

Other Topics in Quantitative Finance

These notes introduce other important topics in finance. In particular, we will briefly discuss important aspects of (i) corporate finance (ii) securitization (iii) risk management and (iv) algorithmic trading.

Corporate finance is an enormous sub-field of finance with many interesting associated problems including capital budgeting, determining the optimal capital structure of the firm and determining the optimal dividend policy of the firm. We will briefly discuss the famous Modigliani-Miller (MM) theorems which address the latter two questions in an idealized economy. We will also discuss the very important incentive / principal-agent problems that arise throughout corporate finance. Our introduction to securitization will be via a simple one-period collateralized-debt-obligation (CDO). We will also briefly discuss mortgage-backed securities (MBS). Risk management is a key task in any financial company and we will discuss some of the main approaches to assessing risk. These approaches include the calculation of risk measures such as VaR and CVaR as well as scenario analysis. Finally we will end with an overview of algorithmic trading, an aspect of finance that has become increasingly important in recent years. Algorithmic trading is of interest because it relaxes the assumptions that economic agents are price-takers and assumes instead that their trading moves security prices adversely. Accounting for these adverse price moves is vital for both high-frequency trading as well as lower-frequency trading where large trades need to be executed in a relatively short period of time.

1 An Introduction to Corporate Finance

We begin this section with some examples of incentive problems that can arise between the equity-holders and debt-holders of a firm. These problems are very important in practice.

1.1 Incentive Problems in Corporate Finance

Recall Merton's structural lattice model of default that we saw earlier when we discussed credit modeling. Our example there assumes the following parameters: $V_0 = 1,000$, $T = 7$ years, $\mu = 15\%$, $\sigma = 25\%$, $r = 5\%$ and the # of time periods $= 7$. The face value of the debt was 800 and the coupon on the debt was zero. The lattice of firm values is displayed below with asterisk next to those values at time T corresponding to firm default.

```
Firm Price Lattice
                                                          6940.6
                                                  5262.6  3990.2
                                          3990.2  3025.5  2294.0
                                  3025.5  2294.0  1739.4  1318.9
                          2294.0  1739.4  1318.9  1000.0  758.2*
                  1739.4  1318.9  1000.0  758.2   574.9   435.9*
          1318.9  1000.0  758.2   574.9   435.9   330.5   250.6*
1000.0    758.2   574.9   435.9   330.5   250.6   190.0   144.1*

t = 0     t = 1   t = 2   t = 3   t = 4   t = 5   t = 6   t = 7
```

We used this model to price the equity and debt of the firm and obtained values of 499.7 and 500.3, respectively. The lattices are displayed below.

Equity Lattice

							6140.6
						4501.6	3190.2
					3266.4	2264.5	1494.0
				2336.9	1570.2	978.4	518.9
			1640.8	1054.7	603.6	258.0	0.0
		1127.8	687.1	358.4	128.3	0.0	0.0
	758.6	435.7	207.1	63.8	0.0	0.0	0.0
499.7	269.9	117.4	31.7	0.0	0.0	0.0	0.0
t = 0	t = 1	t = 2	t = 3	t = 4	t = 5	t = 6	t = 7

Debt Lattice

							800.0
						761.0	800.0
					723.9	761.0	800.0
				688.6	723.9	761.0	800.0
			653.2	684.7	715.3	742.0	758.2*
		611.6	631.7	641.6	630.0	574.9	435.9*
	560.3	564.3	551.1	511.1	435.9	330.5	250.6*
500.3	488.3	457.5	404.2	330.5	250.6	190.0	144.1*
t = 0	t = 1	t = 2	t = 3	t = 4	t = 5	t = 6	t = 7

Turning Down Good Investments

Suppose now the firm is offered a great(!) investment opportunity. The fair value of the investment is $X = 100$ but the cost to the firm will only be 90 which is substantially less than X. The firm has no cash currently available, however, and would have to raise the cash, i.e. 90 dollars, from the current equity holders.

Question: Will the equity holders invest?

Answer: This is clearly an excellent deal since X is the *fair* risk-neutral value of the deal and yet it is available for only 90. We can model this situation by first adding X to the initial value of the firm and computing the resulting firm-value lattice. We obtain (with the time T default values again market with an asterisk):

Firm Value Lattice

							7634.7
						5788.8	4389.3
					4389.3	3328.1	2523.4
				3328.1	2523.4	1913.3	1450.7
			2523.4	1913.3	1450.7	1100.0	834.1
		1913.3	1450.7	1100.0	834.1	632.4	479.5*
	1450.7	1100.0	834.1	632.4	479.5	363.6	275.7*
1100.0	834.1	632.4	479.5	363.6	275.7	209.0	158.5*
t = 0	t = 1	t = 2	t = 3	t = 4	t = 5	t = 6	t = 7

We can then compute the equity lattice in the usual manner:

```
Equity Lattice
                                                              6834.7
                                                     5027.8   3589.3
                                            3665.4   2567.1   1723.4
                                   2639.5   1799.6   1152.4    650.7
                          1868.4   1224.8   726.9    339.0      34.1
                 1296.5   809.4    441.4    176.2     16.9       0.0
        881.1    520.9    261.0     91.5      8.4      0.0       0.0
586.9   327.7    151.3     47.4      4.2      0.0      0.0       0.0

t = 0   t = 1    t = 2    t = 3    t = 4    t = 5    t = 6     t = 7
```

Note that that equity value is now 586.9 which is an increase of $586.9 - 499.7 = 87.2$ dollars. Unfortunately, this increase is less than the 90 required to make the investment and so the equity-holders, if rational, would *not* make the investment! Of course what has happened is that some of the value has gone to the debt-holders (whose lattice we did not display). Their situation has improved since the injection of new equity has made default less likely than before.

This problem (of not making investments that make good economic sense from the overall firm's perspective) is known as the **debt-overhang** problem and can be very pronounced for firms that are relatively close to default.

Exercise 1 *(a) Suppose instead the current equity-holders raised the cash by issuing \$90 of new* non-voting *equity to* outside *investor. These new equity holders would own* $90/(90 + 499/7) = 15.26\%$ *of the firm. Would the original equity holders now vote to make the investment?*

(b) Is there anything wrong with the scenario in part (a)?

(c) What demands would rational outside investors make in order for them to inject \$90 of new equity into the firm so that the good investment could be made? Would the current equity holders agree to those demands?

(d) What would have to happen in order for the equity holders to make the new investment?

Taking on Bad Investments

In addition to the debt overhang problem, there are other incentive problems that can arise. Suppose for example, the fair value of the investment is again 100 but that it costs 110. Clearly from the perspective of the overall value of the firm this is a bad investment and should not be made. And yet from the point of view of the equity-holders it may actually be rational to do this if the investment increases the volatility of the firm. This makes sense if we recall that the equity-holders own a call option on the value of the firm and that the value of an option increases with volatility. It is possible that the increase in equity value due to the increase in volatility will exceed the decrease in equity value due to the poor quality of the investment. In that case it makes sense for the equity holders (who control the firm) to make the investment.

Exercise 2 *How might you model this situation using Merton's structural model?*

Other Incentive Problems in Corporate Finance

Incentive problems such as those outlined above can be and often are mitigated in part by the presence of *debt covenants*. Covenants are contracts that specify certain performance criteria that must be met. They often specify a maximum debt-to-asset ratio, for example, or restrict the size of dividend payments relative to working capital etc. Failure to satisfy a covenant will generally make the debt come due immediately. Nonetheless, incentive problems remain and can be found throughout corporate finance.

In the world of banking, for example, it is widely believed that governments will bail out any "too big-to-fail" institutions that are in financial distress. Governments, i.e. tax-payers, therefore provide an implicit subsidy to these institutions as it lowers their cost of capital. Moreover the size of this subsidy increases as the debt-to-equity ratio increases. This as well as the debt-tax-shield (see below) encourages[1] banks to have very high leverage ratios which of course makes them much more likely to fail.

Incentive problems can also arise between the management and the equity-holders of a firm. Corporate governance laws are such that it can be very difficult for the equity holders (who actually own the firm) to exert control over the management of the firm. This has often been blamed for the very high levels of executive compensation in the US where there is often very little relationship between compensation and performance.

From a modeling perspective, these incentive problems are typically handled using game theory / contract theory and other tools from micro-economics.

1.2 Optimal Capital Structure and Dividend Policy

One of the most important problems in corporate finance is that of determining the optimal capital structure of a firm, i.e. determining the optimal combination of debt and equity that the firm uses to fund its operations and investments. A related problem is that of determining the optimal dividend / payout policy of the firm. In the 1950's Modigliani and Miller[2] (MM) provided answers to these questions in an *idealized* world.

Modigliani Miller Theorems

Consider an idealized world where there are no market frictions such as taxes, asymmetric information, bankruptcy costs etc. Then MM showed the value of the firm was independent of the capital structure. If we consider the fundamental equation of accounting

$$V = D + E \tag{1}$$

then the MM result states that the left-hand-side of (1) does not depend on the debt-equity split. This also implies that the investment decisions of the firm can be separated entirely from the financing decisions of the firm. The proof of this result is very straightforward and simply amounts to making the following argument. Suppose a firm is all-equity financed but that an investor would prefer to invest in a firm with identical cash-flows, i.e. assets, but that instead was partly financed with debt. The investor can achieve this desired investment by doing the borrowing herself rather than relying on the firm to do so. In particular, she can partly fund her investment in the all-equity financed firm by borrowing. The resulting cash-flows obtained by the investor will be identical to the cash-flow from a position in the firm partly financed with debt. It is then clear that the value of the firm should not depend on how it finances itself.

Similar simple arguments also imply that investors are indifferent to the dividend policy of the firm. Likewise a firm does not need to engage in risk management as individual investors can do so themselves. Of course these results only imply in an idealized world that never holds in practice. So why do people care about these MM results? The answer is that these results help to clarify our thinking and understand what real-world market frictions imply for the optimal capital structure. For example, a very important and somewhat controversial feature of the U.S tax code is that debt interest payments are paid *pre* (corporate) tax whereas dividend payments are made *after* tax. This has the effect of making debt more attractive than equity and this favorable treatment[3] of debt can encourage firms, e.g. banks(!), to be very highly leveraged. In order to account for this favorable tax treatment it is common to rewrite (1) as

$$V_L = V + \text{PV(Debt tax shield)} \tag{2}$$

where V is the value of the firm in a no-tax world and PV(Debt tax shield) is simply the present-value of the benefits to the firm of not having to pay tax on interest payments.

Other market frictions include bankruptcy costs. These costs are also significant and in addition to the direct costs, e.g. legal and administrative costs of bankruptcy, they also include indirect costs such as loss of business

[1]Bankers have therefore lobbied very hard against regulations that require them to hold much higher levels of equity. See the recent book *"The Bankers' New Clothes"* by Admati and Hellwig for a discussion and compelling arguments as to why banks should be required to hold much higher levels of equity.

[2]Both of whom were Nobel Prize winners.

[3]This reason alone explains why many commentators have called for an end to this favorable treatment of debt.

etc. resulting from bankruptcy or the fear of bankruptcy. In contrast to taxes, this friction encourages firms to hold more equity and therefore to be less leveraged. It is worth noting that in contrast to the Merton and Black-Cox models for pricing corporate debt, later models due to Leland, Leland and Toft etc, included both taxes and bankruptcy costs in evaluating corporate debt.

1.3 Capital Budgeting

Capital budgeting is one of the key problems in corporate finance. Simply put, capital budgeting is the process associated with figuring out what investments a firm should make and there are three key aspects to this: (i) determining the fair value of each potential investment (ii) determining what subset of these investments the firm should make and (iii) determining how to fund the chosen subset of investments. In an idealized MM world step we could argue that steps (ii) and (iii) are irrelevant and can therefore be ignored. Of course we do not live in an MM world and therefore steps (ii) and (iii) are often important. For example, it may be the case that raising outside funds is too expensive for a firm in which case it may only be able to invest in a subset of the investments that satisfy some budget constraint.

Step (i) is always necessary and there are many approaches. The martingale approach we used to price real options in an earlier set of lecture notes is certainly valid. In corporate finance it is common to focus on the appropriate choice of discount factor rather than an EMM, Q, and we will follow that line here. While the theory states that a stochastic discount factor is appropriate it is common to instead apply a deterministic discount factor. But what one? The risk-free rate is almost always inappropriate as it does not account for the riskiness of the cash-flows. It is common to use the *weighted average cost of capital* or WACC. The WACC is the expected return on a portfolio of all the companies securities and is given (in a MM world with no taxes) by

$$r \;=\; r_D \frac{D}{V} + r_E \frac{E}{V} \tag{3}$$

where r_D and r_E are the expected returns to the debt and equity, respectively. In a world with taxes where the tax regime favors debt, the WACC is adjusted accordingly

$$r^* \;=\; r_D(1-T_c)\frac{D}{V} + r_E \frac{E}{V} \tag{4}$$

where T_c is the marginal corporate tax rate. In practice we cannot observe the WACC, r. Instead we observe r_D and estimate r_E and then use (4) to estimate r^*. It is important to note that WACC is an (approximately) appropriate discount factor only for projects / investments that are similar to the overall business of the firm. In particular, investments that are "riskier" (less risky) than the firm's overall business should be valued with a higher (lower) discount factor than the WACC.

---

2 Securitization

Securitization is the name given to the process of constructing new securities from the cash-flows generated by a pool of underlying securities. The main economic rationale behind securitization is that it enables the construction of new securities with a broad range of risk profiles. Different types of investors may therefore be interested in these new securities even if they had no interest in the underlying securities. If this is the case then there will be an increased demand for the underlying cash-flows and so the cost-of-capital is reduced for the issuers of the underlying securities.

Collateralized debt obligations (CDOs) are a particular class of securities that are structured this way. CDOs are constructed from an underlying pool of fixed-income securities and they were first issued by banks in the mid-1990s. In contrast to the motivation given above, the original motivation for introducing CDO's was *regulatory arbitrage*. By keeping the equity tranche (see below) of the CDO a bank could effectively keep the entire economic risk of the underlying portfolio but because the notional principal of the tranch was much smaller, the bank faced must lighter capital requirements.

2.1 Collateralized Debt Obligations (CDO's)

We want to find the expected losses in a simple 1-period CDO with the following[4] characteristics:

- The maturity is 1 year.
- There are $N = 125$ bonds in the reference portfolio.
- Each bond pays a coupon of one unit after 1 year if it has not defaulted.
- The recovery rate on each defaulted bond is zero.
- There are 3 tranches of interest: the equity, mezzanine and senior tranches with attachment points 0-3 defaults, 4-6 defaults and 7-125 defaults, respectively.

We make the simple assumption that the probability, q, of defaulting within 1 year is identical across all bonds X_i is the normalized asset value of the i^{th} credit, i.e. bond, and we assume

$$X_i = \sqrt{\rho}M + \sqrt{1-\rho}\,Z_i \tag{5}$$

where $M, Z_1, \ldots, Z_N$ are IID normal random variables. Note that the correlation between each pair of asset values is identical. We assume also that the i^{th} credit defaults if $X_i \le \bar{x}_i$. Since the probability of default, q, is identical across all bonds we must therefore have

$$\bar{x}_1 = \cdots \bar{x}_N = \Phi^{-1}(q). \tag{6}$$

It now follows from (5) and (6) that

$$\begin{aligned}
\mathsf{P}(\text{Credit } i \text{ defaults} \,|\, M) &= \mathsf{P}(X_i \le \bar{x}_i \,|\, M) \\
&= \mathsf{P}(\sqrt{\rho}M + \sqrt{1-\rho}\,Z_i \le \Phi^{-1}(q) \,|\, M) \\
&= \mathsf{P}\left(Z_i \le \frac{\Phi^{-1}(q) - \sqrt{\rho}M}{\sqrt{1-\rho}} \;\middle|\; M \right).
\end{aligned}$$

Therefore conditional on M, the total number of defaults is Binomial(N, q_M) where

$$q_M := \Phi\left(\frac{\Phi^{-1}(q) - \sqrt{\rho}M}{\sqrt{1-\rho}} \right).$$

That is,

$$p(k \,|\, M) = \binom{N}{k} q_M^k (1-q_M)^{N-k}.$$

The unconditional probabilities can be computed by integrating numerically the binomial probabilities with respect to M so that

$$\mathsf{P}(k \text{ defaults}) = \int_{-\infty}^{\infty} p(k \,|\, M)\phi(M)\,dM$$

where $\phi(\cdot)$ is the standard normal PDF. We can now compute the expected (risk-neutral) loss on each of the three tranches according to

$$\begin{aligned}
\mathrm{E}_0^Q\,[\text{Equity tranche loss}] &= 3 \times \mathsf{P}(3 \text{ or more defaults}) + \sum_{k=1}^{2} k\,\mathsf{P}(k \text{ defaults}) \\
\mathrm{E}_0^Q\,[\text{Mezzanine tranche loss}] &= 3 \times \mathsf{P}(6 \text{ or more defaults}) + \sum_{k=1}^{2} k\,\mathsf{P}(k+3 \text{ defaults}) \\
\mathrm{E}_0^Q\,[\text{Senior tranche loss}] &= \sum_{k=1}^{119} k\,\mathsf{P}(k+6 \text{ defaults}).
\end{aligned}$$

Figure 1: Expected Tranche Losses As a Function of q and ρ

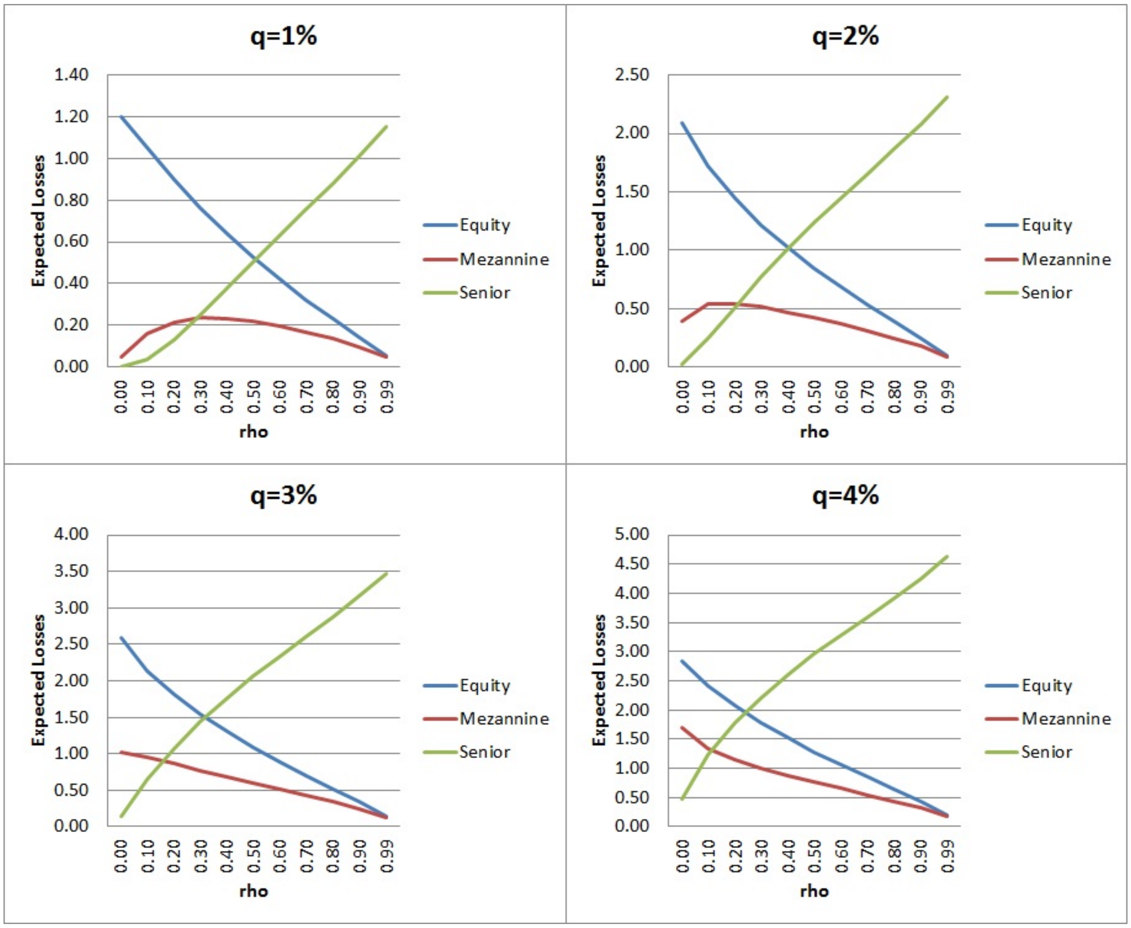

Results for various values of ρ and q are displayed in the figure below. Regardless of the individual default probability, q, and correlation, ρ, we see

$$\mathrm{E}_0^Q\left[\%\ \text{Equity tranche loss}\right] \geq \mathrm{E}_0^Q\left[\%\ \text{Mezzanine tranche loss}\right] \geq \mathrm{E}_0^Q\left[\%\ \text{Senior tranche loss}\right].$$

We also note that the expected losses in the equity tranche are always decreasing in ρ while mezzanine tranches are often relatively insensitive[5] to ρ. The expected losses in senior tranches (with upper attachment point of 100% or 125 units in our example) are always increasing in ρ.

Exercise 3 *How does the total expected loss in the portfolio vary with ρ?*

Remark 1 *The* dependence structure *we used in (5) to link the default events of the various bonds is the famous Gaussian-copula model.*

Remark 2 *In practice CDO's are multi-period securities and many CDO's are actually synthetic CDO's where the underlying portfolio is a synthetic portfolio of bonds. Synthetic CDO tranches trade like a CDS where one party sells protection on a tranche and another party buys protection.*

[4]This example is taken from "The Devil is in the Tails: Actuarial Mathematics and the Subprime Mortgage Crisis", by C. Donnelly and P. Embrechts in ASTIN Bulletin 40(1), 1-33.

[5]This has important implications when it comes to model calibration, an issue we will not pursue further here.

2.2 Mortgage-Backed Securities

We briefly provide another example of how securitization can work by considering a standard level-payment mortgage. We assume the initial mortgage principal is $M_0 := M$ and that equal periodic payments of size B dollars are made. After n such payments the mortgage principal and interest will have all been paid in full. Each payment, B, therefore[6] pays both interest and some of the principal. If we assume the coupon rate is c per period then we can solve for B. In particular, let M_k denote the mortgage principal remaining after the k^{th} period. Then

$$M_k \;=\; (1+c)M_{k-1} - B \quad \text{for} \;\; k = 1, 2, \ldots, n \tag{7}$$

with $M_n = 0$. We can iterate (7) to obtain

$$\begin{aligned} M_k &= (1+c)^k M_0 - B \sum_{p=0}^{k-1} (1+c)^p \\ &= (1+c)^k M_0 - B \left[\frac{(1+c)^k - 1}{c} \right]. \end{aligned} \tag{8}$$

But $M_n = 0$ and so we obtain

$$B = \frac{c(1+c)^n M_0}{(1+c)^n - 1}. \tag{9}$$

We can now substitute (9) back into (8) and obtain

$$M_k \;=\; M_0 \, \frac{(1+c)^n - (1+c)^k}{(1+c)^n - 1}.$$

Suppose now that we wish to compute the present value of the mortgage assuming a deterministic world with no possibility of defaults or prepayments. Then assuming a risk-free interest rate of r per period, we obtain that the fair mortgage value as

$$\begin{aligned} F_0 &= \sum_{k=1}^{n} \frac{B}{(1+r)^k} \\ &= \frac{c(1+c)^n M_0}{(1+c)^n - 1} \times \frac{(1+r)^n - 1}{r(1+r)^n}. \end{aligned} \tag{10}$$

Note that if $r = c$ then (10) immediately implies that $F_0 = M_0$, as expected. In general, however, we will have $r < c$, to account for the possibility of default, prepayment, servicing fees, bank profits etc.

Scheduled Principal and Interest Payments

Since we know M_{k-1} we can compute the interest

$$I_k \;:=\; cM_{k-1} \tag{11}$$

on M_{k-1} that would be due in the next period, i.e. period k. This also means we can interpret the k^{th} payment as paying

$$P_k :=\; B - cM_{k-1} \tag{12}$$

of the remaining principal, M_{k-1}. In any time period, k, we can therefore easily break down the payment B into a scheduled principal payment, P_k, and a scheduled interest payment, I_k. Indeed we can take a large pool of these mortgages and assign all the interest payments (given by (11)) to an *interest-only* (IO) mortgage-backed security (MBS), and all the principal payments (given by (12)) to a *principal-only* (PO) MBS. There are many other classes of MBS including for example sequential CMO's (collateralized mortgage obligations), PAC CMO's etc.

[6] Such a mortgage is said to be fully amortizing.

Prepayment Risk

It is important to note, however, that there is in fact uncertainty in the interest and principal payments. This is due to possible default by the mortgage holder as well as the possibility of *prepayments*. Many mortgage-holders in the US are allowed to pre-pay the mortgage principal earlier than scheduled. In fact payments in excess of the *scheduled* principal payments are called prepayments. There are many possible reasons for prepayments including default (in which case an agency-insured mortgage will prepay the entire mortgage), opportunities to refinance the mortgage at a better rate, requirements to prepay when selling the home underlying the mortgage etc. Prepayment modeling is therefore an important feature of pricing MBS. Indeed the value of some MBS's such as IO and PO MBS are very sensitive to prepayment modeling and assumptions.

Exercise 4 *What do you think happens to the value of a PO / IO security when interest rates (i) increase (ii) decrease?*

Pricing Mortgage-Backed Securities

MBS are fixed-income securities with considerable exposure to interest-rates and (why?) interest-rate volatility. To price them it is therefore necessary to use a term-structure model that has been properly calibrated to an appropriate set of liquid fixed-income derivatives prices as well (of course) as the term structure of interest rates. The calibrated term structure model is then combined with a prepayment model (and possibly a mortgage default model) to construct (typically via Monte-Carlo simulation) MBS prices and associated risk numbers.

2.3 Other Asset-Backed Securities

Securities that are constructed by reassigning the cash-flows of some underlying group of securities are generally referred to as *asset-backed securities* (ABS). CDOs and MBSs are examples of these securities. Other common forms of ABS include collateralized loan obligations (CLOs) and securities constructed from pools of credit card loans, students loans or car loans. In fact it is possible to form more complex ABS from an underlying pool of other ABS. This leads, for example, to ABS CDOs where the CDO is constructed from an underlying pool of ABS, or CDO-squared's where a CDO is constructed from an underlying pool of CDO tranches. The structured credit / ABS market was at the heart of the 2008 financial crisis and came in for considerable criticism. As a result, many of the more complex ABS are no longer actively traded.

---

3 Risk Management

Risk management is an extremely important topic and we can only briefly describe here some of the main approaches.

3.1 Factor Sensitivity Measures or the "Greeks"

A factor sensitivity measure gives the change in the value of the portfolio for a given change in the "risk factors". Commonly used examples include the Greeks of an option portfolio or the duration and convexity of a bond portfolio. These measures are often used to set position limits on trading desks and portfolios. Consider for example an option that is written on an underlying security with price process S_t. We assume the time t price, C, is a function of only S_t and the implied volatility, σ_t. Then a simple application of Taylor's Theorem yields

$$\begin{aligned} C(S+\Delta S, \sigma+\Delta\sigma) &\approx C(S,\sigma) + \Delta S \frac{\partial C}{\partial S} + \frac{1}{2}(\Delta S)^2 \frac{\partial^2 C}{\partial S^2} + \Delta\sigma \frac{\partial C}{\partial \sigma} \\ &= C(S,\sigma) \ + \ \Delta S\, \delta \ + \ \frac{1}{2}(\Delta S)^2\, \Gamma \ + \ \Delta\sigma \text{ vega} \end{aligned}$$

where we have omitted the dependence of the various quantities on t. We therefore obtain

$$\begin{aligned} \text{P\&L} &\approx \delta \Delta S + \frac{\Gamma}{2}(\Delta S)^2 + \text{vega } \Delta\sigma \\ &= \text{delta P\&L} + \text{gamma P\&L} + \text{vega P\&L} . \end{aligned} \tag{13}$$

When $\Delta\sigma = 0$, we obtain obtain the well-known **delta-gamma approximation** which is often used, for example, in historical Value-at-Risk (VaR) calculations. Note that we can also write (13)

$$\begin{aligned} \text{P\&L} &\approx \delta S \left(\frac{\Delta S}{S}\right) + \frac{\Gamma S^2}{2}\left(\frac{\Delta S}{S}\right)^2 + \text{vega } \Delta\sigma \\ &= \text{ESP} \times \text{Return} + \$ \text{ Gamma} \times \text{Return}^2 + \text{vega } \Delta\sigma \end{aligned} \tag{14}$$

where **ESP** denotes the **equivalent stock position** or **"dollar" delta**. Note that it is easy to extend this calculation to a portfolio of options on the same underlying security. It is also straightforward to extend these ideas to derivatives portfolios written with many different underlying securities.

Depending on the particular asset class, investors / traders / risk managers should always know their exposure to the Greeks, i.e. dollar delta, dollar gamma and vega etc. It is also very important to note that approximations such as (14) are **local** approximations as they are based (via Taylor's Theorem) on "small" moves in the risk factors. These approximations can and indeed do break down in violent markets where changes in the risk factors can be very large.

3.2 Scenario Analysis

The scenario approach defines a number of scenarios where in each scenario the various risk factors are assumed to have moved by some fixed amounts. For example, a scenario might assume that all stock prices have fallen by 10% and all implied volatilities have increased by 5 percentage points. Another scenario might assume the same movements but with an additional steepening of the volatility surface. A scenario for a credit portfolio might assume that all credit spreads have increased by some fixed absolute amount, e.g. 100 basis points, or some fixed relative amount, e.g. 10%. The risk of a portfolio could then be defined as the maximum loss over all of the scenarios that were considered. A particular advantage of this approach is that it does not depend on probability distributions that are difficult to estimate.

Underlying	SPX Index								
	Underlying and Volatility Stress Table								
Sum of PnL	Vol Stress								
Underlying Stress	-10	-5	-2	-1	0	1	2	5	10
-20	13,938	11,774	10,631	10,277	9,936	9,608	9,293	8,419	7,183
-10	6,109	4,946	4,436	4,291	4,158	4,035	3,922	3,634	3,296
-5	1,831	1,652	1,637	1,643	1,654	1,670	1,689	1,766	1,946
-2	(314)	89	356	447	539	631	723	1,001	1,461
-1	(920)	(338)	15	132	248	363	478	816	1,361
0	(1,463)	(714)	(280)	(139)	0	137	273	668	1,293
1	(1,939)	(1,035)	(527)	(363)	(203)	(45)	110	559	1,259
2	(2,346)	(1,300)	(723)	(539)	(359)	(182)	(9)	489	1,258
5	(3,125)	(1,744)	(1,003)	(769)	(541)	(318)	(102)	518	1,460
10	(2,921)	(1,297)	(423)	(146)	123	385	641	1,372	2,483
20	2,344	3,559	4,272	4,506	4,738	4,967	5,194	5,860	6,919

Figure 2: P&L for an Options Portfolio on the S&P 500 under Stresses to Underlying and Implied Volatility

Figure 2 shows the P&L under various scenarios of an options portfolio with the S&P 500 as the underlying security. The vertical axis represents *percentage* shifts in the price of the underlying security, i.e. the S&P 500, whereas the horizontal axis represents *absolute* changes in the implied volatility of each option in the portfolio. For example, we see that if the S&P 500 were to fall by 20% and implied volatilities were to all rise by 5 percentage points, then the portfolio would gain 8.419 million dollars (assuming that the numbers in Figure 2 are expressed in units of $1,000$ dollars). When constructing scenario tables as in Figure 2 we can use approximations like (14) to check for internal consistency and to help identify possible bugs in the software.

While scenario tables are a valuable source of information there are many potential pit-falls associated with using them. These include:

1. *Identifying the relevant risk factors*

   While it is usually pretty clear what the main risk factors for a particular asset class are, it is quite possible that a portfolio has been constructed so that it is approximately neutral to changes in those risk factors. Such a portfolio might then only have (possibly very large) exposures to secondary risk factors. It is important then to include shifts in these secondary factors in any scenario analysis. The upshot of this observation is that the the relevant risk factors *depend* on the specific portfolio under consideration rather than just the asset class of the portfolio.

2. *Identifying "reasonable" shifts for these risk factors*

   For example, we may feel that a shift of -10% is plausible for the S&P 500 because we know from experience that such a move, while extreme, is indeed possible in a very volatile market. But how do we determine plausible shifts for less transparent risk factors? The answer typically lies in the use of statistical techniques such as PCA, extreme-value theory, time series methods, common sense(!) etc.

A key role of any risk manager then is to understand what scenarios are plausible and what scenarios are not. For example, in a crisis we would expect any drop in the price of the underlying security to be accompanied by a rise in implied volatilities. We would therefore pay considerably less attention to the numbers in the upper left quadrant of Figure 2.

3.3 Value-at-Risk

Many risk measures such as value-at-risk (VaR) are based on the loss distribution of the portfolio. Working with loss distributions makes sense as the distribution contains all the information you could possibly wish to know about possible losses. A loss distribution implicitly reflects the benefits of netting and diversification. Moreover it is easy to compare the loss distribution of a derivatives portfolio with that of a bond or credit portfolio, at least when the same time horizon is under consideration. However, it must be noted that it may be very difficult to estimate the loss distribution. This may be the case for a number of reasons including a lack of historical data, non-stationarity of risk-factors and poor model choice among others.

When we discuss the *distribution* of losses L_{t+1} (over the interval t to $t+1$) it is important to clarify exactly what we mean. In particular, we need to distinguish between the *conditional* and *unconditional* loss distributions. Consider the series $\mathbf{X_t}$ of risk factor changes and assume that they form a stationary[7] time series with stationary distribution $F_{\mathbf{X}}$. We also let $\mathcal{F}_t$ denote all information available in the system at time t, including $\{\mathbf{X_s} : s \leq t\}$ in particular. We then have the following two definitions.

Definition 1 *The* **unconditional** *loss distribution is the distribution of L_{t+1} given the time t composition of the portfolio and assuming the CDF of $\mathbf{X_{t+1}}$ is given by $F_{\mathbf{X}}$.*

Definition 2 *The* **conditional** *loss distribution is the distribution of L_{t+1} given the time t composition of the portfolio and conditional on the information in $\mathcal{F}_t$.*

For relatively short horizons, e.g. 1 day or 10 days, then the conditional loss distribution is clearly the appropriate distribution for risk management purposes. This would be particularly true in times of high market volatility when the unconditional distribution would bear little resemblance to the true conditional distribution.

[7] A time series, $\mathbf{X_t}$, is *strongly stationary* if $(\mathbf{X_{t_1}}, \ldots, \mathbf{X_{t_n}})$ is equal in distribution to $(\mathbf{X_{t_1+k}}, \ldots, \mathbf{X_{t_n+k}})$ for all $n, k, t_1, \ldots, t_n \in \mathbb{Z}^+$. Most risk factors are assumed to be stationary.

Value-at-Risk

Value-at-Risk (VaR) is the most widely used risk measure in the financial industry. Despite the many weaknesses of VaR, financial institutions are required to use it under various regulatory frameworks. In addition, many institutions routinely report their VaR numbers to shareholders or investors. VaR is a risk measure based on the loss distribution and our discussion will not depend on whether we are dealing with the conditional or unconditional loss distribution. Nor will it depend on whether we are using the true loss distribution or some approximation to it. We will assume that the horizon, Δ, has been fixed, e.g. 1 day or 10 days, and that the random variable L represents the loss on the portfolio under consideration over the time interval Δ. We will use $F_L(\cdot)$ to denote the cumulative distribution function (CDF) of L. We first define the *quantiles* of a CDF.

Definition 3 *Let $F : \mathbb{R} \rightarrow [0,1]$ be an arbitrary CDF. Then for $\alpha \in (0,1)$ the α-quantile of F is defined by*

$$q_\alpha(F) \ := \ \inf\{x \in \mathbb{R} \ : \ F(x) \geq \alpha\}.$$

Note that if F is continuous and strictly increasing, then $q_\alpha(F) = F^{-1}(\alpha)$. For a random variable L with CDF $F_L(\cdot)$, we will write $q_\alpha(L)$ instead of $q_\alpha(F_L)$.

Definition 4 *Let $\alpha \in (0,1)$ be some fixed confidence level. Then the VaR of the portfolio loss at the confidence level, α, is given by VaR$_\alpha := q_\alpha(L)$, the α-quantile of the loss distribution.*

Example 1 (The Normal and t Distributions)
Because the normal and t CDFs are both continuous and strictly increasing, it is straightforward to calculate their VaR$_\alpha$. If $L \sim \mathsf{N}(\mu, \sigma^2)$ then

$$\mathsf{VaR}_\alpha \ = \ \mu + \sigma\Phi^{-1}(\alpha) \quad \text{where } \Phi \text{ is the standard normal CDF.} \tag{15}$$

By the previous lemma, this follows if we can show that $F_L(\mathsf{VaR}_\alpha) = \alpha$. But this follows immediately from (15).

If $L \sim \mathsf{t}(\nu, \mu, \sigma^2)$ so that $(L - \mu)/\sigma$ has a standard t distribution with $\nu > 2$ degrees-of-freedom, then

$$\mathsf{VaR}_\alpha \ = \ \mu + \sigma t_\nu^{-1}(\alpha) \quad \text{where } t_\nu \text{ is the CDF for the t distribution with } \nu \text{ degrees-of-freedom.}$$

Note that in this case we have $\mathsf{E}[L] = \mu$ and $\mathrm{Var}(L) = \nu\sigma^2/(\nu - 2)$. ∎

VaR has several weaknesses:

1. VaR attempts to describe the entire loss distribution with a single number and so significant information is not captured in VaR. This criticism does of course apply to all scalar risk measures. One way around this is to report VaR$_\alpha$ for several different values of α.

2. There is significant model risk attached to VaR. If the loss distribution is heavy-tailed, for example, but a normal distribution is assumed, then VaR$_\alpha$ will be severely underestimated as α approaches 1.

3. VaR is not a *sub-additive* risk measure so that it doesn't lend itself to *aggregation*. For example, let $L = L_1 + L_2$ be the total loss associated with two portfolios, each with respective losses, L_1 and L_2. Then

$$q_\alpha(F_L) \ > \ q_\alpha(F_{L_1}) \ + \ q_\alpha(F_{L_1}) \quad \text{ is possible.} \tag{16}$$

   In the risk literature this is viewed as being an undesirable property as we would expect some diversification benefits when we combine two portfolios together. Such a benefit would be reflected by the combined portfolio having a *smaller* risk measure than the sum of the two individual risk measures.

4 Algorithmic Trading

The key motivation for algorithmic trading is that investors are no longer *price-takers* but instead are adversely affected by their trading. In particular, when they wish to buy (sell) securities the price begins to rise (fall) and so they end up paying more (receiving less) than the original price that prevailed at the beginning of their trading. Accounting for these adverse price moves by modeling the *temporary* and *permanent* price impact is important for both high-frequency trading as well as lower-frequency trading where large trades need to be executed in a relatively short period of time. The size of these price impacts depends on the *liquidity* of the security. While there is no strict definition of liquidity, we think of a security as being liquid if we can quickly trade large quantities of the security with little price impact, i.e. low *slippage*. Characteristics of the security itself, as well as other factors such as general market conditions, time of day etc. determine how liquid a security is.

4.1 Optimal Execution of a Single Stock

We consider a simple model for selling a total of X shares over the course of a single day and in order to execute this trade we will break the day up into a total of T time periods. For example, if we take the duration of a time period to be 5 minutes, then we will have $T = 78$ time periods if we are trading on the NYSE[8]. We let n_j denote the number of shares sold in the j^{th} period and we let $\mathbf{n} := (n_1, \ldots, n_T)$ denote the execution sequence for the T periods. We define $x_k := X - \sum_{j=1}^{k} n_j$ for $k = 1, \ldots, T$, so that x_k denotes the number of shares that have yet to be sold after the k^{th} period. We set $x_0 = X$ and note that we must have $X = \sum_{j=1}^{T} n_j$.

We let S_k denote the pre-trade price for the k^{th} period and let $\widehat{S}_k$ denote the *realized* price-per-share that is obtained for the n_k shares sold at that time. We assume a temporary price impact function, $h(n)$, so that

$$\widehat{S}_k = S_k - h(n_k). \tag{17}$$

We model the permanent price impact via the function, $g(n)$, so that

$$S_{k+1} = S_k + \sigma z_k - g(n_k) \tag{18}$$

where the z_k's are IID standard normal random variables and σ is a volatility parameter. Note that the temporary price impact, $h(n_k)$, in period k only affects the realized price in period k. In particular, it is clear from (18) that $h(n_k)$ does not influence prices in later periods.

Using (17) and (18) we see that the total realized revenue of the execution strategy satisfies

$$\begin{aligned}
\sum_{k=1}^{T} \widehat{S}_k n_k &= \sum_{k=1}^{T} \left(S_k - h(n_k)\right) n_k \\
&= \sum_{k=1}^{T} \Big(S_1 + \sum_{j=1}^{k-1} \left(\sigma z_j - g(n_j)\right)\Big) n_k - \sum_{k=1}^{T} h(n_k) n_k && (19) \\
&= S_1 X + \sigma \sum_{k=1}^{T} z_k x_k - \sum_{k=1}^{T} g(n_k) x_k - \sum_{k=1}^{T} h(n_k) n_k. && (20)
\end{aligned}$$

Note that in going from (19) to (20) we have switched the order of summation for the middle two terms and also used $x_T = 0$ so that the middle two sums in (20) can run from $k = 1$ to T rather than $T - 1$. The expected cost of the execution strategy is

$$C(\mathbf{n}) = X S_1 - \sum_{k=1}^{T} g(n_k) x_k - \sum_{k=1}^{T} h(n_k) n_k.$$

[8] A trading day on the NYSE lasts 6.5 hours with the opening at 9.30am and the close at 4pm.

If we measure riskiness by variance, then the risk of the execution strategy is $V(\mathbf{n}) := \sigma^2 \sum_{k=1}^{T} x_k^2$. We can now formulate the optimal execution optimization problem as

$$\min_{\mathbf{n} \geq 0} \; C(\mathbf{n}) + \rho V(\mathbf{n}) \tag{21}$$

where ρ is a parameter we choose to tradeoff cost versus risk. Before solving (21) we also need to specify the price impact functions. A typical choice for g is a linear permanent price impact so that $g(v) = \gamma v$, for some constant γ. The temporary price impact function is typically chosen to be non-linear. One possibility would be to use the Kissel-Glantz function

$$h(n) = a_1 \left(\frac{100\,|n|}{V} \right)^{\beta} + a_2 \sigma + a_3$$

where V is the average daily traded volume and the a_i's are parameters that can be estimated via regression. The problem (21) can now be solved via standard non-linear optimization methods.

There are many other important features that can be incorporated into this execution problem. For example, we assumed T is fixed but one can also include it as a decision variable to be optimized. In practice many of the parameters / impact functions are time-dependent. For example, it is well-known that there is considerably more market liquidity at the open and close every day and so trading at those times should incur a smaller price impact. It is also possible to include some *predictability* in the price dynamics in which case the resulting problem will then become a dynamic program (DP) rather than a static optimization problem. One can also formulate *portfolio* execution problems where the goal is to trade many different securities. Now trading in one security can alter the dynamics of the other securities and this is modeled via *cross-price impacts*. Estimating these cross-price impacts can be very difficult, however, and so it is common to collapse the portfolio execution problem into a series of single-stock execution problems.

4.2 Limit Order Books and Dark Pools

The execution model in Section 4.1 only accounts for one level of the overall trading strategy. Once we have determined n_k, the number of shares to be sold in period k, we must decide how to actually sell these shares. In particular, we need to construct another strategy for how to sell these shares in the current period. Such a strategy can consist of *limit* orders and *market* orders executed via the *limit-order book* but it could also include sending the trades to so-called *dark-pools*. The goal at all times it to obtain the best execution and to prevent other market participants from learning about your strategy and therefore *front-running* your trades.

A limit-order book is a database that keeps buy orders, i.e. bids, and sell orders, i.e. offers, on a *price-time-priority* basis. A limit order has an associated quantity and price at which the order (buy or sell) can be executed. Limit orders arrive regularly to the order book but the execution of a limit order is uncertain and may never occur.

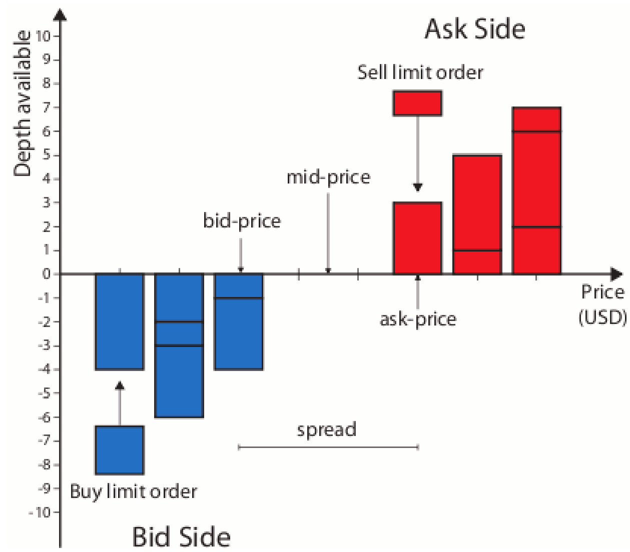

Taken from the article "Limit Order Books" (2013) by Gould, Porter, Williams, McDonald, Fenn and Howison.

This is in contrast to a market order. Consider, for example, a market order to buy Q shares and let p_o be the best, i.e. lowest, offer price in the limit order book. Let Q_o be the number of shares offered at p_o. Then the market order will purchase $\min(Q, Q_o)$ shares at a price p_o. If $Q > Q_o$, then the remaining $Q - Q_o$ shares will be executed at successively higher offer prices in the order-book until all of the shares are purchased. Market orders to sell shares are transacted similarly by "hitting" the best bids in the order-book. The advantage of a market order over a limit order is that execution is guaranteed and immediate. The disadvantage is that the execution price is not as good and it is necessary to pay (at least) the bid-ask spread, which is the difference between the best bid and best offer in the order book.

In contrast to limit-order books, dark pools are trading venues where blocks of shares can be bought or sold without revealing either the size of the trade or the identity of the agents until the trade is filled. In contrast, regular exchanges (with limit-order books) are called *lit* pools. Consequently, dark pool trading hopes to avoid market impact but the volume of shares executed is uncertain. There is also some controversy associated with dark pools as they are said to hinder the "price-discovery" process.

Optimal execution in limit-order books and dark pools is currently a very active research area. It is also currently the focus of many market regulators who seek to understand and prevent adverse events such as the flash-crash of May 2010.