Jon Danielsson
London School of Economics

Financial Risk Forecasting

Seminar Questions

financialriskforecasting.com

2025

Version 4.0

Chapter 1

Financial data

1. Which of the following statements about price-weighted versus value-weighted indices is correct?

   a Price-weighted indices give more weight to companies with higher market capitalization

   b Value-weighted indices are less representative of overall market performance

   c A stock split will decrease a stock's weight in a price-weighted index but not affect its weight in a value-weighted index

   d Price-weighted indices are more commonly used than value-weighted indices

2. Explain the concept that "risk is latent" and why this creates challenges for risk forecasting models that rely on historical data.

3. A stock's price moves from \$100 to \$120 on day 1, then from \$120 to \$108 on day 2.

   i Calculate the simple returns for both days.

   ii Calculate the continuously compounded returns for both days.

   iii What is the total return over the two days using simple returns?

   iv What is the total return over the two days using continuously compounded returns?

   v Explain why continuously compounded returns demonstrate the "symmetry" property.

4. Explain what volatility clustering is, describe how it can be tested statistically, and discuss its implications for risk management.

5. i Consider the autocorrelation function (ACF) plots for daily returns and squared returns of a financial asset, and what to they capture?

   ii How do these ACF patterns relate to market efficiency?

6. Suppose you calculate the first-order autocorrelation of both returns and returns squared. Which of these two would you expect to be more strongly statistically significant?

7. Explain the concept of "fat tails" in financial returns.

   i Define fat tails formally.

   ii Provide an intuitive explanation of what this means for extreme events.

   iii How can fat tails be measured or detected?

   iv What distribution is commonly used to model fat tails?

8. You observe the following statistics for daily S&P 500 returns: Mean = 0.05%, Standard deviation = 1.2%, Skewness = -0.3, Kurtosis = 5.2.

   i What do these statistics tell you about the distribution of returns?

   ii How does this compare to a normal distribution?

   iii What test could you use to formally test for normality?

9. During the 2008 financial crisis, correlations between major stock indices increased dramatically. For example, correlations between the S&P 500, FTSE 100, and Nikkei 225 rose from around 0.4-0.6 in normal times to over 0.8 during the crisis.

   i What stylized fact does this illustrate?

   ii Why is this problematic for portfolio diversification?

   iii How does this relate to the assumption of joint normality?

10. What is the difference between a price weighted and a value weighted stock market index? Which type is more useful for financial analysis? Give one example of each.

11. What are simple and continuously compounded returns?

12. What are the advantages and disadvantages of using continuously compounded returns?

13. What are the three main stylized facts of financial returns?

14. A Japanese investor has JPY 1,000,000 invested in the TOPIX index. The historical annual volatility is 20% and the annual mean return is 2%. The annual Japanese inflation is 0.1%. Based on this information, would you expect the long-run real investment performance to be less than 2%, equal to 2%, or more than 2% annually?

Chapter 2

Volatility forecasting

1. List three major problems with the moving average (MA) volatility model and explain why EWMA is preferred to the MA models.

2. Consider an EWMA model with $\lambda = 0.94$. If yesterday's return was $y_{t-1} = 0.02$ and yesterday's volatility forecast was $\hat{\sigma}_{t-1} = 0.015$, calculate today's volatility forecast $\hat{\sigma}_t$.

3. Suppose you estimate a normal ARCH model for a stock, finding the parameter estimates for day t to be

$\hat{\alpha}$	0.5
$\hat{\omega}$	0.00001

   Making all necessary assumptions, what is the conditional kurtosis?

4. Describe the relationship between the EWMA model and GARCH, and explain why EWMA has undefined unconditional volatility.

5. A GARCH(1,1) model for daily stock returns has the following parameter estimates: $\hat{\omega} = 0.00001$, $\hat{\alpha} = 0.12$, $\hat{\beta} = 0.85$. Yesterday's return was $y_t = -0.03$ and yesterday's conditional variance was $\hat{\sigma}_t^2 = 0.0009$. Calculate today's conditional variance forecast and determine if this GARCH model is stationary.

6. A GARCH(1,1) model has parameter estimates $\hat{\omega} = 0.00001$, $\hat{\alpha} = 0.1$, and $\hat{\beta} = 0.9$. Yesterday's return was $y_t = -0.1$ and yesterday's conditional variance was $\hat{\sigma}_t^2 = 0.0002$. Calculate both the unconditional and conditional volatility, and explain the implications when $\hat{\alpha} + \hat{\beta} = 1$.

7. For a GARCH(1,1) model with parameters $\omega = 0.000001$, $\alpha = 0.08$, and $\beta = 0.90$, calculate the unconditional volatility.

8. Contrast conditional and unconditional volatility and kurtosis in volatility models. Explain how ARCH/GARCH models can produce fat-tailed returns while assuming conditionally normal residuals.

9. Explain why the parameter restriction $\alpha + \beta < 1$ is not always imposed in GARCH estimation, despite it being required for stationarity.

10. In a GARCH(1,1) model, distinguish between the roles of α and β in terms of "news" and "memory".

11. Calculate the half-life of volatility shocks for a GARCH(1,1) model with $\alpha = 0.05$ and $\beta = 0.92$.

12. Explain the difference between conditional and unconditional fat tails in the context of tGARCH versus normal GARCH models.

13. In the APARCH model $\sigma_t^2 = \omega + \alpha(|y_{t-1}| - \zeta y_{t-1})^\delta + \beta\sigma_{t-1}^\delta$, explain how the parameters ζ and δ capture leverage and power effects. Also, explain the economic intuition behind the leverage effect in equity markets.

14. Explain why the likelihood function for GARCH models starts from $t = 2$ rather than $t = 1$.

15. In maximum likelihood estimation, distinguish between local and global maxima problems and suggest how to address them.

16. Given the optimization challenges in volatility model estimation, explain why simpler models like GARCH(1,1) are often preferred over more complex alternatives.

17. Compare and contrast the concepts of unconditional and conditional volatility.

Chapter 3

Multivariate volatility

1. Why must a covariance matrix be positive definite and what happens if this condition is violated?
2. How do dimensionality and positive definiteness interact in practice?
3. A portfolio manager needs to estimate the covariance matrix for 50 stocks. How many unique parameters need to be estimated and what is this problem called?
4. The EWMA model is one of the most widely used approach for multivariate volatility modeling in practice. Write down the EWMA model for forecasting the covariance matrix for two assets.
5. Why does EWMA automatically ensure positive definiteness ?
6. A portfolio has weights $w = [0.6, 0.4]'$ and covariance matrix $\Sigma = \begin{pmatrix} 0.04 & 0.02 \\ 0.02 & 0.09 \end{pmatrix}$.

   i Calculate the portfolio variance.

   ii What is the portfolio volatility?
7. Identify the main advantage and main limitation of the EWMA model.
8. Compare the CCC and DCC models for multivariate volatility modeling, covering their theoretical foundations, estimation procedures, and practical applications.

   Explain what each acronym stands for and the fundamental difference between the models.
9. Why is the DCC model generally preferred over CCC and what are the trade-offs?
10. Write down the DCC model equations and explain the two-step estimation procedure.
11. How does crisis behavior in prices affect the performance of standard multivariate volatility models?
12. What practical approaches can risk managers use to address problems caused by crisis dynamics?

Chapter 4

Risk measures

1. Explain why risk is considered a "latent variable" and how this differs from measuring observable quantities like stock prices.
2. Consider three assets A, B, and C, all with identical means ($\mu = 0.08$) and variances ($\sigma^2 = 0.04$) but different distribution shapes. Asset A has normally distributed returns, Asset B has positively skewed returns, and Asset C has negatively skewed returns with fat tails. Explain why different investors might prefer different assets despite the MV model treating them as equivalent.
3. Does ES satisfies the monotonicity axiom of coherent risk measures? Consider two portfolios: A has returns that are always \$100 lower than B for every possible outcome. Which portfolio should have higher risk according to monotonicity?
4. Test whether volatility satisfies the translation invariance axiom of coherent risk measures. If an asset has volatility $\sigma_A = 0.03$ and we add a constant $c = 0.001$ to all returns, what happens to the volatility?
5. Demonstrate positive homogeneity for VaR. If a portfolio worth \$5 million has a 1% daily VaR of \$100,000, what is the 1% VaR if the portfolio size doubles to \$10 million?
6. Show how VaR can violate subadditivity using the example from Chapter 4. Two independent assets A and B each have a 4.9% probability of losing 100 and 95.1% probability of no loss. Calculate the 5% VaR for each asset individually and for an equally weighted portfolio.
7. Using the square root of time scaling rule, calculate the 10-day VaR from a 1-day VaR of \$50,000. Explain when this scaling rule is accurate and when it may fail.
8. Explain how a trader might manipulate VaR using derivatives.

Chapter 5

Implementing risk forecasts

1. A portfolio manager uses HS to calculate 1% VaR with 500 days of return data. The sorted returns (from smallest to largest) show the 5th smallest return is -0.032 and the 6th smallest is -0.029. For a portfolio worth $25 million, calculate the 1% VaR and explain the key assumption underlying this method.

2. A portfolio manager has collected the following 20 daily returns (in %) for a stock over the past month:

-2.1	1.3	-0.8	2.4	-1.5	0.9
-3.2	0.7	1.8	-2.8	0.4	-1.1
2.2	-0.6	1.4	-4.1	0.3	1.7
-1.9	-0.5				

   For a portfolio value of $50,000, calculate the 1-day 5% VaR and ES using HS, and comment on the reliability of these estimates given the sample size.

3. Outline one advantage and two disadvantages of using HS for forecasting risk.

4. A stock is currently priced at $100. Calculate the 1-day 5% VaR using both simple returns and continuously compounded returns, assuming daily volatility is 2% and returns are normally distributed. Comment on the difference.

5. Calculate the 1-day 5% VaR for a $50,000 portfolio with normally distributed returns, zero mean, and daily volatility of 1.8%.

6. Suppose you estimate a t-GARCH model for a stock with the following parameter estimates for day t:

$\hat{\alpha}$	0.1
$\hat{\omega}$	0.00001
$\hat{\beta}$	0.85
$\hat{\nu}$	∞
y_t	-0.1
$\hat{\sigma}_t^2$	0.0002

What is the 1-day 1% VaR forecast for day $t+1$? State all necessary assumptions.

7. Using the t-GARCH parameter estimates from Question 6, calculate the 2-day 1% VaR forecast using the square-root-of-time rule.

8. A portfolio worth JPY 1,000,000 has annual volatility of 20% and annual mean return of 2%. Calculate the 1-month and 1,000-year VaR using the square-root-of-time rule, and explain what fundamental issues the long-term VaR result reveals about the limitations of VaR as a risk measure. Use 1% probability level and assume IID normally distributed returns.

9. Explain what "model risk" means in the context of VaR estimation and provide two specific examples of how it can manifest in practice.

10. Using the t-GARCH parameter estimates from Question 6, suppose instead that $\hat{\nu} = 3$ degrees of freedom. Calculate the 1-day 1% VaR forecast for day $t+1$.

11. For a portfolio with normally distributed returns having mean $\mu = 0$ and standard deviation $\sigma = 0.03$, calculate both the 5% VaR and 5% ES. Explain why ES is always greater than VaR and what this difference represents.

12. Why is the mean often ignored in daily VaR calculations?

13. A risk manager compares VaR estimates from three different models for the same portfolio:

Model	1-day 5% VaR
HS	\$450,000
Normal GARCH	\$420,000
t-GARCH ($\nu = 4$)	\$510,000

Explain why these estimates differ and which model would be most appropriate for stress testing purposes.

14. Explain what "burn-in time" means in the context of EWMA volatility models and why it is necessary for reliable risk forecasting.

15. Explain why financial risk measurements are only accurate to about 2 significant digits. If a risk model outputs a VaR of 7.8432%, what should be reported and why?

Chapter 6

Analytical methods

1. Briefly explain the problem of asymmetry in bond risk analysis?
2. What does the modified duration of a bond measure?
3. What factors affect the accuracy of duration-normal VaR for bonds?
4. What computational issue arises when using the Delta-Gamma method for option VaR?
5. Why are the methods in this chapter not recommended for most applications?
6. Explain why bond prices are more sensitive to interest rate changes for longer maturities. Include the concept of duration in your explanation.

Chapter 7

Simulation methods

1. Explain the trade-off between doing too few and too many simulations in Monte Carlo analysis. How should a risk manager determine the optimal number of simulations?

2. What is the benefit of simulations over analytical methods in quantitative financial analysis?

3. A risk manager wants to use Monte Carlo simulation to estimate the 1% VaR for a portfolio. Explain the key components of a pseudo-random number generator, discuss the concept of period in RNGs, and calculate how many simulations are theoretically possible using the Mersenne Twister RNG before the sequence repeats.

4. A portfolio consists of 2,000 shares of stock (current price $35, daily volatility 2.5%) and 100 call options on the same stock (strike $40, 2 months to expiration, Black-Scholes price $1.80). Calculate the 1% VaR using Monte Carlo simulation, explaining how options modify the standard procedure.

5. Consider using Monte Carlo methods for forecasting risk. Outline how we forecast the risk of a portfolio consisting of two stocks and an option on one of them.

Chapter 8

Backtesting

1. Why is it important to evaluate the quality of a risk forecast model, and why is it difficult to use operational criteria for this evaluation?
2. Define the following key backtesting terms: estimation window, testing window, VaR violation, and violation ratio.
3. Consider the coverage test used in backtesting VaR models. Explain how it works, identify its main advantage and main disadvantage.
4. Explain why backtesting ES is significantly more challenging than backtesting VaR. Provide specific reasons and discuss potential approaches to address these challenges.
5. What is the independence test in backtesting and what is its main disadvantage?
6. Identify one reason why backtesting ES is more challenging than backtesting VaR.
7. How would we apply backtesting to evaluate the quality of volatility forecast models?
8. Consider four GARCH(1,1) parameter estimates for daily stock returns:

Set	ω	α	β	$\alpha + \beta$
A	0.000010	0.05	0.90	0.95
B	0.000015	0.15	0.84	0.99
C	0.000030	0.08	0.90	0.98
D	0.000010	0.08	0.90	0.98

   Which parameter sets would be most appropriate for (a) the 2008 financial crisis and (b) the calm period of 2004-2006? Explain your reasoning for each choice.

9. Consider two GARCH(1,1) parameter sets: Set A ($\omega = 0.000010$, $\alpha = 0.05$, $\beta = 0.90$) and Set B ($\omega = 0.000015$, $\alpha = 0.15$, $\beta = 0.84$). Calculate the unconditional variance and half-life of volatility shocks for each set. What do these calculations reveal about different market regimes, and what are the implications for backtesting?
10. A bank's VaR model produces the following 1% daily VaR forecasts and actual P&L outcomes over 10 consecutive trading days (in millions):

Day	VaR Forecast	Actual P&L	Violation?
1	-2.5	1.2	No
2	-2.8	-3.1	Yes
3	-2.3	-0.8	No
4	-2.7	-2.9	Yes
5	-2.4	0.5	No
6	-2.6	-1.8	No
7	-2.5	-4.2	Yes
8	-2.9	-0.3	No
9	-2.2	-2.8	Yes
10	-2.8	1.8	No

Calculate the violation rate, perform the coverage test at 5% significance level, and interpret the results.

Chapter 9

Extreme Value Theory

1. State and briefly describe the three types of tail distributions.
2. What is the formal definition of fat tails?
3. What are the two common approaches to POT (peaks over thresholds)? Which approach would you prefer?

Chapter 10

Endogenous risk

1. Define the terms endogenous and exogenous risk, and provide examples of each.

2. The former general manager of the BIS, Andrew Crockett stated in 2000:

   > "The received wisdom is that risk increases in recessions and falls in booms. In contrast, it may be more helpful to think of risk as increasing during upswings, as financial imbalances build up, and materialising in recessions."

   How does his view on risk relate to endogenous risk, and what is the implication for governments' policies on financial stability?

3. It has been said that financial risk models are least reliable when we need them the most. Explain this statement.

4. You run a Japan-oriented hedge fund with a target leverage of five. You own stock in only one company, where the price of one stock is ¥20,000. The number of stocks in your portfolio is 10. Your hedge fund is very small and you have no price impact when you trade. Suppose the price of the stock falls by ¥1,000. Making all necessary assumptions, how many stocks would be left in your portfolio after the necessary rebalancing?

5. Starting with the previous question, suppose instead that your fund is large and exerts significant pricing power. In particular, for every ¥1,000,000 you trade, the price of the stock moves by ¥2,000 in the direction of the trade (up if you buy, down if you sell). Suppose the price of the stock falls for exogenous reasons by the same amount as in the previous question. Making all necessary assumptions, how many stocks would be left in your portfolio after the necessary rebalancing?

6. Consider the price of a typical stock and the VIX index. Are they mean reverting? Explain your reasoning and discuss the implications for risk management.

7. Recall the case of the Long-Term Capital Management (LTCM) hedge fund and its investment strategies related to volatility. In the months before LTCM's default in 1998, the VIX was steadily rising, causing increasing distress for LTCM. Describe in

detail why the rising VIX caused the default of LTCM and analyze this as an example of endogenous risk.

Chapter 11

Regulation

1. Explain the key changes introduced by Basel III for market risk measurement compared to Basel II, focusing on the shift from Value-at-Risk to ES.

2. The Basel Accords are overseen by the Basel Committee on Banking Supervision (BCBS) and implemented by G20 countries. Explain the institutional framework and describe the key components of the Basel III capital requirements.

3. A bank's trading book has the following VaR statistics over the past 250 trading days: current day 1% VaR = \$8 million, 60-day average VaR = \$10 million, and 6 violations occurred in the testing period. Calculate the market risk capital requirement under Basel II regulations and explain the traffic light system.

4. The Swiss bank UBS failed in 2008 due to \$19 billion in losses on collateralized debt obligations (CDOs) composed of U.S. sub-prime mortgages. Despite these massive losses, UBS's VaR models showed zero risk for these positions. Analyze this case study and explain how this illustrates the limitations of VaR as a risk measure.

5. Basel III introduced significant reforms to the trading book regulations, replacing the 10-day 99% VaR with 97.5% ES calculated over various holding periods. Explain the key changes from Basel II to Basel III in market risk measurement and analyze the advantages and challenges of these reforms.

6. Explain why financial institutions are more heavily regulated than most other private firms. Discuss the specific externalities and information asymmetries that justify this regulatory treatment.

7. Capital serves as a buffer against unexpected losses and as a limit to leverage. A commercial bank has the following simplified balance sheet (in billions):

Assets		Liabilities & Equity	
Cash	\$50	Deposits	\$800
Loans	\$900	Bonds	\$200
Securities	\$100	Equity	\$50
Total	\$1,050	Total	\$1,050

The bank expects loan losses of $20 billion over the next year, but actual losses could range from $10 billion to $40 billion. Analyze how capital serves as a buffer against unexpected losses and calculate the bank's leverage ratio.