0.960933725 model-years
Protection factor	0.011647093 per unit notional

The premium-leg present value is:

$$PV_0^{\text{prem}} = 0.01(1{,}000{,}000)(0.960933725) \approx \text{USD } 9{,}609.34$$

The protection-leg present value is:

$$PV_0^{\text{prot}} = 1{,}000{,}000(0.011647093) \approx \text{USD } 11{,}647.09$$

The present value of the accrued premium is USD 48.29. It is small in this example. The model calculates it at every possible default time; it does not approximate it by one half of a period.

Solve the exact par spread

With the same inputs, the par spread is:

$$s^\star = \frac{0.011647093}{0.960933725} \approx 0.012120600 = 121.206 \text{ bp/year}$$

At the unrounded par spread, the present value of each leg is approximately USD 11,647.09. The notional cancels in the par spread, because both legs are proportional to the same notional.

7.1.8 Generalizing the premium leg: an extinguishing swap

This lesson uses the name **extinguishing swap** for the construction below. The name belongs to this course; it is not a standard market term or textbook term. An extinguishing swap is any schedule of cash flows that stops permanently at the default of the reference entity. The premium leg above has this property.

Consider the same dates $0 = t_0 < t_1 < \cdots < t_n = T$ and the same reference entity. Replace the scheduled premium $sN\alpha_i$ at each date with an arbitrary deterministic, signed net cash flow c_i for one named perspective. Under the signed cash-flow convention of this course, a positive c_i is received and a negative c_i is paid. As with the premium leg, a scheduled amount is paid only on survival to its date. Every later scheduled amount is extinguished at the default time. In this lesson, c_i is an input and deterministic; it does not depend on a floating rate. If a cash flow accrues continuously between dates, an accrued-at-default term can be added with the same construction as the cds-accrued-premium-annuity. This lesson does not repeat that term.

The expectation from disjoint events, used for both CDS legs above, gives the cds-extinguishing-swap-present-value of the scheduled cash flows. It is the same survival-weighted discounting that gave A_0^{sched}:

$$PV_0^{\text{ext}} = \sum_{i=1}^{n} c_i \, D(0,t_i) \, S^{\mathbb{Q}}(0,t_i) \tag{7.1.17}$$

With $c_i = sN\alpha_i$ for every i, the present value is sNA_0^{sched}, the scheduled-payment term of the premium leg. So the premium leg is the special case of an extinguishing swap whose cash flows are one running spread paid on a fixed notional.

Example 7.1.2: A two-date extinguishing swap

A two-date extinguishing swap

Use the model of the "One-year exact legs" example: $\kappa = r + \lambda = 0.06$ per model-year, with semiannual dates $t_1 = 0.5$ and $t_2 = 1$. The discounted-survival factors are:

$$D(0,t_1)S(0,t_1) = \exp(-0.06\times 0.5) \approx 0.970445534, \qquad D(0,t_2)S(0,t_2) = \exp(-0.06\times 1) \approx 0.941764534$$

Suppose this extinguishing swap's schedule is a USD 30,000 receipt at t_1 and a USD 12,000 payment at t_2, so $c_1 =$ USD 30,000 and $c_2 = -$USD 12,000. The present value of the extinguishing swap is:

$$PV_0^{\text{ext}} = 30{,}000(0.970445534) + (-12{,}000)(0.941764534) \approx \text{USD } 17{,}812.19$$

Both terms use the same $D(0,t_i)S(0,t_i)$ factors as the premium leg. Only the signed cash flows c_i are different from the premium leg's $sN\alpha_i$.

7.1.9 Why there is no volatility input here

Risk-neutral valuation does not remove volatility. A change from the real-world measure to $\mathbb{Q}$ changes the pricing probabilities, and in common diffusion models it changes the drift. It does not change the size of the random fluctuations. Shreve shows that the Brownian volatility term remains after the risk-neutral change of measure[10].

This CDS model needs no separate volatility parameter for two reasons. First, its risk-free rate, hazard rate, and recovery rate are deterministic. Second, both legs are expectations that are linear in the cash flows. In other cases, the distributions, the dependence, and the volatility of the inputs could matter: if rates, the hazard rate, or recovery were stochastic, or if the product were a CDS option, had wrong-way risk, or had nonlinear exposure. The risk-neutral measure gives the pricing weights. The restrictive model assumptions, not the measure, remove the need for a separate volatility input.

7.1.10 Both legs for one contract

Figure 7.1.1 shows the present value of each leg for one contract in this model. The explorer integrates the accrued premium and the protection payment over the exact modeled default time.

7.1.11 Check your understanding

The assessment provides direct and transfer evidence for leg perspective, premium-leg magnitude, protection-leg magnitude, and par spread.

Knowledge check 7.1.1: CDS legs and par spread

1. Under the lesson sign convention, both CDS legs are reported as positive present-value magnitudes. Which expression gives the protection buyer's signed net value?

   (a) Protection-leg magnitude minus premium-leg magnitude

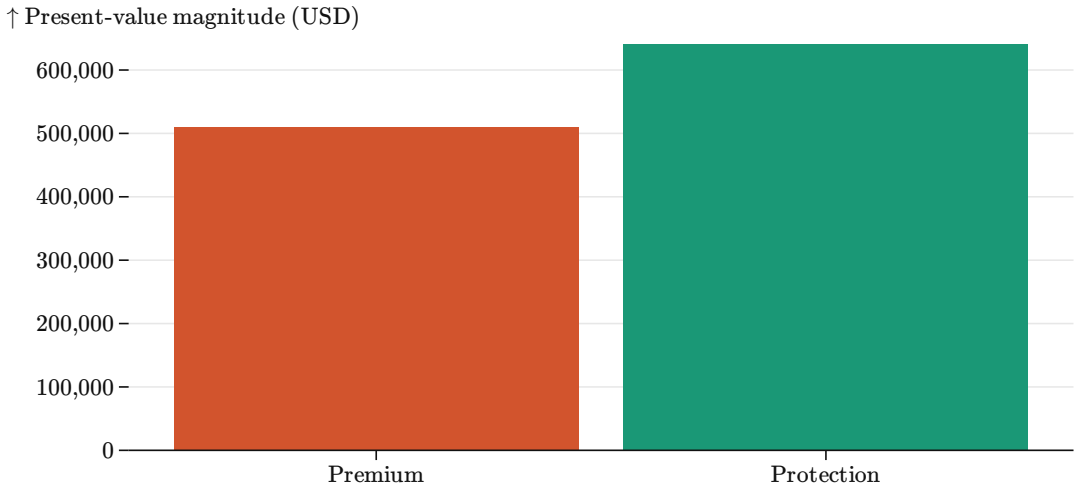


Figure 7.1.1 **CDS leg explorer.** The present values of the premium leg and the protection leg, as positive magnitudes in USD, for a notional of USD 10,000,000 and four equal model periods a year. The protection buyer's value is the protection leg minus the premium leg. Shown for a contractual spread of 120.0 bp/year, a constant risk-neutral hazard rate of 2.5%/year, a recovery rate of 40%, a continuously compounded risk-free rate of 4% and a maturity of 5 years: premium leg USD 509,703, protection leg USD 640,321.

(b) Premium-leg magnitude minus protection-leg magnitude

(c) Protection-leg magnitude plus premium-leg magnitude

2. A CDS has a positive protection-leg present-value magnitude of USD 35 and a positive premium-leg present-value magnitude of USD 28. What is the protection buyer's signed net value in USD?

   (a) +7

   (b) -7

   (c) +63

3. In the lesson's exact default-time model, a one-year annual-pay CDS has scheduled-premium annuity 0.941764533584 model-years and exact accrued-on-default annuity 0.009608857782 model-years. Calculate the positive premium-leg present-value magnitude in USD for notional USD 1,000,000 and running spread 100 basis points (0.01 per year).

4. A two-year quarterly CDS has notional USD 3,000,000 and running spread 125 basis points (0.0125 per year). Exact integration under the input constant-hazard model gives scheduled-premium annuity 1.829782217887 model-years and accrued-on-default annuity 0.011512762522 model-years. Calculate the positive premium-leg magnitude in USD.

5. For a one-year CDS, exact integration over modeled default time gives a discounted risk-neutral default-density integral of 0.019411822139. Calculate the positive protection-leg magnitude in USD for notional USD 1,000,000 and recovery rate 0.40.

6. Use the exact flat-hazard protection formula $N(1-R)\lambda[1-\exp(-(r+\lambda)T)]/(r+\lambda)$. Calculate the positive protection-leg magnitude in USD for N = USD 3,000,000, $R = 0.35$, $\lambda = 0.05$ per year, $r = 0.03$ per year, and $T = 2$ years.

7. Calculate the zero-upfront par spread in basis points per year when the exact premium annuity is 0.951373391366 model-years and the exact protection factor is 0.011647093283 per unit notional.

8. A different exact default-time model produces scheduled-premium annuity 1.829782217887 model-years, accrued-on-default annuity 0.011512762522 model-years, and protection factor 0.060066585732 per unit notional. Calculate its zero-upfront par spread in basis points per year.

Answers: *end of the chapter.*

7.1.12 Model boundary and review note

The valuation is exact only inside this simplified model. The premium leg and the protection leg are derived for a general, possibly time-varying risk-neutral hazard rate and a general dated discount curve. The closed forms, the numeric examples, the explorer, and the assessment use a special case: one constant risk-neutral hazard rate, one deterministic continuously compounded risk-free rate, a deterministic recovery rate, and settlement at the default time. The extinguishing swap is a construction of this lesson that generalizes the premium leg; it is not a named market instrument. The model omits calendar schedules, day counts, business-day rules, payment delays, term-structure calibration, auctions, deliverables, credit-event legal definitions, counterparty risk, collateral, funding, restructuring, and transaction costs. All materially changed content remains ‘draft‘ pending independent human review.

References

[1] Tuckman & Serrat, Fixed Income Securities: Tools for Today’s Markets (4th ed., 2022). §14.5, printed pp. 361–362.

[2] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). §24.1, printed pp. 548–550.

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

[4] Tuckman & Serrat, Fixed Income Securities: Tools for Today’s Markets (4th ed., 2022). §14.6, printed pp. 368–370, Table 14.10 and eqs. 14.6–14.7.

[5] Tuckman & Serrat, Fixed Income Securities: Tools for Today’s Markets (4th ed., 2022). §14.7, printed p. 371.

[6] Tuckman & Serrat, Fixed Income Securities: Tools for Today’s Markets (4th ed., 2022). printed pp. 367–368 and 505, eqs. 14.4–14.5 and A14.1–A14.2.

[7] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). §24.2, printed pp. 551–554, Tables 24.1–24.4.

[8] Tuckman & Serrat, Fixed Income Securities: Tools for Today’s Markets (4th ed., 2022). §14.6, printed pp. 368–370, Table 14.10.

[9] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). §14.6, printed pp. 370–371, and Appendix A14.2, printed pp. 506–507.

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

Answers to the knowledge checks

Knowledge check 7.1.1: CDS legs and par spread

1. **(a)** Protection-leg magnitude minus premium-leg magnitude. The protection leg is an inflow and the premium leg is an outflow for the protection buyer, so signed net value is protection minus premium.

2. **(a)** +7. From the protection buyer's perspective, net value is 35 - 28 = +7.

3. 9513.73391366 (accepted within ±0.01). The exact total premium annuity is 0.941764533584 + 0.009608857782 = 0.951373391366 model-years. The positive premium magnitude is 0.01 × 1,000,000 × 0.951373391366 = USD 9,513.73391366.

4. 69048.56176534 (accepted within ±0.01). The total exact premium annuity is 1.841294980409 model-years. Multiplying 0.0125 × 3,000,000 × 1.841294980409 gives USD 69,048.56176534. The accrued factor comes from integrating default time, not a half-period replacement.

5. 11647.0932834 (accepted within ±0.01). Loss given default is 1 - 0.40 = 0.60. The exact protection magnitude is 1,000,000 × 0.60 × 0.019411822139 = USD 11,647.0932834.

6. 180199.757197 (accepted within ±0.01). Here r + λ = 0.08. Substitution gives 3,000,000 × 0.65 × 0.05 × [1 - exp(-0.16)] / 0.08 = USD 180,199.757197. This discounts over actual modeled default time through the integral's closed form.

7. 122.423996602 (accepted within ±0.001). The par annual decimal spread is 0.011647093283 / 0.951373391366 = 0.0122423996602. Multiplying by 10,000 gives 122.423996602 basis points per year; notional cancels.

8. 326.219244453 (accepted within ±0.001). The exact total premium annuity is 1.841294980409. The par spread is 0.060066585732 / 1.841294980409 = 0.0326219244453 per year, or 326.219244453 basis points per year.

Chapter 7.2

CDS market-standard quote and upfront

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/cds/market-standard-quote-and-upfr

7.2.1 What you will be able to do

After this lesson, you should be able to:

- distinguish a conventional spread quote from the contract's fixed running coupon;
- identify which quantity is an input and which is solved in each quotation direction;
- calculate the signed upfront amount that makes protection-buyer inception value zero;
- reverse the simplified conversion without incorrectly freezing a quote-dependent premium annuity.

7.2.2 Three quantities with three different roles

Standardized CDS quotation separates a fixed running coupon from a market-determined upfront amount. Textbook treatments describe how to convert a quoted spread, treated as a par spread, into the value of a contract with a fixed coupon. The official converter specification accepts either the spread or the upfront and returns the other. [1] [2] [3]

This lesson uses the name **market-standard quote** (**MSQ**) for this conventional quoted spread. The name belongs to this course. The books and the official specification more commonly say *quoted spread*, *conventional spread*, or *spread*. This lesson does not claim that "MSQ" is a universal abbreviation.

The three quantities have different roles, but they are not three independent inputs:

- the cds-market-standard-quote s_{MSQ} is treated as a zero-upfront cds-par-spread and calibrates this lesson's flat Constant hazard rate λ;
- the cds-standard-coupon c_{std} determines the running premium cash flows;
- the cds-upfront-amount U_0 offsets the difference in value at inception.

The phrase "upfront fee" is also common. This lesson uses **upfront amount**, because the amount is signed. The sign of U_0 means:

- $U_0 > 0$: the protection buyer pays;
- $U_0 = 0$: no upfront cash changes hands in the toy model;
- $U_0 < 0$: the protection buyer receives $-U_0$.

In the formulas, rates are decimals per year. Convert a basis-point quote to decimal form by dividing by 10,000.

7.2.3 First decide which quantities are inputs and which are solved

The equation that connects the par spread and the hazard rate can be used in two directions. So a quantity can be an input in one calculation and solved in another. Table 7.2.1 shows the inputs and the solved quantities of each calculation.

Table 7.2.1 **Inputs and solved quantities of each calculation.** For the model par spread and for each direction of the quote conversion, the quantities taken as inputs and the quantities solved or calculated.

Calculation	Input	Solved or calculated
Model par spread	Hazard rate, recovery, discount input, maturity, and cash-flow conventions	Par spread $s^\star$
MSQ to upfront	Market quote s_{MSQ}, standard coupon c_{std}, and the converter conventions	Flat hazard rate λ, then signed upfront U_0
Upfront to MSQ	Market upfront U_0, standard coupon c_{std}, and the converter conventions	Flat hazard rate λ, then s_{MSQ}

If the hazard rate is an input, the par spread is the output:

$$s^\star = \frac{PV_0^{\mathrm{prot}}(\lambda)}{NA_0^{\mathrm{prem}}(\lambda)} \tag{7.2.1}$$

If the MSQ is the input instead, (7.2.1) is a calibration equation: the quoted spread is the input, and the flat hazard rate is the unknown. If the upfront amount is the input, the upfront balance is the calibration equation. The market participant or the price source provides the market level. The standard converter specifies how the spread and the upfront are converted into each other under common assumptions. [4]

So do not fix both s_{MSQ} and λ independently. When the other model inputs are fixed, s_{MSQ} and λ must together satisfy the zero-upfront condition.

7.2.4 Why one upfront amount is enough at inception

For fixed model inputs, one positive cds-protection-leg-present-value summarizes all future protection cash flows. One positive cds-premium-leg-present-value summarizes all fixed-coupon cash flows. The difference of the two present values is one currency amount. So one cash transfer at time zero can offset that difference exactly. The upfront balance is:

$$PV_0^{\text{prot}} = PV_0^{\text{prem}} + U_0 \tag{7.2.2}$$

Equivalently, the value to the protection buyer at inception, after the signed upfront amount, is zero:

$$0 = PV_0^{\text{prot}} - PV_0^{\text{prem}} - U_0 \tag{7.2.3}$$

This is a value identity, not a claim that the CDS becomes riskless. After inception its value can change when credit, interest-rate, recovery, liquidity, or other inputs change.

7.2.5 Risk-neutral does not mean risk-free

The calibration uses **risk-neutral pricing probabilities**. They are not necessarily real-world default forecasts. Under risk-neutral valuation, probabilities are pricing weights, chosen consistently with traded prices and risk-free discounting. "Risk-free" describes the discount input, not the probability measure. So the hazard rate inferred from prices can differ from an estimated real-world default intensity. [5] [6]

A change of probability measure does not remove volatility. This converter has no separate volatility input, for two reasons. Its discount rate, hazard rate, recovery rate, and schedule are deterministic. Its leg values are expectations that are linear in the one-default cash flows. A model with stochastic rates, stochastic default intensity, stochastic recovery, wrong-way dependence, or CDS options can require volatility and correlation inputs. The risk-neutral measure gives pricing weights; it does not remove those risks.

7.2.6 Convert MSQ to signed upfront

The conversion must use one internally consistent model. For each candidate constant hazard rate, calculate the positive cds-premium-annuity $A_0^{\text{prem}}(\lambda)$ and the protection-leg present value. First, solve the zero-upfront quote condition for λ:

$$s_{\text{MSQ}} N A_0^{\text{prem}}(\lambda) = PV_0^{\text{prot}}(\lambda) \tag{7.2.4}$$

The root λ of (7.2.4) is the flat risk-neutral hazard rate implied by the MSQ, under the assumptions of this lesson on the schedule, discounting, recovery, and accrual at default. The premium-leg present value with the fixed coupon is then:

$$PV_0^{\text{prem}}(c_{\text{std}}; \lambda) = c_{\text{std}} N A_0^{\text{prem}}(\lambda) \tag{7.2.5}$$

Substituting (7.2.4) and (7.2.5) into the upfront balance (7.2.2) gives the upfront amount:

$$\begin{aligned} U_0 &= PV_0^{\text{prot}}(\lambda) - c_{\text{std}} N A_0^{\text{prem}}(\lambda) \\ &= N A_0^{\text{prem}}(\lambda) \left(s_{\text{MSQ}} - c_{\text{std}} \right) \end{aligned} \tag{7.2.6}$$

The substitution in (7.2.6) is valid only because every term uses the same model and the same calibrated hazard rate. The premium annuity and the notional are positive. So an MSQ above the fixed coupon gives a positive upfront amount, which the buyer pays. An MSQ below the fixed coupon gives a negative upfront amount, which the buyer receives.

7.2.7 Convert signed upfront back to MSQ

When the input is a signed upfront amount, do not rearrange the last line of (7.2.6) with the annuity held fixed. In this flat-hazard converter, both the protection-leg present value and the premium annuity change with the unknown hazard rate. Instead, solve this equation for the hazard rate:

$$U_0 = PV_0^{\text{prot}}(\lambda) - c_{\text{std}} N A_0^{\text{prem}}(\lambda) \tag{7.2.7}$$

Solve (7.2.7) for a non-negative λ. Then (7.2.8) gives the MSQ from this λ:

$$s_{\text{MSQ}} = \frac{PV_0^{\text{prot}}(\lambda)}{N A_0^{\text{prem}}(\lambda)} \tag{7.2.8}$$

If a curve and its annuity have already been calibrated and are held fixed, rearranging (7.2.6) gives a shortcut:

$$s_{\text{MSQ}} = c_{\text{std}} + \frac{U_0}{N A_0^{\text{prem}}(\lambda)} \tag{7.2.9}$$

This shortcut is not the full conversion from upfront to quote when the quote itself determines λ, because then the denominator changes with the quote.

7.2.8 When MSQ and par spread are equal

Within one internally consistent risk-neutral model, the MSQ and the par spread are equal. After the converter has inferred one λ, from the spread or from the upfront amount, the following quantities are equal:

$$\begin{aligned} s_{\text{MSQ}} &= s^\star \\ &= \frac{PV_0^{\text{prot}}(\lambda)}{N A_0^{\text{prem}}(\lambda)} \\ &= c_{\text{std}} + \frac{U_0}{N A_0^{\text{prem}}(\lambda)} \end{aligned} \tag{7.2.10}$$

The equality in (7.2.10) is a risk-neutral pricing identity. No real-world default probability enters it. The equality requires the same pricing probabilities, discount inputs, recovery, schedule, default-accrual treatment, and settlement convention in every term.

The phrase **par spread from another curve** needs more care. The observed trade with a standard coupon gives one combined price:

$$U_0 = PV_0^{\text{prot}} - c_{\text{std}} N A_0^{\text{prem}} \tag{7.2.11}$$

This price does not show the protection-leg present value and the premium annuity separately. The standard converter separates them by assuming its constant-hazard model. A different risk-neutral term structure can reprice the same contract, with its fixed coupon and upfront amount, but with a different premium annuity. That term structure then gives a different par coupon to a hypothetical zero-upfront contract. The official converter itself distinguishes its single constant hazard rate from the term structure of a fuller credit-curve model. [7]

This difference is not an arbitrage by itself, because the observed fixed-coupon CDS and the hypothetical zero-upfront CDS have different premium cash flows. No arbitrage requires identical tradable cash flows to have the same price. But suppose a zero-upfront CDS that pays s_{MSQ} could

be traded on otherwise identical terms. Then a curve calibrated consistently to that traded contract would have to give it zero value. [8]

Example 7.2.1: Quote and upfront conversions

MSQ to buyer-paid upfront

Consider a five-year CDS with quarterly premiums and notional N = USD 10,000,000, 2% continuously compounded risk-free rate, 40% recovery, MSQ $s_{\text{MSQ}} = 0.015$ per year (150 basis points per year), and fixed coupon $c_{\text{std}} = 0.010$ per year (100 basis points per year).
Under the flat-hazard model of this lesson, with exact default times, the calibration to the quote gives:

$$\lambda = 0.024937617 \text{ per year}, \qquad A_0^{\text{prem}}(\lambda) = 4.466915290 \text{ model-years}$$

Because the calibrated quote is a par spread, the protection-leg present value is:

$$PV_0^{\text{prot}} = 0.015(10{,}000{,}000)(4.466915290) = \text{USD } 670{,}037.29$$

The premium-leg present value with the fixed coupon is:

$$PV_0^{\text{prem}} = 0.010(10{,}000{,}000)(4.466915290) = \text{USD } 446{,}691.53$$

The upfront amount is the difference of the two present values:

$$U_0 = 670{,}037.29 - 446{,}691.53 = \text{USD } 223{,}345.76$$

The upfront amount is positive, so the protection buyer pays it. As a check, $670{,}037.29 - 446{,}691.53 - 223{,}345.76 = 0$, up to the rounding of the displayed values.

Why reverse conversion is a solve

Suppose the input for the same contract is the signed upfront amount U_0 = USD 223,345.76, and the MSQ is unknown. A solver changes the flat λ and recalculates both the protection-leg present value and the premium annuity, until (7.2.7) holds.
The root is approximately $\lambda = 0.024937617$ per year. Only with this root does (7.2.8) give $s_{\text{MSQ}} \approx 0.015$ (150 basis points per year).
The shortcut with the annuity fixed at 4.466915290 gives the same answer here, but only because this annuity was already calculated at the solution's hazard rate. Before the solve, this annuity is unknown.

7.2.9 What a real standard converter adds

The official converter uses dates and market conventions. Its inputs include the trade date and the maturity date, the notional, the standard coupon, the recovery rate, the discount inputs, and either the spread or the upfront. Its outputs include the other quotation and a cash-settlement amount. [3]

This lesson is not an implementation of that standard model. It uses model-year periods and a single flat hazard rate. A production conversion must also handle calendar schedules, day-count and business-day conventions, curve construction and interpolation, settlement timing, the accrued amount, and other contract details. The identities of this lesson explain why the quote, the coupon, and the upfront amount differ. But the numbers in this lesson are not cash-settlement instructions for a trade.

7.2.10 Check your understanding

The assessment separates quote interpretation, signed upfront arithmetic, and the reverse flat-hazard solve. Each competency has direct and transfer evidence.

Knowledge check 7.2.1: CDS market-standard quote and upfront

1. This lesson uses MSQ as a local abbreviation for market-standard quote. Which statement matches its simplified conversion model?

   (a) MSQ is the conventional quoted spread treated as a zero-upfront par spread for calibrating the flat hazard rate

   (b) MSQ is always the fixed running coupon paid on the contract

   (c) MSQ is the signed upfront cash amount divided by notional

2. A simplified converter is given an MSQ of 180 basis points per year and a standard coupon of 100 basis points per year. Which rate determines the scheduled running premium cash flows?

   (a) The 100-basis-point standard coupon

   (b) The 180-basis-point MSQ

   (c) The 80-basis-point difference

3. Calculate the signed upfront amount in USD when notional is USD 10,000,000, the quote-implied premium annuity is 4.0 model-years, MSQ is 150 basis points per year, and the fixed coupon is 100 basis points per year. Positive means paid by the protection buyer.

4. Calculate the signed upfront amount in USD when notional is USD 5,000,000, the quote-implied premium annuity is 2.5 model-years, MSQ is 300 basis points per year, and the fixed coupon is 500 basis points per year. Positive means paid by the protection buyer; enter a negative number if the buyer receives cash.

5. For a curve that is already calibrated and held fixed, calculate the market-standard quote in basis points per year from notional USD 8,000,000, premium annuity 4.0 model-years, fixed coupon 100 basis points per year, and signed upfront USD 160,000 paid by the protection buyer.

6. Use the lesson's one-year flat-hazard model with one annual period, zero continuously compounded risk-free rate, recovery 40%, notional USD 10,000,000, and fixed coupon 1% per year. For hazard rate lambda, A(lambda) = (1 - exp(-lambda))/lambda and the par quote is (1 - 0.40)lambda. A signed upfront of USD 195,082.302 is paid by the protection buyer. Root solve U0/N = A(lambda)[(1 - 0.40)lambda - 0.01], then report the quote in basis points per year.

Answers: *end of the chapter.*

7.2.11 Model boundary and review note

This AI-assisted lesson and all cited locators remain `draft` pending human editorial and quantitative review. All input values are invented for teaching. The converter assumes one deterministic discount rate, one deterministic recovery rate, one constant risk-neutral hazard rate, regular model-year periods, exact default-time accrued premium inside that toy schedule, and an upfront payment at time zero. It omits the contract details and numerical methods listed above, as well as stochastic rates, stochastic intensity or recovery, volatility and correlation calibration, counterparty risk, collateral, funding, liquidity, and transaction costs.

References

[1] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.6, printed pp. 366-371; Appendix A14.2, printed pp. 506-507, eqs. A14.5-A14.7.

[2] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.4, printed pp. 556-557.

[3] International Swaps and Derivatives Association, ISDA CDS Standard Model and Standard CDS Contract Converter Specification. Standard CDS Contract Converter Specification, version May 5, 2009, Functionality and Specification, printed p. 1. https://www.cdsmodel.com/documentation.html

[4] International Swaps and Derivatives Association, ISDA CDS Standard Model and Standard CDS Contract Converter Specification. Standard CDS Contract Converter Specification, version May 5, 2009, Functionality and Specification, printed pp. 1 and 4. https://www.cdsmodel.com/documentation.html

[5] 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; Ch. 14 §14.7, printed p. 371.

[6] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 23 §23.5, printed pp. 528-530.

[7] International Swaps and Derivatives Association, ISDA CDS Standard Model and Standard CDS Contract Converter Specification. Standard CDS Examples, version October 26, 2012, printed p. 2 (the converter uses a single flat hazard rate rather than a term structure of flat spreads). https://www.cdsmodel.com/documentation.html

[8] 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, replication, and arbitrage enforcement).

Answers to the knowledge checks

Knowledge check 7.2.1: CDS market-standard quote and upfront

1. **(a)** MSQ is the conventional quoted spread treated as a zero-upfront par spread for calibrating the flat hazard rate. MSQ is a lesson-local label for the conventional quoted spread. The toy converter treats it as a zero-upfront par spread to infer one flat risk-neutral hazard rate; the contract can have a different fixed running coupon.

2. **(a)** The 100-basis-point standard coupon. The fixed standard coupon determines running premiums. The MSQ calibrates the pricing model, and the upfront amount bridges the resulting protection value and fixed-coupon premium value.

3. 200000 (accepted within ± 0.01). Convert the rates to 0.015 and 0.010. Then U0 = N A (sMSQ - cstd) = 10,000,000 $\times$ 4.0 $\times$ 0.005 = USD 200,000. It is positive, so the protection buyer pays it.

4. -250000 (accepted within ± 0.01). U0 = 5,000,000 $\times$ 2.5 $\times$ (0.03 - 0.05) = -USD 250,000. The high fixed coupon makes the running premium leg exceed protection value, so the protection buyer receives the negative signed upfront.

5. 150 (accepted within ± 0.001). With the calibrated annuity held fixed, sMSQ = cstd + U0/(N A) = 0.01 + 160,000/(8,000,000 $\times$ 4.0) = 0.015, or 150 basis points per year. A full upfront-to-quote conversion cannot hold A fixed while the quote-implied hazard rate changes.

6. 300 (accepted within ± 0.001). The root is lambda = 0.05 per year: A = (1 - exp(-0.05))/0.05 = 0.97541151 and A $\times$ (0.60 $\times$ 0.05 - 0.01) = 0.0195082302, matching U0/N. The quote is 0.60 $\times$ 0.05 = 0.03, or 300 basis points per year.

Chapter 7.3

Credit curve and market observables

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/cds/credit-curve-and-market-observ

7.3.1 What you will be able to do

After this lesson, you should be able to:

- classify each quantity in a standard CDS valuation as market-observed, contractual, conventional, or model-solved;
- explain why the traded upfront alone is enough to determine the pricing hazard rate, and why that is a consequence of risk-neutral pricing rather than an extra assumption;
- read a credit curve as tenor marks together with the survival curve fitted to them, and tell a change of quotation from a curve fit;
- calculate the survival probability implied by fitted hazard segments and check that a tenor reprices to its mark.

7.3.2 A traded price is an expectation

The risk-neutral lesson established that, in an arbitrage-free model with the cash account as numeraire, a claim's present-value is the expectation under the risk-neutral-probability-measure $\mathbb{Q}$ of its discounted payoff; see (1.7.2). Shreve states the general discounted-payoff formula [1].

A CDS is such a claim. Let the local cds-buyer-discounted-cash-flows Π^{buyer} be the random sum of the protection buyer's discounted signed cash flows over the life of the contract, before any upfront amount. It is the protection received at the modeled default time minus the scheduled and accrued premium paid. Each cash flow is discounted from its own payment time to valuation-time 0. The present value of the CDS is the risk-neutral expectation of this sum:

$$PV_0 = \mathbb{E}^{\mathbb{Q}}\left[\Pi^{\text{buyer}}\right] \tag{7.3.1}$$

Now add the one fact that the market gives. A new standard contract is traded at its market upfront amount: the buyer pays the cds-upfront-amount U_0 at time zero and receives the claim valued by (7.3.1). A trade at the market level has zero present value at inception. So the upfront amount equals the present value:

$$U_0 = \mathbb{E}^{\mathbb{Q}}\left[\Pi^{\text{buyer}}\right] = PV_0^{\text{prot}} - PV_0^{\text{prem}} \tag{7.3.2}$$

The second equality follows from the linearity of expectation: the expectations of the protection cash flows and of the premium cash flows are calculated separately. They are the two positive leg present values of the preceding lessons. Equation (7.3.2) is the same identity as (7.2.2), read in the other direction. In the preceding lesson, the model parameters were inputs and the upfront was the output. Here, the upfront is observed, and the equation constrains the model.

Why one equality is enough. The standard contract has no cash flow at inception other than the upfront amount. So the trade gives exactly one number. In the model of this course, discounting and recovery are deterministic, and both legs are linear expectations over the default time. So the only unknown in (7.3.2) is the timing of default under the pricing measure: the survival-probability curve. Setting the expected values of the two legs equal, with the upfront amount on the premium side, is therefore not a modeling choice added to risk-neutral pricing. It is what risk-neutral pricing says about a contract that has just been traded at the market level. Tuckman and Serrat describe the market convention in these steps. First, find the hazard rate that makes the expected discounted fee leg and contingent leg (in this course, the premium leg and the protection leg) equal. Then value the standard coupon at that hazard rate [2].

One number determines one unknown. With a single flat Constant hazard rate λ per contract, as in the preceding lesson, one traded price is exactly enough. With a term structure of hazard rates, one traded price per tenor is exactly enough only if the curve has one free parameter per tenor. The piecewise-constant curve below has this property. With fewer prices than parameters, the curve is under-determined, and a convention must fix the remaining parameters. With more prices than parameters, the curve can in general only fit the prices approximately (a best fit).

7.3.3 What the market shows, the model assumes, and the model solves

If every quantity is called an input, it is not clear which quantities the market can contradict. Table 7.3.1 classifies the quantities of a standard contract at one tenor.

The risk-neutral interpretation has two consequences.

- **The solved probabilities are pricing weights.** Hull describes implying default probabilities from CDS quotes as the analogue of implying volatilities from option prices. Hull states that CDS valuation needs risk-neutral, not real-world, default probabilities [10]. A survival curve fitted to quotes is a statement about prices, not a forecast of default frequencies. Tuckman and Serrat make the same point for the price-implied hazard rate [11].
- **Recovery must be assumed because one price cannot identify two unknowns.** The protection payoff is proportional to LGD, and the implied hazard rate is approximately proportional to 1/LGD. So a contract valued with the same recovery rate that implied its hazard rate is not very sensitive to the choice of recovery rate. A binary CDS, whose payoff does not depend on recovery, would add the second equation needed to separate the two [12].

7.3.4 The credit curve is two objects

A discount curve is a set of dated discount factors in one currency and on one funding basis; it is not one yield (see the rates lesson). A credit curve is the same idea for one reference entity. But a

Table 7.3.1 **Observed, contractual, convention, and solved quantities.** Each quantity of a standard contract at one tenor: whether the market shows it, the contract fixes it, a convention assumes it, or the model solves it, and where it comes from.

Quantity	Status	Where it comes from
Signed upfront U_0, or the cds-market-standard-quote it stands for	**Observed**	Dealer bid and offer for the standard contract; a mid is taken as the mark. [3]
Risk-free discount-factor $D(0,t)$	**Observed** in another market, then an input	Built from the rates market, not from the CDS; the converter takes it as a discount input. [4]
cds-standard-coupon c_{std}, cds-notional N, maturity, credit events, reference obligations	**Contractual**	Written into the standard contract; the coupon is 100 or 500 basis points and maturities fall on standard dates. [5] [6]
recovery-rate R and hence loss-given-default LGD	**Convention**	Assumed by the converter, commonly 40% but 20% or 25% for some contract types. It is not observed in the quote. [7]
Shape of the hazard rate between and beyond tenors, accrual treatment	**Convention**	The converter's single flat hazard rate, or this lesson's piecewise-constant curve. [8]
Constant hazard rate λ, the survival curve, the cds-par-spread $s^\star$, the other quotation	**Solved**	Root of (7.3.2) given everything above. [9]

credit curve has two layers, and this course keeps them separate.

The marked curve. Use the local cds-tenor-index k across the cds-tenor-count K quoted tenors with maturities $T_1 < \cdots < T_K$. At each tenor, the market shows the price of the standard contract. A mark is the recorded price at one tenor. If the trading desk records the price as a conventional spread, the mark is the local cds-tenor-quote s_k. If the desk records the traded price itself, the mark is the local cds-tenor-upfront-fraction u_k. This course calls a curve whose marks are market-standard quotes an **MSQ-marked curve**, and one whose marks are upfront fractions an **upfront-marked curve**. Both labels are local to this course (`NEEDS_SOURCE`: the desk vocabulary for marking conventions is not documented in the registered sources). Tuckman and Serrat tabulate five-year spreads for several reference entities. They also plot the term structure of spreads across tenors: its slope is normally upward, but it is inverted in periods of stress [13]. Hull lists the tenors that commonly trade; the five-year tenor is the most liquid [14].

The fitted curve. The legs lesson allows a deterministic time-varying hazard rate. It shows that the survival probability is the exponential of minus the integral of the hazard rate. The fitted credit curve is the special case in which the hazard rate is constant on each segment between consecutive tenors: the local cds-segment-hazard-rate λ_k on $(T_{k-1}, T_k]$, with $T_0 = 0$. The survival probability is then a product with one factor per segment:

$$S^{\mathbb{Q}}(0,T_k) = S^{\mathbb{Q}}(0,T_{k-1}) \exp\Big(-\lambda_k \left(T_k - T_{k-1}\right)\Big), \qquad S^{\mathbb{Q}}(0,0) = 1 \tag{7.3.3}$$

Inside a segment, the survival probability decreases at the constant hazard rate of that segment, as in the constant-hazard lesson. Hull notes that when spreads for different maturities are available, the hazard rate is calculated as a step function [15].

The marked curve is data. The fitted curve is a model object, chosen so that the model reprices the data. The next distinction depends on keeping the two layers separate.

7.3.5 Transforming marks is not fitting the curve

Curve transformation changes the quotation of each mark and does not build a term structure. For each tenor separately, the standard converter takes s_k and solves for a flat hazard rate for that tenor alone. This course writes this rate as the local cds-flat-tenor-hazard-rate λ_k^{flat}. The converter then values the fixed coupon at that rate and returns u_k. The converter can also do the same steps in the reverse direction. This is the conversion of the preceding lesson, applied once per tenor [16]. The flat hazard rates of different tenors are solutions of separate equations. They are not one survival curve, and the official converter itself distinguishes its single flat hazard rate from a term structure [17]. The same tenor-by-tenor conversion also gives a third quotation: the par spread of a hypothetical zero-upfront contract. The name **curve transformation** belongs to this course (`NEEDS_SOURCE` for the phrase itself).

Fitting the curve finds one term structure that reprices every mark at once. With the marks written as upfront fractions, the fit condition at tenor k is the buyer-value identity (7.3.2), evaluated with the segment hazard rates up to k:

$$PV_0^{\text{prot}}(T_k; \lambda_1, \dots, \lambda_k) - c_{\text{std}} N A_0^{\text{prem}}(T_k; \lambda_1, \dots, \lambda_k) = u_k N \tag{7.3.4}$$

Both legs are the exact default-time integrals of the legs lesson, with the piecewise-constant hazard rate substituted. Each premium period lies inside one segment. So each period is a one-period flat-hazard model that starts from the survival probability and the discount factor at the start of the period. Solve (7.3.4) in maturity order. The first tenor determines λ_1. Then λ_1 is held fixed, and the second tenor determines λ_2, and so on. Each step is a one-dimensional root solve for λ_k. The buyer value of a contract increases with the hazard rate of any segment inside the life of the contract, so each step has at most one root. A zero-upfront par mark is the special case $u_k = 0$, with the coupon set to the quoted par spread. The result is a curve under which every quoted tenor reprices exactly. The survival probabilities of this curve at all tenors are consistent with each other.

The same survival curve then prices every other contract on the same reference entity: a maturity between the quoted tenors, an existing contract with a different coupon, or a hypothetical zero-upfront contract. This is the purpose of a fitted curve. The fit has two limits. First, the fitted par spread at a tenor is in general not equal to the MSQ mark of that tenor. The reason is that the fitted premium annuity differs from the flat annuity of the converter, although both reprice the same traded upfront amount. This is the "par spread from another curve" point of the preceding lesson. Second, a solution need not exist. A later mark that is cheaper than the zero-hazard value implied by the earlier segments would need a negative λ_k. The model rejects a negative hazard rate, because a survival probability cannot increase with time.

Example 7.3.1: A credit curve from market marks

Transform the marks

Take an MSQ-marked curve for one reference entity: standard contracts with quarterly premiums and a fixed coupon $c_{\text{std}} = 0.010$ per year (100 basis points per year), continuously compounded risk-free rate 2% per year, recovery $R = 0.40$, and marks $s_1 = 0.008$, $s_2 = 0.012$, $s_3 = 0.016$ at maturities $T_1 = 1$, $T_2 = 3$, $T_3 = 5$ model-years. Notional is USD 10,000,000.
For each tenor separately, the flat converter of the preceding lesson gives:

Tenor k	s_k (bp/yr)	λ_k^{flat}, per yr	Flat A_0^{prem}, model-yrs	u_k	U_0, USD	Flat $S(0,T_k)$
1	80	0.013300046	0.981077842	−0.001962156	−19,621.56	0.986788009
2	120	0.019950083	2.820140305	0.005640281	56,402.81	0.941905574
3	160	0.026600129	4.449093452	0.026694561	266,945.61	0.875464525

Each row is the conversion of the preceding lesson. For example, at five years, $u_3 = 4.449093452 \times (0.016 - 0.010) = 0.026694561$. The one-year mark is below the coupon, so the buyer receives the upfront amount. The three flat hazard rates are solutions of three separate equations. The last column is not one survival curve: for example, it does not define the survival probability at two years.

Fit the curve

Now fit a piecewise-constant hazard curve to the three upfront fractions in maturity order.

Tenor k	Segment, model-yrs	λ_k, per yr	Fitted $S(0,T_k)$	Fitted A_0^{prem}, model-yrs	Fitted PV_0^{prot}/N	Fitted $s^\star$ (bp/yr)
1	(0, 1]	0.013300046	0.986788009	0.981077842	0.007848623	80.0000
2	(1, 3]	0.023463765	0.941550234	2.829254583	0.033932826	119.9356
3	(3, 5]	0.037572831	0.873389867	4.500220799	0.071696769	159.3183

Every tenor reprices. For example, at five years, $0.071696769 - 0.010 \times 4.500220799 = 0.026694561 = u_3$. In the tested implementation, the residuals at all three tenors are below 10^{-12}.

Compare this table with the transformation. The hazard rate of the first segment equals the one-year flat rate, because a one-year contract depends only on the first segment. The hazard rates of the later segments are higher than the flat rates. The five-year flat rate, 0.026600, is an average over five years. The fitted hazard rate on the last two years is 0.037573, because the shorter tenors, with lower hazard rates, already price the earlier years. So the fitted five-year survival probability, 0.873390, is below the flat survival probability, 0.875465. The fitted par spreads at three and five years, 119.94 and 159.32 basis points, are not equal to the marks of 120 and 160 basis points. The reason is that the fitted curve reproduces the same traded upfront amounts with different premium annuities.

The fitted curve also defines the survival probability between tenors. By (7.3.3), the survival probability at two years is $0.986788009 \exp(-0.023463765 \times 1) = 0.963903771$, and at four years it is $0.941550234 \exp(-0.037572831 \times 1) = 0.906829881$.

The marks are prices of particular contracts

Now read the same three numbers, 80, 120, and 160 basis points, as zero-upfront par spreads, not as MSQ marks with a coupon of 100 basis points. So $u_k = 0$, and the coupon of each tenor equals its quote. The fit then gives segment hazard rates of 0.013300046, 0.023480220, and 0.037863368 per year, and a five-year survival probability of 0.872853786 instead of 0.873389867. The first segment is unchanged, because the converter also treats the one-year MSQ of 80 basis points as the par spread of a zero-upfront one-year contract. The later segments change. A standard contract with a coupon of 100 basis points and an upfront amount is a different stream of cash flows from a zero-upfront contract with a coupon of 120 basis points. The two prices are written with the same number, but they are prices of different contracts. A mark is the price of a particular contract, and the fit must know which contract.

When no curve fits

Suppose the one-year zero-upfront par mark is 300 basis points and the three-year mark is 100 basis points. The first segment prices the one-year contract. Then the three-year contract at

100 basis points is worth more to the seller than the first year alone implies, even with zero default risk after year one. So its mark is below the zero-hazard bound, and the fit reports that the mark would need a negative segment hazard rate. In the tested implementation, a three-year mark of 110 basis points has a solution, with a hazard rate of 0.001573 per year in the second segment.
Steeply inverted marks are not necessarily an arbitrage, but they indicate a problem. Either the marks are out of date or not at market levels, or the piecewise-constant shape and the deterministic recovery of the model cannot describe them.

7.3.6 A note on jargon

The vocabulary of CDS trading desks mixes contract terms, quotation conventions, and risk measures. Table 7.3.2 gives this course's reading of the most common phrases, with the registered source where one supports it. Entries marked `NEEDS_SOURCE` are recorded as usage this course adopts locally; a human should confirm or replace them before review.

7.3.7 Check your understanding

The assessment separates the observed-versus-solved classification, the transformation-versus-fit distinction, and the survival and repricing arithmetic of a fitted curve. Each competency has direct and transfer evidence.

Knowledge check 7.3.1: Credit curve and market observables

1. In the standard conversion of a five-year contract's quote into an upfront, which quantity is solved from the traded price rather than observed in the market or assumed by convention?

   (a) The flat risk-neutral hazard rate

   (b) The fixed running coupon

   (c) The recovery rate

2. A dealer shows five-year protection at 160 basis points bid and 165 basis points offer, and the converter uses a 40% recovery rate. Which statement is correct?

   (a) The 40% recovery is a convention assumed by the converter; the quote does not observe it

   (b) The quote reveals that the market expects 40% recovery

   (c) The dealer is quoting the hazard rate directly

3. An MSQ-marked curve lists 80, 120, and 160 basis points at one, three, and five years. Converting each mark to an upfront with the standard converter, one tenor at a time, is best described as which of the following?

   (a) A change of quotation using a separate flat hazard rate per tenor; it does not build one term structure

   (b) A fit of one piecewise-constant hazard curve that reprices all three tenors together

   (c) A direct observation of the survival probabilities at one, three, and five years

Table 7.3.2 **CDS desk jargon.** Common phrases of CDS trading desks with this course's reading of each and the registered source that supports it. Entries marked NEEDS_SOURCE are local usage for a human to confirm.

Phrase	Meaning in this course
Reference entity, credit event, notional	The named borrower, the contractually defined trigger, and the amount protected. [18] [19]
Buying protection, selling protection; short credit, long credit	A protection buyer pays premium and gains on default, like a short seller of the bond; a protection seller resembles a bond holder. [20]
Index buyer and seller	For index CDS the words are reversed: the index buyer receives premium and pays on default, like a portfolio holder. [21]
Spread, quoted spread, conventional spread, MSQ	The annualized premium of a hypothetical zero-upfront contract, used as the conventional quote; this course's local label is market-standard quote. [22] [14]
Par spread	The running spread that makes the two legs equal with zero upfront under a stated curve; see the cds-par-spread card.
Standard coupon, points upfront, price	The fixed 100 or 500 basis point running coupon and the upfront that balances it; the upfront per 100 notional is quoted as points, or as a price of 100 minus points. [23] [24]
Risky annuity, risky duration, RPV01	The premium annuity: the present value of one unit of running spread per unit notional, survival-weighted; Hull calls it the duration of the CDS payments. The label RPV01 is `NEEDS_SOURCE`. [24]
Risky DV01, spread DV01, CS01	The change in a CDS value for a one basis point change in the CDS spread, found by re-solving the hazard rate and revaluing; the next lesson uses it. The label CS01 is `NEEDS_SOURCE`. [25]
Duration times spread, DTS	A spread-risk measure for bonds that scales spread duration by the spread level, motivated by spreads moving proportionally. [25]
On-the-run, off-the-run	The most liquid current standard contract of a tenor versus earlier contracts with older maturity dates. [26]
IMM dates, roll	Standard maturity dates on the twentieth of March, June, September, and December; contracts and index series move to new dates on a schedule. [23] [27]
Recovery convention	The assumed recovery in the converter: commonly 40%, with 20% or 25% for some contract types. [7]
Credit curve, spread curve, term structure of spreads	Marks across tenors and the survival curve fitted to them; upward-sloping in calm markets and inverted in stress. [21]
MSQ-marked curve, upfront-marked curve, curve transformation, curve fitting	Local labels defined in this lesson (`NEEDS_SOURCE` for desk usage).
Equivalent notional, equivalent ratio	The liquid-tenor notional whose quote sensitivity offsets a position, and its ratio to the position notional; defined in the next lesson (`NEEDS_SOURCE` for desk usage).
CDS-bond basis, negative basis trade	The CDS spread minus a bond spread measure; a negative basis trade buys the bond and buys protection. [28] [29]
CDS-equivalent bond spread, hazard-adjusted duration	The CDS spread corresponding to the hazard rate that reprices a bond, and a bond duration computed under the hazard model. [30]
Cheapest-to-deliver, auction, final price	The delivery option among reference obligations and

4. Under a fitted piecewise-constant curve, the three-year contract's par spread is 119.94 basis points while its MSQ mark is 120 basis points, and the fit reprices the three-year upfront exactly. Why can both numbers be right?

   (a) They are par spreads under two different survival curves that reprice the same traded upfront with different premium annuities

   (b) The fit failed to reprice the three-year tenor

   (c) The difference is an arbitrage between the two quotations

5. Under the fitted curve, the five-year tenor has protection factor 0.071696769 per unit notional and premium annuity 4.500220799 model-years, and its standard coupon is 100 basis points per year. Calculate the upfront fraction of notional, paid by the protection buyer, that the fit reproduces.

6. A fitted curve has a segment hazard rate of 0.02 per year on (0, 2] and 0.05 per year on (2, 5] model-years. Calculate the survival probability at four years.

Answers: *end of the chapter.*

7.3.8 Model boundary and review note

This AI-assisted lesson and all cited locators remain `draft` pending human editorial and quantitative review. All input values are invented for teaching. The fitted curve is piecewise constant with knots at quoted tenors, extended flat, under one deterministic discount rate, one deterministic recovery rate, and exact default-time accrual on a regular model-year schedule. A production curve also handles calendar schedules, day counts, business-day rules, settlement timing, accrued premium at inception, an interpolation and extrapolation policy chosen by the trading desk, and a discount curve built from the rates market. The lesson does not model bid-offer, liquidity, dealer inventory, stochastic rates or intensity, recovery risk, counterparty risk, collateral, funding, or transaction costs, and it makes no claim that pricing-model survival probabilities forecast realized defaults.

References

[1] Shreve, Stochastic Calculus for Finance II: Continuous-Time Models (2004). Ch. 5 §5.2.4, printed pp. 218-219, eqs. 5.2.29-5.2.31.

[2] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.6, printed pp. 368-370.

[3] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.1, printed p. 549 (market-maker bid and offer in basis points per year).

[4] International Swaps and Derivatives Association, ISDA CDS Standard Model and Standard CDS Contract Converter Specification. Standard CDS Contract Converter Specification, version May 5, 2009, Functionality and Specification, printed p. 1. https://www.cdsmodel.com/documentation.html

[5] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.6, printed pp. 366-367 (standard maturity dates on the twentieth of March, June, September, and December; 100 or 500 basis point coupons with a market-determined upfront).

[6] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.4, printed p. 556 (a coupon and a recovery rate are specified for each underlying and maturity).

[7] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.6, printed p. 368, n. 15.

[8] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.6, printed p. 370 (the quote-to-upfront translation is a convention like price and yield; participants need not believe in constant hazard or a particular recovery, but all use the convention).

[9] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.4, printed pp. 556-557 (a hazard rate is implied from the quoted spread by iterative search).

[10] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.2, printed p. 554.

[11] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.7, printed p. 371.

[12] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.2, printed pp. 554-555.

[13] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.5, printed pp. 362-363, Table 14.9 and Figure 14.5.

[14] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.1, printed p. 549.

[15] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.2, printed p. 554, n. 6.

[16] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.4, printed pp. 556-557.

[17] International Swaps and Derivatives Association, ISDA CDS Standard Model and Standard CDS Contract Converter Specification. Standard CDS Examples, version October 26, 2012, printed p. 2 (the converter uses a single flat hazard rate rather than a term structure of flat spreads). https://www.cdsmodel.com/documentation.html

[18] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.1, printed pp. 548-549.

[19] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.5, printed p. 361.

[20] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.5, printed pp. 364-365.

[21] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.5, printed p. 363.

[22] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.5, printed p. 362.

[23] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.6, printed p. 366.

[24] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.4, printed p. 557.

[25] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.10, printed p. 379.

[26] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.6, printed pp. 366-367.

[27] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.3, printed pp. 555-556.

[28] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.8, printed pp. 374-375.

[29] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.1, printed pp. 550-551.

[30] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.7, printed p. 371; §14.9, printed p. 377.

[31] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.1, printed pp. 549 and 551.

[32] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.5, printed p. 363; §14.13, printed p. 387.

[33] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.3, printed p. 555.

[34] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.6, printed p. 367, n. 14.

[35] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.1, printed p. 550.

Answers to the knowledge checks

Knowledge check 7.3.1: Credit curve and market observables

1. **(a)** The flat risk-neutral hazard rate. The coupon is contractual and the recovery rate is a convention assumed by the converter. The hazard rate is the unknown that the one traded price identifies through the equality of expected discounted leg values.

2. **(a)** The 40% recovery is a convention assumed by the converter; the quote does not observe it. One traded price cannot identify both recovery and default timing, so the converter assumes recovery and solves the hazard rate. The bid and offer are the observed market levels; the recovery is not among them.

3. **(a)** A change of quotation using a separate flat hazard rate per tenor; it does not build one term structure. The tenor-by-tenor conversion is a curve transformation. Each tenor's flat hazard rate answers a separate question, so the three flat rates are not one survival curve; fitting the curve is a separate step that solves segment hazards in maturity order.

4. **(a)** They are par spreads under two different survival curves that reprice the same traded upfront with different premium annuities. The MSQ mark is the par spread under the converter's single flat hazard for that tenor; the fitted par spread uses the term structure's premium annuity. Both reproduce the same traded upfront, and the hypothetical zero-upfront contract is not itself traded, so there is no arbitrage.

5. 0.026694561 (accepted within ± 0.000001). The fit condition per unit notional is protection factor minus coupon times premium annuity: 0.071696769 - $0.010 \times 4.500220799 = 0.026694561$, which equals the five-year mark.

6. 0.869358235 (accepted within ± 0.000001). Survival compounds segment by segment: exp(-0.02 $\times$ 2) $\times$ exp(-0.05 $\times$ 2) = exp(-0.14) = 0.869358235. Using only the second segment's rate over four years would be wrong.

Chapter 7.4

Equivalent notional and quote risk

Interactive version: https://danieltuzes.github.io/credit-and-derivatives/cds/equivalent-notional-and-quote-

7.4.1 What you will be able to do

After this lesson, you should be able to:

- calculate a CDS position's signed value change for a one basis point increase in the liquid tenor's quote, re-solving the hazard rate rather than holding the premium annuity fixed;
- explain why, in a one-factor model, risk measured against the market-standard quote, the upfront, or the par spread is the same risk;
- calculate the equivalent notional and side of the most liquid contract that offsets a position, and its equivalent ratio;
- state what an equivalent-ratio hedge does not offset.

7.4.2 One market level moves everything

The preceding lessons value every contract on a reference entity with one flat Constant hazard rate λ. The quote lesson showed how the market level determines λ: the cds-market-standard-quote s_{MSQ} of the standard contract is treated as a par spread, and λ is the root of (7.2.4). In this lesson, that contract is the most liquid one, with the local cds-liquid-tenor-maturity T_L. In practice, the most liquid tenor is usually five years [1] [2].

A position is one contract on the same reference entity with its own maturity, fixed coupon, and positive local cds-position-notional N_{pos}. Its legs are the exact default-time expectations of the legs lesson, evaluated at the same λ. Let the local cds-position-side σ be $+1$ for bought protection and -1 for sold protection. The local cds-position-value is the protection-buyer value of the legs lesson, multiplied by the side:

$$V^{\mathrm{pos}} = \sigma \left(PV_0^{\mathrm{prot}} - PV_0^{\mathrm{prem}} \right) \tag{7.4.1}$$

Hull marks a CDS to market in the same way, with opposite signs for buyer and seller [3]. Any upfront exchanged at inception is a sunk cash flow and is left out of V^{pos}; it does not change any sensitivity.

For the liquid contract, the local cds-liquid-unit-value is the buyer value per unit of notional:

$$v_L = \frac{PV_0^{\text{prot}}(T_L)}{N} - c_{\text{std}} A_0^{\text{prem}}(T_L) \tag{7.4.2}$$

This value is the fair upfront fraction of the quote lesson. So at the current market level, it equals the upfront amount of the liquid contract per unit of notional [4].

Because λ is the only state variable, every quotation of the liquid contract is a function of λ. Write the local cds-liquid-quotation q for any one of them: the market-standard quote, the upfront fraction, or the par spread. In this model, each quotation is strictly increasing in λ, so each quotation is invertible. So a change of the quote is a change of the hazard rate, whichever quotation the trading desk uses for its marks. The model has one risk factor, expressed in three quotations. In this model, with a flat hazard rate and a flat interest rate, the par spread is also the same at every maturity. So the par spread of a position and the liquid quote are the same number.

7.4.3 Risky DV01 by bump and reprice

Tuckman and Serrat define the risky DV01 of a CDS as the change in its value for a one basis point change in the CDS spread, computed by finding the hazard rate that moves the spread by one basis point and revaluing the contract [5]. This lesson uses that bump-and-reprice definition with the local cds-quote-bump $h = 0.0001$ per year, one basis-point of market-standard quote.

Write $\lambda(s_{\text{MSQ}})$ for the hazard rate re-solved from a quote. The local cds-position-quote-sensitivity is the change in the position value when the quote increases by h:

$$\text{DV01}^{\text{pos}} = V^{\text{pos}}\Big(\lambda\left(s_{\text{MSQ}} + h\right)\Big) - V^{\text{pos}}\Big(\lambda\left(s_{\text{MSQ}}\right)\Big) \tag{7.4.3}$$

The local cds-liquid-unit-quote-sensitivity is the same difference, for one unit of bought liquid protection:

$$\text{DV01}^{L} = v_L\Big(\lambda\left(s_{\text{MSQ}} + h\right)\Big) - v_L\Big(\lambda\left(s_{\text{MSQ}}\right)\Big) \tag{7.4.4}$$

Three properties of the definition are important in practice.

- **Both legs change.** A higher hazard rate increases the protection leg and decreases the premium annuity. So the sensitivity is not the annuity multiplied by one basis point. Holding the annuity fixed is the shortcut that the quote lesson warned against.
- **The sign depends on the side.** Bought protection gains value when the quote increases, so its sensitivity is positive. Sold protection has a negative sensitivity. DV01^{L} is defined for bought protection and is positive.
- **The bump is finite.** The value is a nonlinear function of the quote. So the result of a one basis point bump differs slightly from the derivative. The section “What an equivalent-ratio hedge does not do” discusses this difference.

7.4.4 The liquid-tenor equivalent

A hedge that removes the quote risk is a position in the liquid contract whose notional makes the sensitivity of the hedged position zero. Let the local cds-equivalent-notional N_{eq} be the notional of bought liquid protection. The hedge condition and its solution are:

$$\mathrm{DV01}^{\mathrm{pos}} + N_{\mathrm{eq}}\,\mathrm{DV01}^{L} = 0 \qquad \Longleftrightarrow \qquad N_{\mathrm{eq}} = -\frac{\mathrm{DV01}^{\mathrm{pos}}}{\mathrm{DV01}^{L}} \tag{7.4.5}$$

A positive N_{eq} means: buy protection in the liquid tenor. A negative N_{eq} means: sell protection in the liquid tenor. Because $\mathrm{DV01}^{L}$ is positive, the sign of N_{eq} is always opposite to the sign of the position's sensitivity. The local cds-equivalent-ratio is the size of the hedge relative to the position:

$$\rho_{\mathrm{eq}} = \frac{|N_{\mathrm{eq}}|}{N_{\mathrm{pos}}} \tag{7.4.6}$$

This hedge is the DV01-neutral hedge that Tuckman and Serrat describe for bonds. There, the face amount of a liquid hedge is minus the ratio of the two DV01s, multiplied by the position amount. The instrument with the larger sensitivity is traded in the smaller quantity [6]. The names **equivalent notional** and **equivalent ratio**, for the hedge amount and its ratio, belong to this course (`NEEDS_SOURCE` for desk usage; the registered sources define the DV01 hedge but not these names).

The choice of quotation does not change the hedge. A trading desk can measure the position's risk against the market-standard quote, the upfront, or a par spread. In this model, the choice does not change the equivalent notional. In the limit of a small bump, both sensitivities are derivatives with respect to the same quotation q. The chain rule through the one state variable λ gives:

$$N_{\mathrm{eq}} = -\frac{dV^{\mathrm{pos}}/dq}{dv_L/dq} = -\frac{(dV^{\mathrm{pos}}/d\lambda)\,(d\lambda/dq)}{(dv_L/d\lambda)\,(d\lambda/dq)} = -\frac{dV^{\mathrm{pos}}/d\lambda}{dv_L/d\lambda} \tag{7.4.7}$$

The factor $d\lambda/dq$, which converts a change of the hazard rate into a change of the quotation, cancels. So MSQ risk, upfront risk, and par spread risk are one risk, and the same equivalent notional offsets all three. The sensitivities themselves differ across quotations, because each sensitivity is measured per unit of a different quotation. The hedge amount is the same. With a finite bump, the three equivalent ratios differ only because of the curvature, as the worked example shows.

Example 7.4.1: Equivalent-notional hedges

A seven-year position hedged with five-year

The liquid contract is a five-year standard contract with quarterly premiums, $c_{\mathrm{std}} = 0.010$ per year (100 basis points per year) and market-standard quote $s_{\mathrm{MSQ}} = 0.016$ per year (160 basis points per year). The risk-free rate is 2% per year, continuously compounded, and the recovery rate is 40%. Calibration to the quote gives $\lambda = 0.026600129$ per year. By (7.4.2), one unit of bought five-year protection is worth $v_L = 0.026694561$ before the upfront amount.

The position is bought protection on a seven-year contract with the same coupon of 100 basis points and $N_{\mathrm{pos}} =$ USD 10,000,000. At the calibrated hazard rate, the protection-leg present value of the position is USD 953,285.99, and its premium-leg present value is USD 595,803.74. So by (7.4.1), $V^{\mathrm{pos}} =$ USD 357,482.24.

Bump the liquid quote by $h = 0.0001$, to 161 basis points. The re-solved hazard rate is 0.026766381 per year. Revaluing with (7.4.3) and (7.4.4) gives the two sensitivities:

$$\mathrm{DV01}^{\mathrm{pos}} = 363{,}240.36 - 357{,}482.24 = \text{USD } 5{,}758.12 \text{ per basis point}$$

$$\mathrm{DV01}^{L} = 0.000434070474 \text{ per unit notional per basis point}$$

By (7.4.5), the equivalent notional is:

$$N_{\text{eq}} = -\frac{5{,}758.117764}{0.000434070474} = -\text{USD } 13{,}265{,}398.37$$

So the hedge sells protection on a notional of USD 13,265,398 of the five-year contract. By (7.4.6), $\rho_{\text{eq}} = 1.3265$. The seven-year contract has the larger sensitivity per unit of notional. So the hedge notional in the five-year contract is larger than the notional of the position. The tested implementation reports a hedged sensitivity of zero to arithmetic precision.

A sold three-year position at a 500 basis point coupon

Keep the same liquid contract and market level. The position is sold protection on a three-year contract with coupon 500 basis points per year and $N_{\text{pos}} = \text{USD } 20{,}000{,}000$, so $\sigma = -1$. Because the coupon of 500 basis points is higher than the market level of 160 basis points, the seller's value is positive: $V^{\text{pos}} = \text{USD } 1{,}899{,}072.40$.
After the same bump of one basis point, the seller's value decreases to USD 1,893,025.81. So the sensitivity is:

$$\text{DV01}^{\text{pos}} = -\text{USD } 6{,}046.59 \text{ per basis point}$$

The seller loses value when the quote increases. With the same $\text{DV01}^{L} = 0.000434070474$, the equivalent notional is:

$$N_{\text{eq}} = -\frac{-6{,}046.591823}{0.000434070474} = +\text{USD } 13{,}929{,}977.23$$

So the hedge buys protection on a notional of USD 13,929,977 of the five-year contract, and $\rho_{\text{eq}} = 0.6965$. The three-year contract has a smaller sensitivity per unit of notional than the five-year contract. So a five-year notional smaller than the position notional offsets the position. This holds although the high coupon gives the position a large value.

Three quotations, one hedge

Repeat the seven-year hedge, but bump the upfront fraction of the liquid contract by 0.0001 (one basis point of notional) instead of its quote. Re-solve the hazard rate from the bumped upfront fraction. The equivalent ratio is 1.32670, compared with 1.32654 from the quote bump. With both bumps one thousand times smaller, both ratios are 1.32674, as (7.4.7) predicts. So the difference at one basis point is a result of the curvature of a finite bump, not of a different risk.
As a check, apply the construction to the liquid contract itself, with bought protection on USD 10,000,000. The equivalent ratio is exactly one, on the opposite side: the hedge is the offsetting trade.

7.4.5 What an equivalent-ratio hedge does not do

The hedge sets one sensitivity to zero, at one market level, in one model. It leaves at least the following exposures, and a trading desk manages each of them separately.

- **Default.** The seven-year hedge above sells protection on USD 13,265,398 against bought protection on USD 10,000,000. If default occurs immediately, with 40% recovery, the position receives $\text{LGD} \times 10{,}000{,}000 = \text{USD } 6{,}000{,}000$ and the hedge pays $0.6 \times 13{,}265{,}398.37 = \text{USD } 7{,}959{,}239.02$. The net loss is about USD 1,959,239, before accrued premium. Equal

quote sensitivities do not imply equal default payoffs. So the jump-to-default exposure of a hedged portfolio must be measured separately.

- **Curve shape.** The model has one hazard rate. Under the fitted curve of the credit-curve lesson, a seven-year position also depends on the segments beyond five years, on which the five-year contract does not depend. A change of the five-year quote alone would change the value of only part of the position. Measuring the sensitivity separately for each tenor is the standard extension (`NEEDS_SOURCE`: the registered sources describe single-factor DV01 hedging only).
- **Curvature.** Both sensitivities change when the quote changes. So the hedge stops being exact and must be rebalanced. Tuckman and Serrat make the same point for bond DV01 hedges [7].
- **Inputs held fixed.** The recovery rate, the discount curve, and the accrual and settlement conventions are constant in the model. A change in any of them changes the values of the two contracts by different amounts.
- **Costs.** The hedge pays the bid-offer spread of the liquid contract. The lower liquidity of the position's own tenor is the reason why the hedge is not the offsetting trade in that tenor [8].

7.4.6 Check your understanding

The assessment separates the bump-and-reprice sensitivity, including its sign for a seller, from sizing the equivalent notional and ratio. Each competency has direct and transfer evidence.

Knowledge check 7.4.1: Equivalent notional and quote risk

1. A bought-protection position is worth USD 357,482.24 at the liquid five-year quote of 160 basis points and USD 363,240.36 after the quote is bumped to 161 basis points with the flat hazard rate re-solved. Calculate the position's risky DV01 in USD per basis point.
2. A trader has sold protection on a three-year contract. Valued for a protection buyer, the contract is worth USD -1,899,072.40 at the liquid quote and USD -1,893,025.81 after a one basis point bump with the hazard rate re-solved. Calculate the seller's risky DV01 in USD per basis point, using the sign convention that a positive number is a gain for the seller.
3. A bought-protection position has risky DV01 of USD 5,758.117764 per basis point. One unit of notional of bought protection in the liquid five-year contract has risky DV01 of 0.000434070474 per basis point. Calculate the magnitude in USD of the liquid-tenor equivalent notional that offsets the position.
4. A sold-protection position with notional USD 20,000,000 has risky DV01 of USD -6,046.591823 per basis point for the seller. One unit of notional of bought protection in the liquid five-year contract has risky DV01 of 0.000434070474 per basis point. Calculate the equivalent ratio of the liquid-tenor hedge to the position notional.

Answers: end of the chapter.

7.4.7 Model boundary and review note

This AI-assisted lesson and all cited locators remain `draft` pending human editorial and quantitative review. All input values are invented for teaching. The sensitivity is a one basis point bump of the liquid tenor's market-standard quote, under one flat hazard rate, one deterministic discount rate, one deterministic recovery rate, and a regular quarterly model-year schedule. The hedge is sized at one market level and is not rebalanced. The lesson omits curve shape risk, default and recovery risk beyond the illustration above, bid-offer, transaction costs, calendar and settlement conventions, counterparty risk, collateral, funding, and any stochastic model of rates or intensity.

References

[1] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.1, printed p. 549.

[2] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.5, printed p. 365.

[3] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.2, printed pp. 553-554.

[4] Hull, Options, Futures, and Other Derivatives (8th ed., 2012). Ch. 24 §24.4, printed p. 557.

[5] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.10, printed p. 379.

[6] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 4 §4.3, printed pp. 110-111, eqs. 4.6-4.9.

[7] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 4 §4.2, printed p. 109 (DV01 as a local measure).

[8] Tuckman & Serrat, Fixed Income Securities: Tools for Today's Markets (4th ed., 2022). Ch. 14 §14.6, printed pp. 366-367 (unwinding through the on-the-run contract and managing the maturity and spread mismatch).

Answers to the knowledge checks

Knowledge check 7.4.1: Equivalent notional and quote risk

1. 5758.12 (accepted within ± 0.01). The bump-and-reprice sensitivity is the value after the bump minus the value before it: 363,240.36 - 357,482.24 = USD 5,758.12 per basis point. It is positive because bought protection gains when the quote rises.

2. -6046.59 (accepted within ± 0.01). The seller's value is minus the buyer's value: USD 1,899,072.40 before and USD 1,893,025.81 after the bump. The seller's sensitivity is 1,893,025.81 - 1,899,072.40 = USD -6,046.59 per basis point; a seller loses when the quote rises.

3. 13265398.4 (accepted within ± 1). $N_{\text{eq}} = \frac{-5{,}758.117764}{0.000434070474} = -$USD 13,265,398.4. The negative sign means the hedge sells protection in the five-year contract; its magnitude is USD 13,265,398.4.

4. 0.6965 (accepted within ± 0.0001). $N_{\text{eq}} = \frac{-(-6{,}046.591823)}{0.000434070474} = +\text{USD } 13{,}929{,}977.2$, so the hedge buys protection in the five-year contract. The equivalent ratio is 13,929,977.2 / 20,000,000 = 0.6965.

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.

$\mathbf{1}_{\{k=n\}}$ **Final payment indicator**. Equals one on the final payment date, where the payment index equals the number of payments, and zero on every other date. It places the face value on the final date only. *In Chapter 2.1 only.*

c **Annual coupon rate**. Contractual annual rate used to determine a fixed-rate bond's coupon payments; applied to face value, not to market price. *Units:* decimal rate per year.

CF_k^{bond} **Bond cash flow**. Promised amount paid to the bondholder on one scheduled payment date; a positive receipt for the bondholder, with default excluded in the bond lessons. *Units:* stated currency at payment time.

m_{B} **Bond payment frequency**. Number of scheduled coupon payments per year in the simplified regular bond, a positive integer. *Units:* scheduled coupon payments per year.

k **Bond payment index**. Labels one remaining scheduled bond payment in increasing time order; it selects a payment and is not itself a time or currency amount. *Units:* dimensionless schedule index.

C **Coupon payment**. Level periodic cash amount promised by the simplified fixed-rate bond; a positive receipt for the bondholder in these lessons. *Units:* stated currency per coupon date.

F **Face value**. Contractual reference amount used to determine coupons and principal redemption. *Units:* stated currency.

T **Maturity time**. Final scheduled time when principal is redeemed in the simplified bond. *Units:* model-years from the valuation date.

n **Number of bond payments**. Counts the remaining regular coupon dates including maturity; a positive integer for the simplified regular bond schedule. *Units:* scheduled payment dates.

P_0 **Bond price**. Present value of the simplified bond's promised payments at valuation time; the amount paid by the buyer, shown as a positive value in the bond lessons. *Units:* stated currency at valuation time.

$y^{(m_B)}$ **Yield to maturity**. Single nominal annual rate that reproduces the simplified bond price. *Units:* nominal annual decimal rate compounded at the bond payment frequency.

ΔP **Price change**. New toy-model price minus old toy-model price for a stated yield shock. This finite change is not yet a duration or convexity approximation. *Units:* stated currency per bond at valuation time. *In Chapter 2.4 only.*

$$\Delta P = P(y + \Delta y) - P(y)$$

Δy **Yield shock**. A lesson-local finite change in the input yield. The sign is new yield minus old yield; the unit must be decimal, percentage points, or basis points as explicitly labeled. *Units:* decimal rate per year. *In Chapter 2.4 only.*

$$\Delta y = y_{\text{new}} - y_{\text{old}}$$

$P(y)$ **Price-yield curve**. Bond price as a function of yield while promised positive fixed cash flows remain constant and only the yield varies. *Units:* stated currency per bond.

n_{elapsed} **Coupon period elapsed days**. Actual calendar days from the previous coupon date through settlement under the lesson convention. *Units:* days. *In Chapter 2.5 only.*

n_{period} **Coupon period total days**. Actual calendar days from the previous coupon date to the next coupon date. *Units:* days. *In Chapter 2.5 only.*

AI **Accrued interest**. Coupon amount attributed to the interval from the previous coupon date through settlement under the stated day-count convention; a positive amount added to the clean price to obtain the dirty invoice price. *Units:* stated currency at settlement.

P^{clean} **Clean bond price**. Bond price quoted without accrued interest under the stated settlement convention; a positive quoted price to the buyer before accrued interest is added. *Units:* stated currency at settlement.

P^{dirty} **Dirty bond price**. Full cash or invoice price paid for the bond, equal to clean price plus accrued interest; a positive cash price paid by the buyer under the stated settlement convention. *Units:* stated currency at settlement.

I_0^{pre} **Predelivery coupon present value**. Valuation-time present value of coupons paid to the current bond owner before forward delivery. *Units:* stated currency at valuation time per bond. *In Chapter 2.6 only.*

$F_{0,T_{\text{fwd}}}^{B}$ **Bond forward price**. Fair dirty delivery price fixed at valuation time for delivery of the named bond at the forward date; positive delivery cash paid by the long under the lesson's no-arbitrage assumptions. *Units:* stated currency at forward delivery per bond.

K_{fwd} **Forward delivery price**. Contractual currency amount per unit of underlying paid by the long at delivery and received by the short. *Units:* delivery-time currency per unit of underlying.

T_{fwd} **Forward delivery time**. Future model time when the forward counterparties exchange the underlying and delivery payment; a contract date shared by the long and short, not the underlying asset's maturity. *Units:* model-years from the stated valuation time.

T **Forward discount end time**. Later model time at which one currency unit is paid. *Units:* model-years from valuation time. *In Chapter 3.1 only.*

t **Forward discount start time**. Earlier future model time from which the later payment is valued. *Units:* model-years from valuation time. *In Chapter 3.1 only.*

$Z(t,T)$ **Forward discount factor**. Value at future model time t of one currency unit paid at later model time T under the stated deterministic curve, with t strictly before T. *Units:* currency at current time per currency-unit at future time.

N **Forward contract quantity**. Positive number of underlying units delivered under the forward contract. Contract value and payoff scale linearly with this quantity. *Units:* units of underlying. *In Chapter 4.1 only.*

$$N > 0$$

$\Pi_{T_{\text{fwd}}}^{\text{long}}$ **Long forward expiry payoff**. Signed delivery-time value received by the long after exchanging the fixed delivery payment for the underlying. Per unit it equals underlying value at delivery minus the contractual delivery price. *Units:* stated currency at forward delivery time. *In Chapter 4.1 only.*

S_t **Derivative underlying value**. Value at model time t of one unit of the asset or claim named as the derivative's underlying; a positive quoted value, while a position in the underlying carries its own signed quantity. *Units:* stated currency per unit of underlying at model time.

V_t^{fwd} **Forward contract value**. Signed current value of an existing forward from the named counterparty's perspective; positive to the long when the current forward price exceeds the contract's fixed delivery price under the lesson model. *Units:* stated currency at model time.

$F_{t,T_{\text{fwd}}}$ **Forward price**. Delivery price that would give a newly struck forward for the stated delivery time zero current value; a quoted contract rate rather than a cash amount received at quotation time. *Units:* delivery-time currency per unit of underlying.

I_0 **Underlying income present value**. Valuation-time value of deterministic cash income paid by the underlying after valuation and before the contract horizon, forward delivery or the matched option expiry. It is subtracted from spot value because the forward or option position does not receive that income. *Units:* stated currency at valuation time per unit of underlying.

$\Pi^{\text{call}}_{T_{\text{opt}}}$ **Long call expiry payoff**. Non-negative amount received by the holder of a European call at expiry. Per unit it is the positive part of underlying expiry value minus strike. *Units:* stated currency at option expiry. *In Chapter 4.2 only.*

$\Pi^{\text{put}}_{T_{\text{opt}}}$ **Long put expiry payoff**. Non-negative amount received by the holder of a European put at expiry. Per unit it is the positive part of strike minus underlying expiry value. *Units:* stated currency at option expiry. *In Chapter 4.2 only.*

$G_{T_{\text{opt}}}$ **Option holder expiry profit**. Expiry payoff less the premium carried to expiry under the explicitly stated financing convention. In the simplest examples financing is ignored, so the initially paid premium is subtracted directly. *Units:* stated currency at option expiry. *In Chapter 4.2 only.*

π_0 **Option premium paid**. Positive valuation-time amount paid by the option holder to acquire the contractual right. Its financing convention must be stated before converting an expiry payoff into profit. *Units:* stated currency at premium payment time. *In Chapter 4.2 only.*

c_t **European call value**. Current non-negative value to the holder of a European call under the stated model, before any financing or transaction costs. *Units:* stated currency at model time.

T_{opt} **Option expiry time**. Future model time when a European option's exercise decision and payoff are determined; a contract date shared by holder and writer, distinct from a bond maturity. *Units:* model-years from the stated valuation time.

K **Option strike price**. Contractual price per unit of underlying used to determine the option's exercise payoff; a positive contractual amount, with payoff signs depending on call or put and holder or writer perspective. *Units:* expiry-time currency per unit of underlying.

p_t **European put value**. Current non-negative value to the holder of a European put under the stated model, before any financing or transaction costs. *Units:* stated currency at model time.

$\widetilde{F}_{0,T_{\text{opt}}}$ **Prepaid forward value**. Current value of receiving the underlying at expiry with no delivery payment then. Under the known-income model it equals spot value less the present value of income paid before expiry. *Units:* stated currency at valuation time per unit of underlying. *In Chapter 4.3 only.*

V_0 **Binomial claim value**. Time-zero value obtained identically from the replicating portfolio and risk-neutral expected payoff. *Units:* stated currency at valuation time. *In Chapter 4.4 only.*

X^{down} **Binomial down claim payoff**. Signed expiry payoff of the claim when the underlying takes its down-state value. *Units:* stated currency at expiry. *In Chapter 4.4 only.*

S^{down} **Binomial down underlying value**. The lower of the two possible underlying values at the one-period expiry. *Units:* stated currency per unit of underlying at expiry. *In Chapter 4.4 only.*

Δ **Binomial hedge units**. Signed units of underlying in the portfolio that replicates the two claim payoffs. *Units:* units of underlying per claim. *In Chapter 4.4 only.*

$M_{T_{\text{opt}}}$ **Binomial risk-free expiry cash**. Signed expiry cash amount held with the underlying units to reproduce both claim payoffs. *Units:* stated currency at expiry. *In Chapter 4.4 only.*

q **Binomial risk-neutral up weight**. Model-implied pricing probability assigned to the up state by the no-arbitrage tree. It is dimensionless and lies between zero and one when the tree is arbitrage-consistent. *Units:* probability between zero and one. *In Chapter 4.4 only.*

X^{up} **Binomial up claim payoff**. Signed expiry payoff of the claim when the underlying takes its up-state value. *Units:* stated currency at expiry. *In Chapter 4.4 only.*

S^{up} **Binomial up underlying value**. One of the two possible underlying values at the one-period expiry. *Units:* stated currency per unit of underlying at expiry. *In Chapter 4.4 only.*

$CF_{i+1,j}$ **Lattice next cash flow**. Signed cash paid at a successor node immediately before its continuation value is measured. *Units:* stated currency at the successor time. *In Chapter 4.5 only.*

$d_{i,j}$ **Lattice node discount factor**. One-period discount factor at a lattice node: currency at the node per one unit of currency at either successor node one step later. *Units:* current-node currency per next-time currency.

$q_{i,j}$ **Lattice node up weight**. Pricing probability assigned to the up successor, conditional on the current node; in a lattice with issuer default, also conditional on survival to the next lattice time. *Units:* probability between zero and one.

$V_{i,j}$ **Lattice node value**. Claim value at time row i and state node j obtained by one-period backward induction from its successor nodes, conditional on reaching that node under the input pricing lattice. *Units:* stated currency at the node time.

j **Lattice state index**. Integer label for one state node within a time row of a finite recombining lattice; a bookkeeping label under the stated successor ordering, not a probability or state value. *Units:* dimensionless integer index.

i **Lattice time index**. Integer label for one time row in a finite recombining valuation lattice; a bookkeeping label, not a model-year time or currency amount. *Units:* dimensionless integer index.

$X^{B,\text{call}}_{T_{\text{opt}}}$ **Bond call expiry payoff**. Non-negative holder payoff from buying the bond at the dirty strike when economically beneficial. *Units:* stated currency per option at expiry. *In Chapter 5.1 only.*

$B^{\text{dirty}}_{T_{\text{opt}}}$ **Bond option expiry value**. Ex-coupon dirty cash value of the named bond at the European exercise instant. *Units:* stated currency per bond at option expiry. *In Chapter 5.1 only.*

$X^{B,\text{put}}_{T_{\text{opt}}}$ **Bond put expiry payoff**. Non-negative holder payoff from selling the bond at the dirty strike when economically beneficial. *Units:* stated currency per option at expiry. *In Chapter 5.1 only.*

$R^{B}_{i,j}$ **Bond default recovery cash**. Non-negative recovery paid on the underlying bond at the next lattice date after default in the period. *Units:* stated currency per bond at the next time. *In Chapter 5.2 only.*

$H^{B}_{i,j}$ **Surviving bond continuation value**. Next-time scheduled bond cash plus risk-neutral expected ex-cash-flow bond value, conditional on period survival. *Units:* stated currency per bond at the next time. *In Chapter 5.2 only.*

CF^{B}_{i+1} **Surviving bond next cash flow**. Coupon or coupon-plus-redemption paid at the next lattice time only if the issuer survives the period. *Units:* stated currency per bond at the next time. *In Chapter 5.2 only.*

$B^{\text{alive}}_{i,j}$ **Alive-state bond value**. Ex-cash-flow value of the defaultable bond at a lattice node conditional on the issuer still being alive, before the subsequent survival and default branches; scheduled cash at the node has already been paid. *Units:* stated currency per bond at the node time.

$s_{i,j}$ **Conditional node survival probability**. Pricing-model probability that the issuer survives the next lattice period, conditional on being alive at the current node; a risk-neutral input, not an unconditional real-world forecast. *Units:* probability between zero and one.

$O^{\text{KO}}_{i,j}$ **Knockout bond option value**. Alive-node value of a European bond option that is extinguished with zero option rebate by issuer default before exercise; a non-negative holder value conditional on the issuer being alive at the node. *Units:* stated currency per option at the node time.

i **Credit period index**. Selects one interval in the lesson's ordered survival schedule. The index is bookkeeping and is not itself a time or probability. *In Chapter 6.1 only.*

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

Δt **Credit small time increment**. Positive model-time interval used in the conditional definition of hazard. The hazard definition takes this increment down toward zero while holding the starting time fixed. *Units:* model-years. *In Chapter 6.1 only.*

$$\Delta t > 0$$

n **Number of credit periods**. Counts the intervals in the input survival schedule. It is a positive integer count for this finite toy model. *In Chapter 6.1 only.*

$$n \in \{1, 2, \dots\}$$

τ **Default time**. Random model time at which the reference entity first defaults. It is measured in model-years from valuation time under the stated risk-neutral model. *Units:* model-years after valuation time.

$$\tau > 0$$

λ **Constant hazard rate**. Constant conditional default intensity used by the lesson's simplified exponential survival model; a risk-neutral pricing input conditional on survival to the current instant, not a cumulative probability. *Units:* decimal intensity per model-year.

Δq_i **Interval default probability**. Probability assigned by the model to first default during one stated time interval, conditional only through the input survival curve construction. *Units:* probability between zero and one.

$S(0,t)$ **Survival probability**. Probability, under the explicitly stated model measure, that no modeled default has occurred between valuation time and a stated future time. *Units:* probability between zero and one.

F **Recovery par amount**. Positive currency amount promised after survival and used as the recovery base after default. It is a toy-model input and not a market price. *Units:* stated currency at the scheduled payment time. *In Chapter 6.2 only.*

$$F > 0$$

T **Recovery payment time**. Single future time when either the survival-state or recovered-par amount is paid. Default may occur earlier, but this toy convention delays the recovered-par payment to this time. *Units:* model-years from valuation time. *In Chapter 6.2 only.*

$$T > 0$$

LGD **Loss given default**. Fraction of an explicitly stated reference amount not recovered under a deterministic recovery convention, relative to the same reference amount used by the recovery rate. *Units:* decimal fraction between zero and one.

PV_0^{RoP} **Recovery of par present value**. Present value of one maturity payment that is par after survival and a fixed fraction of par after earlier default; a positive asset value to the holder in the one-period recovery-of-par-paid-at-maturity model. *Units:* stated currency at valuation time.

R **Recovery rate**. Fraction of a stated reference amount recovered after a modeled default under an explicitly stated recovery convention; a non-negative fraction whose reference amount, payment timing, and settlement convention are set by the model that uses it. *Units:* decimal fraction between zero and one.

A_0^{accr} **CDS accrued premium annuity**. Present-value factor for exact premium accrual paid at modeled default. It integrates accrued year fraction, discounting, and risk-neutral default density over every possible default time. *Units:* model-years. *In Chapter 7.1 only.*

$a_i(u)$ **CDS accrued year fraction at default**. Year fraction accrued from the preceding premium date to a possible default time in period i. In the regular model-year schedule it equals elapsed model time since the period start. *Units:* model-years of premium accrual. *In Chapter 7.1 only.*

$$a_i(u) = u - t_{i-1}$$

κ **CDS combined decay rate**. Abbreviation for the sum of the deterministic risk-free rate and constant hazard rate. It shortens the closed forms and has no separate market interpretation. *Units:* decimal rate per model-year. *In Chapter 7.1 only.*

$$\kappa = r + \lambda$$

G_i **CDS default event**. Event that modeled default occurs after one premium date and no later than the next. Together with survival through maturity, these events are disjoint and exhaustive in the one-default model. *In Chapter 7.1 only.*

$$G_i = \{t_{i-1} < \tau \leq t_i\}$$

u **CDS default integration time**. Integration variable ranging over every possible default time inside one CDS period, or over the interval used to cumulate a time-varying hazard rate into a survival probability. It is a deterministic dummy variable; the random default time is a different symbol. *Units:* model-years after valuation time. *In Chapter 7.1 only.*

Y **CDS event payoff**. Random discounted payoff used to connect event partitions to a CDS-leg expectation. Its currency sign and perspective must be stated before taking an expectation. *Units:* stated currency at valuation time. *In Chapter 7.1 only.*

y_i **CDS event payoff value**. Constant value taken by the generic payoff on one event in a finite partition. If the payoff varies inside an event, conditional expectation replaces this constant. *Units:* stated currency at valuation time. *In Chapter 7.1 only.*

c_i **CDS extinguishing swap cash flow**. Deterministic signed net cash flow scheduled at one payment date for one named perspective, in the extinguishing-swap construction that generalizes the premium leg to an arbitrary schedule. Positive means received, under the signed cash-flow convention used throughout this course. *Units:* stated currency at the scheduled payment time. *In Chapter 7.1 only.*

PV_0^{ext} **CDS extinguishing swap present value**. Present value, for the named perspective, of a scheduled net cash-flow stream that stops entirely once the reference entity defaults. This is the extinguishing-swap construction generalizing the premium leg. *Units:* stated currency at valuation time. *In Chapter 7.1 only.*

$\lambda(u)$ **CDS general hazard rate**. General risk-neutral default intensity, written as a function of time rather than assumed constant. It generalizes the constant hazard rate used elsewhere in this course to a deterministic function of time; the constant case is recovered by setting it equal to one fixed value for every time. *Units:* decimal intensity per model-year. *In Chapter 7.1 only.*

i **CDS period index**. Selects one default interval and its following scheduled premium date. The index is bookkeeping and is not itself a time, probability, or currency amount. *In Chapter 7.1 only.*

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

B_0^{prot} **CDS protection factor**. Present value of protection per unit CDS notional in the simplified model. It includes loss given default, discounting, and the risk-neutral default-time density. *Units:* present currency-units per notional currency-unit. *In Chapter 7.1 only.*

A_0^{sched} **CDS scheduled premium annuity**. Present-value factor for scheduled premiums per unit notional and unit annual running rate. It includes a scheduled payment only when the name survives to that payment date. *Units:* model-years. *In Chapter 7.1 only.*

r **Continuously compounded risk-free rate**. Deterministic annual decimal rate used to discount this lesson's model cash flows. It is a pricing input and is separate from the hazard rate. *Units:* decimal rate per model-year. *In Chapter 7.1 only.*

$$D(0, u) = \exp(-ru)$$

n **Number of CDS periods**. Counts the finite premium and default intervals in the simplified CDS schedule. It is a positive integer determined by model term and payment frequency. *In Chapter 7.1 only.*

$$n \in \{1, 2, ...\}$$

α_i **CDS accrual year fraction**. Input year fraction that converts an annualized spread into the premium amount for one scheduled period; a positive model input. *Units:* model-years of premium accrual.

s **CDS contract spread**. Annualized premium rate applied to notional and each stated accrual year fraction; a positive rate paid by the protection buyer in the simplified lesson. *Units:* decimal per year in calculations; basis points per year when explicitly quoted.

N **CDS notional**. Reference currency amount that scales the simplified premium and protection legs; a positive amount, not itself a signed leg cash flow. *Units:* stated currency.

$s^\star$ **CDS par spread**. Contractual spread that makes the two positive leg magnitudes equal at valuation time with zero upfront amount, solved under the lesson's input curves, recovery, timing, and accrued-premium convention. *Units:* decimal per year in calculations; basis points per year when explicitly quoted.

A_0^{prem} **CDS premium annuity**. Present-value coefficient that multiplies contractual spread and notional in the simplified premium leg; positive, and includes scheduled premiums and exact accrued premium under the lesson's default-time model. *Units:* model-years of present value per unit notional.

PV_0^{prem} **CDS premium leg present value**. Positive valuation-time magnitude of the simplified protection buyer's premium payments; its signed contribution to protection-buyer net value is negative. *Units:* stated currency at valuation time.

PV_0^{buyer} **CDS protection buyer net present value**. Signed protection-buyer value equal to protection-leg magnitude minus premium-leg magnitude; positive favors the protection buyer and negative favors the protection seller in the two-leg toy model. *Units:* stated currency at valuation time.

PV_0^{prot} **CDS protection leg present value**. Positive valuation-time magnitude of the simplified loss-given-default payment received by the protection buyer after a modeled default. *Units:* stated currency at valuation time.

s_{MSQ} **CDS market-standard quote**. Conventional spread that is the input or the output of the lesson's simplified converter; it equals the zero-upfront par spread within the converter's implied flat-hazard model and need not equal the contract's fixed running coupon. *Units:* decimal per year in calculations; basis points per year when explicitly quoted.

c_{std} **CDS standard coupon**. Fixed annualized rate used to determine the contract's running premium cash flows in the lesson's standard-coupon model; a positive rate paid by the protection buyer on surviving notional and as accrued premium after default. *Units:* decimal per year in calculations; basis points per year when explicitly quoted.

U_0 **CDS upfront amount**. Time-zero cash amount that balances protection and fixed-coupon premium value under the lesson's pricing convention; positive means paid by the protection buyer and negative means received by the protection buyer. *Units:* stated currency at valuation time.

Π^{buyer} **CDS buyer discounted cash flows**. Random sum of the protection buyer discounted signed cash flows over the contract life before any upfront: protection received on default less scheduled and accrued premium paid, each discounted from its own payment time to valuation time. *Units:* stated currency at valuation time. *In Chapter 7.3 only.*

λ_k^{flat} **CDS flat tenor hazard rate**. Single constant hazard rate that the standard converter implies for one tenor on its own when it transforms that tenor market-standard quote into an

upfront. Different tenors give different flat rates that do not form one term structure. *Units:* decimal intensity per model-year. *In Chapter 7.3 only.*

λ_k **CDS segment hazard rate**. Constant risk-neutral hazard rate on the fitted curve segment between two consecutive tenor maturities. It is one piece of a piecewise-constant term structure, not a separate flat model for that tenor. *Units:* decimal intensity per model-year. *In Chapter 7.3 only.*

$$\lambda_k \geq 0$$

K **CDS tenor count**. Counts the quoted tenors on the credit curve. It is a positive integer and does not measure time, money, or probability. *In Chapter 7.3 only.*

$$K \in \{1, 2, ...\}$$

k **CDS tenor index**. Selects one quoted tenor of the credit curve in increasing maturity order. The index is bookkeeping and is not itself a time, a quote, or a probability. *In Chapter 7.3 only.*

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

T_k **CDS tenor maturity**. Maturity of the standard contract quoted at one tenor, measured in model-years from valuation time. Two consecutive tenor maturities bound one hazard segment of the fitted curve. *Units:* model-years from valuation time. *In Chapter 7.3 only.*

s_k **CDS tenor quote**. Market-standard quote marked at one tenor: the conventional spread that the standard converter turns into that tenor traded upfront using a flat hazard rate for that tenor alone. *Units:* decimal per year in calculations; basis points per year when explicitly quoted. *In Chapter 7.3 only.*

u_k **CDS tenor upfront fraction**. Signed upfront fraction of notional exchanged at valuation time for the standard contract at one tenor; positive means paid by the protection buyer. It is the traded price that the tenor market-standard quote stands for. *Units:* fraction of notional at valuation time. *In Chapter 7.3 only.*

$$u_k = U_0/N$$

N_{eq} **CDS equivalent notional**. Signed notional of the liquid contract, bought as protection, whose quote sensitivity cancels the position sensitivity; a negative value means protection is sold in the liquid tenor. *Units:* stated currency. *In Chapter 7.4 only.*

ρ_{eq} **CDS equivalent ratio**. Absolute equivalent notional divided by position notional: how much of the most liquid contract offsets one unit of the position quote risk. It is positive; the side is stated separately. *In Chapter 7.4 only.*

$$\rho_{\text{eq}} = |N_{\text{eq}}|/N_{\text{pos}}$$

q **CDS liquid quotation**. Any one quotation of the liquid contract viewed as a function of the flat hazard rate: its market-standard quote, its upfront fraction, or its par spread. Each is a strictly increasing function of the hazard rate in this model. *In Chapter 7.4 only.*

T_L **CDS liquid tenor maturity**. Maturity of the most liquid standard contract on the reference entity, whose market-standard quote gives the single market level of the flat-hazard model. *Units:* model-years from valuation time. *In Chapter 7.4 only.*

$\mathrm{DV01}^{L}$ **CDS liquid unit quote sensitivity**. Change in the value of one unit of notional of bought liquid-tenor protection when the liquid quote rises by the bump. It is positive, because bought protection gains when the quote rises. *Units:* fraction of notional per basis point. *In Chapter 7.4 only.*

v_L **CDS liquid unit value**. Value at valuation time of one currency unit of notional of bought protection in the liquid contract, before upfront. It is a fraction of notional. *Units:* fraction of notional at valuation time. *In Chapter 7.4 only.*

N_{pos} **CDS position notional**. Positive notional of the position being hedged, in currency units. It carries no sign; the position side does. *Units:* stated currency. *In Chapter 7.4 only.*

$\mathrm{DV01}^{\mathrm{pos}}$ **CDS position quote sensitivity**. Change in the signed position value when the liquid tenor market-standard quote rises by the bump, with every contract revalued at the re-solved flat hazard rate. Market usage calls this a risky DV01 or spread DV01. *Units:* stated currency per basis point. *In Chapter 7.4 only.*

σ **CDS position side**. Sign of the position side: plus one for bought protection and minus one for sold protection. It converts protection-buyer value into value for the position holder. *In Chapter 7.4 only.*

$$\sigma \in \{+1, -1\}$$

V^{pos} **CDS position value**. Signed value of the position for its own side at valuation time, before any upfront exchanged at inception: the protection-buyer value of the contract multiplied by the side sign. *Units:* stated currency at valuation time. *In Chapter 7.4 only.*

h **CDS quote bump**. Size of the bump applied to the liquid tenor quote in the bump-and-reprice sensitivity. The lesson uses one basis point of market-standard quote. *Units:* decimal rate per year. *In Chapter 7.4 only.*

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.