BIIRG

VOLUME V

QUANTITATIVE RESEARCH

Analyst Training Manual

Statistics, econometrics, factor models, portfolio construction, risk, time series, machine learning, backtesting, and systematic research.

Baruch Islamic Investment Research Group

TRAINING & RESEARCH SERIES · QUANTITATIVE RESEARCH DIVISION **2026 EDITION**

PUBLICATION INFORMATION

TITLE	Quantitative Research — Analyst Training Manual
SERIES	BIIRG Analyst Training Series, Volume V of V
PUBLISHER	Baruch Islamic Investment Research Group
DIVISION	Quantitative Research Division, Training & Research
EDITION	2026 Edition — August 2026
AUDIENCE	Incoming and continuing BIIRG quantitative analysts

BIIRG ANALYST TRAINING SERIES

A five-volume research curriculum

Volume I Equities Research	Volume II Fixed Income & Sukuk
Volume III Real Assets	Volume IV Macro Research
Volume V Quantitative Research THIS VOLUME	

The volumes are designed to be read in any order and used as reference throughout an analyst's tenure. Quantitative Research is the measurement volume: it supplies the estimation, testing, and portfolio-construction machinery the other four volumes rely on when they make a claim about risk or return.

VOLUME V

Contents

FRONT MATTER

How to Use This Manual

Quantitative research converts financial questions into measurable hypotheses, datasets, models, tests, and portfolio decisions.

The analyst's task is not merely to run regressions or optimize a portfolio. It is to understand the economic question, construct valid data, select a model appropriate to the data-generating process, test whether an apparent relationship survives statistical and economic scrutiny, and determine whether it remains useful after costs, constraints, estimation error, and regime change.

BIIRG STANDARD — THE RESEARCH CHAIN

Question → economic hypothesis → data → transformations → exploratory analysis → model → statistical validation → economic validation → backtest → robustness → portfolio construction → risk → implementation → monitoring.

Every link must be auditable.

The component system

Recurring kinds of material are marked consistently throughout the manual.

DEFINITION
Formal definition of a concept or term.

CORE CONCEPT
Foundational theory the rest of a chapter rests on.

ANALYST NOTE
Practical interpretation and desk-level judgment.

EXAMPLE
Illustrative application of a method.

BIIRG STANDARD
Expected BIIRG methodology or research standard.

KEY TAKEAWAY
The central conclusion of a section.

FORMULA
Relationship, variable definitions, interpretation, application.

WATCH FOR
Analytical mistakes, data pitfalls, model weaknesses.

INVESTMENT IMPLICATION
Translation of statistical output into a portfolio or capital-allocation decision.

Core analytical principle

A statistically significant relationship without an economic mechanism is a candidate for a false discovery, not a strategy. Conversely, a sound economic mechanism whose estimate is drowned in estimation error is not yet a usable model. Quantitative research is the discipline of holding both requirements at once, and of documenting the point at which either one fails.

FRONT MATTER

Learning Outcomes

On completing this volume, a BIIRG quantitative analyst should be able to do each of the following without supervision.

01 Master return mathematics, probability, descriptive statistics, covariance, correlation, distributions, sampling, and hypothesis testing.

02 Understand OLS deeply enough to derive, estimate, interpret, diagnose, and challenge linear regressions.

03 Use time-series concepts including stationarity, autocorrelation, AR/MA/ARIMA, volatility models, cointegration, and regime methods.

04 Understand CAPM, Fama-French and broader factor models, factor exposure, alpha, residual risk, and attribution.

05 Construct and evaluate portfolios using mean-variance methods, risk budgeting, constraints, covariance estimation, and robust alternatives.

06 Measure volatility, drawdown, VaR, expected shortfall, tracking error, information ratio, Sharpe, Sortino, beta, and tail risk.

07 Design backtests that avoid look-ahead bias, survivorship bias, overfitting, leakage, and unrealistic execution assumptions.

08 Understand cross-validation, regularization, tree models, clustering, dimensionality reduction, and appropriate uses of machine learning.

09 Translate models into systematic signals, sizing rules, rebalancing logic, and monitoring frameworks.

10 Maintain reproducible research standards in Excel, Python, versioned datasets, and written research documentation.

PART I

Research Process and Return Mathematics

How a question becomes a falsifiable hypothesis, and the return and frequency conventions every later chapter assumes. Errors made here propagate silently through everything downstream.

01 The Quantitative Research Process

Start with an economic question, state the null, and decide whether the objective is explanation or prediction before choosing a method.

1.1 Start with an economic question

A model should answer a defined question. Does valuation predict cross-sectional returns? Does macro regime alter factor performance? Can volatility be forecast better than a rolling historical estimate? Does a liquidity signal predict drawdowns? A statistically significant relationship without economic rationale is especially vulnerable to data mining.

1.2 Hypothesis structure

TABLE 1.1 — ANATOMY OF A TESTABLE HYPOTHESIS

ELEMENT	EXAMPLE
Economic mechanism	Tighter liquidity raises required risk premia
Observable proxy	Financial-conditions index / credit spreads
Target	Next-month equity excess return
Null	No predictive relationship
Alternative	Tighter conditions predict lower future return
Validation	Out-of-sample, alternative periods, controls, costs

Source: BIIRG Quantitative Research Division.

1.3 Prediction versus explanation

Explanatory models seek interpretable causal or associational structure. Predictive models seek out-of-sample accuracy. A variable can be useful for prediction without being causal, and a causal variable may have weak forecasting power. State the objective before choosing the method.

1.4 Reproducibility

- Version raw and cleaned data.
- Document sources and timestamps.
- Separate data ingestion, transformation, model, and output layers.
- Set random seeds where relevant.
- Record parameters and universe definitions.
- Preserve failed tests; do not retain only successful specifications.

02 Financial Data and Return Mathematics

Simple and log returns, compounding, annualization conventions, and why frequency is never a neutral choice.

2.1 Simple return

Simple return: $R_t = (P_t - P_{t-1} + D_t) / P_{t-1}$. It is intuitive and aggregates cross-sectionally, but multi-period simple returns compound multiplicatively.

2.2 Log return

Log return: $r_t = \ln(P_t/P_{t-1})$ for price-only data. Log returns add across time and are convenient statistically. For large moves they differ materially from simple returns.

2.3 Compounding

Cumulative return = $\Pi(1 + R_t) - 1$. Arithmetic averaging does not produce multi-period compound return.

2.4 Annualization

TABLE 2.1 — ANNUALIZATION CONVENTIONS

METRIC	COMMON ANNUALIZATION
Mean periodic return	Mean × periods per year, an approximation
Volatility	Periodic standard deviation × √(periods per year)
Geometric return	$(\text{Ending} / \text{Beginning})^{1/\text{years}} - 1$
Sharpe	Annual excess return / annual volatility

Source: BIIRG Quantitative Research Division.

2.5 Excess and active returns

Excess return is return above a risk-free or specified hurdle. Active return is portfolio return minus benchmark return. These concepts feed Sharpe and information ratios respectively.

2.6 Data frequency

Daily, weekly, and monthly observations contain different microstructure noise and sample sizes. Never combine volatilities from different frequencies without converting to a common basis and understanding serial dependence.

WATCH FOR

The √(periods) volatility scaling assumes serially uncorrelated returns. Where returns are autocorrelated, illiquid, or appraisal based, that scaling understates or overstates annual risk, and every Sharpe ratio built on it inherits the error.

PART II

Probability and Statistical Foundations

Random variables, distributions, estimation uncertainty, testing, dependence, and the linear algebra that makes portfolio and regression calculations scale. Chapters 3 to 8 are the machinery the rest of the volume calls.

03 Probability and Random Variables

Expectation, conditioning, and Bayesian updating — the language in which every regime and scenario statement is actually made.

3.1 Random variables

A random variable maps outcomes to numerical values. Discrete variables take countable values; continuous variables take values over intervals.

3.2 Expected value and variance

E[X] is the probability-weighted mean. $\mathrm{Var}(X) = E[(X - E[X])^2]$. Standard deviation is the square root of variance.

3.3 Conditional probability

$P(A|B) = P(A \cap B) / P(B)$. Conditional reasoning is central to regime analysis: the probability of a drawdown given recession is different from the unconditional probability.

3.4 Bayes' rule

FORMULA — BAYES' THEOREM

$$P(A|B) = P(B|A) \cdot P(A) / P(B)$$

WHERE	P(A) is the prior, P(B|A) the likelihood of the evidence under A, and P(B) the unconditional probability of the evidence.
INTERPRETATION	Bayes' theorem updates prior beliefs with evidence.
APPLICATION	It provides a disciplined mental model for updating investment probabilities as data arrives, and makes the prior an explicit, criticizable input.

3.5 Law of total expectation

$E[X] = \Sigma\, P(S_i)\, E[X|S_i]$. Scenario-weighted expected returns are an investment application.

PART II · PROBABILITY AND STATISTICAL FOUNDATIONS CHAPTER 04

04 Descriptive Statistics and Distributions

Central tendency, dispersion, higher moments, and the convenience cost of assuming normality in return data.

4.1 Central tendency

- **Mean:** sensitive to outliers.
- **Median:** robust midpoint.
- **Mode:** most frequent observation.

4.2 Dispersion

- Variance and standard deviation.
- Range.
- Interquartile range.
- Median absolute deviation.

4.3 Skewness and kurtosis

Skewness measures asymmetry. Negative skew is common in strategies that earn small gains and occasional large losses. Excess kurtosis captures heavier tails than the normal distribution.

4.4 Normal distribution

The normal distribution is analytically convenient but often understates financial tail risk. Return distributions can be skewed, fat-tailed, heteroskedastic, and serially dependent.

4.5 Percentiles and z-scores

$z = (x − \text{mean}) / SD$. Standardization enables comparison across variables, but z-scores inherit the properties of the estimated mean and volatility.

PART II · PROBABILITY AND STATISTICAL FOUNDATIONS CHAPTER 05

05 Sampling, Estimation, and Confidence Intervals

Every parameter in this volume is an estimate. This chapter is about how much of it is noise.

5.1 Population versus sample

Investment datasets are samples from an evolving process, not fixed populations. Parameter estimates therefore contain uncertainty.

5.2 Standard error

Standard error measures sampling uncertainty of an estimator. For an IID sample mean, $SE = s / \sqrt{n}$, but financial dependence can invalidate IID assumptions.

DEFINITION — CONFIDENCE INTERVAL

A confidence interval gives a range generated by a procedure with a specified long-run coverage rate. It should not be interpreted as a probability statement about a fixed parameter under classical inference.

5.4 Bias-variance tradeoff

Flexible models can reduce in-sample bias but increase estimation variance. Robust research seeks the simplest model that captures economically useful structure.

5.5 Bootstrap

Bootstrapping resamples observations to estimate sampling distributions. For time series, block bootstrap methods can preserve serial dependence better than IID resampling.

PART II · PROBABILITY AND STATISTICAL FOUNDATIONS
CHAPTER 06

06 Hypothesis Testing

What a rejection does and does not establish, and why the search process itself has to enter the inference.

6.1 Null and alternative

The null hypothesis is the benchmark claim tested against an alternative. Rejecting the null means evidence is inconsistent with it at a chosen significance level; it does not prove the alternative with certainty.

6.2 Type I and Type II error

TABLE 6.1 — THE TWO ERRORS

ERROR	MEANING
Type I	Reject a true null; false positive
Type II	Fail to reject a false null; false negative

Source: BIIRG Quantitative Research Division.

6.3 p-values

A p-value is the probability, under the null and model assumptions, of observing a statistic at least as extreme as the one obtained. It is not the probability the hypothesis is true.

6.4 t-test and F-test

A t-test commonly evaluates one coefficient or mean restriction. An F-test evaluates joint restrictions, such as whether several regression coefficients are simultaneously zero.

KEY TAKEAWAY — MULTIPLE TESTING

If hundreds of signals are tested, some will appear significant by chance. Control false discoveries, use economic priors, holdout samples, and account for the research search process.

PART II · PROBABILITY AND STATISTICAL FOUNDATIONS
CHAPTER 07

07 Covariance, Correlation, and Dependence

Diversification is a statement about correlation, and correlation is a noisy, state-dependent, strictly linear measurement.

7.1 Covariance

$\mathrm{Cov}(X,Y) = E[(X - E[X])(Y - E[Y])]$. Positive covariance means variables tend to move in the same direction.

7.2 Correlation

$\rho_{XY} = \mathrm{Cov}(X,Y) / (\sigma_X \sigma_Y)$, bounded from −1 to +1. Correlation measures linear dependence, not causation.

FORMULA — TWO-ASSET PORTFOLIO VARIANCE

$$\sigma_p^2 = w_1^2\sigma_1^2 + w_2^2\sigma_2^2 + 2w_1w_2\sigma_1\sigma_2\rho_{12}$$

WHERE	w are weights, σ are volatilities, and ρ_{12} is the pairwise correlation.
INTERPRETATION	Diversification comes from imperfect correlation: the cross term shrinks as ρ falls.
APPLICATION	Any diversification claim is only as good as the correlation estimate behind it, so state the estimation window.

7.4 Correlation instability

Correlations are state-dependent and often rise during stress. Rolling estimates should be interpreted as noisy measurements, not constants.

7.5 Nonlinear dependence

Pearson correlation can miss nonlinear relationships and tail dependence. Rank correlations, copulas, conditional correlations, and scenario analysis may reveal additional structure.

PART II · PROBABILITY AND STATISTICAL FOUNDATIONS CHAPTER 08

08 Linear Algebra for Quant Finance

Matrix notation is what makes portfolio and econometric calculations scale, and what makes a badly estimated covariance matrix dangerous.

8.1 Vectors and matrices

Portfolio weights, returns, factor exposures, and regression systems are naturally represented with vectors and matrices. Matrix notation allows scalable portfolio and econometric calculations.

8.2 Portfolio return and variance

Expected portfolio return: $\mu_p = w'\mu$.

Portfolio variance: $\sigma_p^2 = w'\Sigma w$, where Σ is the covariance matrix.

8.3 Matrix inverse

The inverse covariance matrix appears in unconstrained mean-variance solutions. Poorly estimated covariance matrices can be nearly singular, producing unstable weights.

8.4 Eigenvalues and eigenvectors

Principal components decompose covariance into orthogonal directions. In rates, principal components often resemble level, slope, and curvature movements.

8.5 Positive semidefinite covariance

A valid covariance matrix must be positive semidefinite. Data errors or ad hoc manipulations can violate this property.

PART III

Regression and Inference

OLS derived, estimated, diagnosed, and challenged. Most of what goes wrong in applied finance regressions is an assumption problem, not an arithmetic one.

PART III · REGRESSION AND INFERENCE · CHAPTER 09

09 Ordinary Least Squares

The estimator, its algebra, and the boundary between association and causal interpretation.

9.1 Linear model

FORMULA — THE OLS ESTIMATOR

$$y = X\beta + \varepsilon \quad \cdot \quad \hat{\beta} = (X'X)^{-1}X'y$$

WHERE	y is the dependent vector, X the regressor matrix, β the coefficient vector, ε the error. The estimator holds when X has full column rank.
INTERPRETATION	OLS chooses coefficients minimizing the sum of squared residuals.
APPLICATION	Rank failure is not an abstraction: perfectly collinear dummy sets or duplicated factor columns will produce it in practice.

9.2 Simple regression

For $y_i = \beta_0 + \beta_1 x_i + u_i$, $\hat{\beta}_1 = \Sigma(x_i - \bar{x})(y_i - \bar{y}) / \Sigma(x_i - \bar{x})^2$ and $\hat{\beta}_0 = \bar{y} - \hat{\beta}_1\bar{x}$.

9.3 Interpretation

β_1 is the expected change in y associated with a one-unit change in x, holding other included regressors constant in multiple regression. Interpretation is causal only under stronger identification assumptions.

9.4 Residuals

Residual $e_i = y_i - \hat{y}_i$. OLS residuals sum to zero when an intercept is included and are orthogonal to included regressors in sample.

9.5 R-squared

$R^2 = 1 - SSE/SST$. It measures in-sample variance explained, not predictive power or causality. Adjusted R^2 penalizes additional regressors.

PART III · REGRESSION AND INFERENCE CHAPTER 10

10 Regression Diagnostics and Inference

Which assumption failures bias the coefficient, which only break the standard error, and which are economic rather than statistical problems.

10.1 Classical assumptions

- Linearity in parameters.
- Zero conditional mean for unbiased causal interpretation.

- Random or otherwise appropriate sampling.
- No perfect multicollinearity.
- Homoskedasticity for classical standard errors.
- No serial correlation where assumed.

10.2 Heteroskedasticity

When residual variance changes with predictors, OLS coefficients can remain unbiased under exogeneity but conventional standard errors are wrong. Use heteroskedasticity-robust errors when appropriate.

10.3 Autocorrelation

Serially correlated residuals are common in time series. Newey-West or other HAC errors, or explicit time-series models, may be required.

10.4 Multicollinearity

Highly correlated regressors inflate coefficient uncertainty and can make signs unstable. It does not inherently bias OLS.

10.5 Omitted variable bias

If an omitted determinant of y is correlated with an included regressor, the estimated coefficient can be biased. Economic reasoning matters as much as statistical diagnostics.

10.6 Endogeneity

Reverse causality, omitted variables, measurement error, or simultaneity can make regressors correlated with the error term. Instrumental variables and natural experiments are potential solutions when valid identification exists.

> **WATCH FOR**
>
> Robust standard errors fix inference, not identification. A HAC-corrected t-statistic on an endogenous regressor is a precisely estimated wrong number.

PART III · REGRESSION AND INFERENCE · CHAPTER 11

11 Multiple Regression, Transformations, and Nonlinearity

Logs, polynomials, interactions, and dummies — how to fit nonlinear effects while keeping a model linear in coefficients.

11.1 Multiple regression

Multiple regression estimates partial associations conditional on other variables. Coefficients depend on model specification and variable scaling.

11.2 Log models

TABLE 11.1 — INTERPRETING LOG SPECIFICATIONS

MODEL	INTERPRETATION
$\ln(y) = a + bx$	One-unit x increase ≈ 100b% change in y for small b
$y = a + b\ln(x)$	1% x increase ≈ b/100 units change in y
$\ln(y) = a + b\ln(x)$	b is elasticity: 1% x increase ≈ b% y change

Source: BIIRG Quantitative Research Division.

11.3 Polynomial terms

Adding x^2, x^3, or interactions allows nonlinear marginal effects while remaining linear in coefficients. Interpret derivatives, not coefficients in isolation.

11.4 Interaction terms

An interaction $x_1 \cdot x_2$ allows the effect of x_1 to depend on x_2. Regime-dependent factor returns are a common application.

11.5 Dummy variables

Indicator variables represent categories or regimes. Avoid the dummy-variable trap by omitting a reference category when an intercept is present.

PART IV

Time Series and Volatility

Once observations are ordered in time, most cross-sectional intuition breaks. Persistence, nonstationarity, volatility clustering, and structural breaks all demand explicit treatment.

12 Time-Series Foundations

Time ordering, autocorrelation, trend, seasonality, and the breaks that quietly invalidate a fitted parameter.

12.1 Time ordering

Time series violate the assumption that observations can be freely shuffled. Lags, trends, seasonality, persistence, and structural breaks matter.

12.2 Autocorrelation

The ACF measures correlation between a series and its lags. The PACF measures correlation at a lag after controlling for intermediate lags.

12.3 White noise

White noise has zero mean, constant variance, and no serial correlation under a basic definition. Financial residuals may still exhibit nonlinear dependence.

12.4 Trend and seasonality

Deterministic trend, stochastic trend, and seasonality require different treatments. Removing a trend mechanically can destroy economically meaningful information.

12.5 Structural breaks

Policy changes, crises, market microstructure, index methodology, and technology can alter model parameters. Stability tests and rolling estimates are essential.

PART IV · TIME SERIES AND VOLATILITY CHAPTER 13

13 AR, MA, ARIMA, and Forecasting

The workhorse univariate models, and the evaluation discipline that decides whether a forecast is worth anything.

13.1 AR models

AR(p): $y_t = c + \phi_1 y_{t-1} + \ldots + \phi_p y_{t-p} + \varepsilon_t$. Persistence and stationarity depend on parameter values.

13.2 MA models

MA(q): $y_t = \mu + \varepsilon_t + \theta_1 \varepsilon_{t-1} + \ldots + \theta_q \varepsilon_{t-q}$. Current values depend on current and past shocks.

13.3 ARMA and ARIMA

ARMA combines AR and MA for stationary series. ARIMA adds differencing to address integrated or nonstationary series.

13.4 Forecast evaluation

CORE CONCEPT — SIX EVALUATION CRITERIA

- Mean absolute error
- Root mean squared error
- Out-of-sample R^2
- Directional accuracy
- Economic utility
- Benchmark comparison

13.5 Expanding versus rolling windows

Expanding windows use all prior data; rolling windows discard old observations. Rolling estimation adapts faster to regime change but uses fewer observations.

PART IV · TIME SERIES AND VOLATILITY CHAPTER 14

14 Volatility Modeling

Historical, exponentially weighted, and conditional-variance estimates, and the difference between implied and realized.

14.1 Historical volatility

Annualized volatility is commonly the sample standard deviation of periodic returns × √(periods per year), assuming that scaling is reasonable.

14.2 EWMA

An exponentially weighted moving average gives greater weight to recent squared returns. A decay parameter controls persistence.

14.3 ARCH and GARCH

GARCH models conditional variance as a function of past shocks and past variance. GARCH(1,1) captures volatility clustering with relatively few parameters.

14.4 Implied versus realized

Implied volatility is inferred from option prices under a pricing model; realized volatility is measured from actual returns. The gap reflects volatility risk premium, expectations, supply and demand, and model effects.

14.5 Volatility regimes

Volatility is persistent but mean-reverting over long horizons and can jump abruptly. Forecast evaluation should emphasize stress periods as well as average error.

PART IV · TIME SERIES AND VOLATILITY CHAPTER 15

15 Stationarity, Unit Roots, and Cointegration

Why regressing levels on levels produces impressive nonsense, and what to do when theory implies a long-run relationship.

DEFINITION — COVARIANCE STATIONARITY

A covariance-stationary process has stable mean, variance, and autocovariance structure over time. Many statistical methods assume stationarity.

15.2 Unit roots

A unit-root process contains a stochastic trend. Regressing unrelated nonstationary levels can produce spurious high R^2 and significance.

15.3 Differencing

First differences can transform an I(1) series into a stationary series, but differencing removes long-run level information.

15.4 Cointegration

Two or more nonstationary series are cointegrated if a stationary linear combination exists, implying a long-run equilibrium relationship.

15.5 Error-correction model

An ECM combines short-run changes with adjustment toward a long-run cointegrating relationship. It is useful when theory implies equilibrium but short-run deviations occur.

PART V

Factors, Portfolios, and Risk

From single-factor beta to multifactor attribution, from the efficient frontier to the estimation error that destabilizes it, and the risk and performance measures a committee actually reads.

16 Factor Models and CAPM

Beta, alpha, and residual risk are all statements about a specified model. Change the model and all three change.

FORMULA — CAPM AND BETA

$$E[R_i] - R_f = \beta_i(E[R_m] - R_f) \quad \cdot \quad \beta_i = \mathrm{Cov}(R_i, R_m) / \mathrm{Var}(R_m)$$

WHERE	R_f is the risk-free rate and R_m the market return; beta measures systematic sensitivity to the market factor under the model.
INTERPRETATION	A beta above 1 implies greater sensitivity to market moves, not necessarily greater total risk.
APPLICATION	Report the estimation window, frequency, and market proxy alongside any beta; all three move the number.

16.3 Alpha

In a regression $R_i - R_f = \alpha + \beta(R_m - R_f) + \varepsilon$, alpha is average return unexplained by the included factor. It is model-dependent.

16.4 Systematic versus idiosyncratic risk

Factor risk is shared across securities; residual risk is security-specific under the model. Diversification can reduce idiosyncratic risk.

16.5 Factor model uses

CORE CONCEPT — SIX APPLICATIONS OF A FACTOR MODEL

- Risk decomposition
- Portfolio construction
- Performance attribution
- Signal neutralization
- Hedging
- Benchmarking

17 Multifactor and Quantamental Models

Systematic factors, macro state variables, and the requirement that every variable in a scoring model have a stated reason to be there.

17.1 Fama-French

The Fama-French framework extends market beta with empirically motivated factors such as size and value; later models add profitability and investment. Factor definitions and data construction must be read carefully.

17.2 Common systematic factors

TABLE 17.1 — SYSTEMATIC FACTORS AND TYPICAL PROXIES

FACTOR	TYPICAL PROXY
Value	Cheap versus expensive fundamentals
Momentum	Past relative performance
Quality	Profitability, balance-sheet strength
Size	Small versus large capitalization
Low volatility	Lower-risk securities
Carry	Yield or income relative to financing
Liquidity	Compensation for illiquidity

Source: BIIRG Quantitative Research Division.

17.3 Macro factors

Rates, inflation, growth, credit, commodities, FX, and liquidity can be used as explanatory or state variables. Because macro factors are often correlated, model stability and interpretation require care.

17.4 Quantamental research

Quantamental research combines systematic measurement with fundamental logic. A scoring model should specify why each variable should predict returns or risk, how it is normalized, and whether it adds incremental information.

KEY TAKEAWAY — FACTOR CROWDING

A factor can underperform when valuation becomes extreme, positioning is crowded, or market structure changes. Historical premium is not a guarantee.

PART V · FACTORS, PORTFOLIOS, AND RISK CHAPTER 18

18 Portfolio Theory and Mean-Variance Optimization

The frontier, the tangency portfolio, the constraint set, and the estimation error that makes an unconstrained optimizer untrustworthy.

18.1 Expected portfolio return

$E[R_p] = w'\mu$ and variance = $w'\Sigma w$. Both are functions of estimated inputs, and both inherit that estimation error.

18.2 Efficient frontier

The efficient frontier contains portfolios with maximum expected return for each variance level, or minimum variance for each expected return.

18.3 Global minimum variance

The GMV portfolio minimizes $w'\Sigma w$ subject to weights summing to one. It requires covariance estimates but not expected returns.

18.4 Tangency portfolio

Under standard assumptions, the tangency portfolio maximizes the Sharpe ratio relative to a risk-free asset.

18.5 Optimization constraints

- Long-only.
- Maximum position.
- Sector and asset-class limits.
- Turnover.
- Tracking error.
- Factor exposure.
- Liquidity.
- Shariah eligibility.

WATCH FOR — ESTIMATION ERROR

Mean-variance optimization is highly sensitive to expected-return estimates. Small changes can create extreme weights, motivating shrinkage, constraints, Bayesian methods, or risk-based portfolios.

PART V · FACTORS, PORTFOLIOS, AND RISK CHAPTER 19

19 Covariance Estimation and Robust Optimization

Methods that accept a little bias in exchange for portfolios that do not fall apart out of sample.

19.1 Sample covariance

Sample covariance is simple but noisy when the number of assets is large relative to observations.

19.2 Shrinkage

Shrinkage combines sample estimates with a structured target to reduce estimation error, trading some bias for lower variance.

19.3 Factor covariance

Factor models estimate covariance through common factor exposures plus idiosyncratic variance, reducing dimensionality.

19.4 Risk parity

Risk parity allocates so components contribute more evenly to portfolio risk. It does not guarantee economic diversification if assets share hidden factors.

19.5 Black-Litterman

Black-Litterman combines an equilibrium prior with investor views and confidence, helping stabilize expected-return inputs.

19.6 Robust optimization

Robust methods explicitly account for uncertainty in inputs, or constrain portfolios away from solutions that rely on fragile estimates.

20 Risk Measurement

Dispersion, drawdown, tail measures, benchmark-relative risk, and the decomposition of portfolio risk onto positions.

20.1 Volatility

Volatility measures dispersion, not direction. It treats upside and downside variation symmetrically.

20.2 Drawdown

Drawdown_t = portfolio value / prior peak − 1. Maximum drawdown is the worst peak-to-trough decline.

20.3 Value at Risk

VaR estimates a loss threshold over a horizon at a confidence level. Parametric, historical, and Monte Carlo methods make different assumptions.

20.4 Expected shortfall

Expected shortfall estimates average loss beyond the VaR threshold and captures tail severity better than VaR.

20.5 Tracking error

Tracking error is the standard deviation of active returns. It measures benchmark-relative risk.

20.6 Marginal and component risk

Marginal contribution measures how portfolio risk changes with a small position change; component contribution allocates total risk across positions.

20.7 Stress testing

Stress tests impose coherent shocks to rates, spreads, equities, commodities, FX, correlations, and volatility. Historical and hypothetical scenarios are complementary.

PART V · FACTORS, PORTFOLIOS, AND RISK CHAPTER 21

21 Performance Measurement and Attribution

Ratios, attribution, and the uncomfortable question of whether a track record is statistically distinguishable from luck.

TABLE 21.1 — PERFORMANCE RATIOS

RATIO	DEFINITION	WHAT IT MEASURES, AND WHAT TO STATE
Sharpe	$(R_p - R_f) / \sigma_p$	Excess return per unit of total volatility
Sortino	Excess return / downside deviation	Uses downside deviation; state the target or minimum acceptable return
Information ratio	$(R_p - R_b)$ / tracking error	Active return per unit of active risk
Calmar	Annualized return / max drawdown	Commonly absolute maximum drawdown; state period and convention

Source: BIIRG Quantitative Research Division; conventions vary by vendor.

21.5 Alpha and attribution

Factor attribution separates returns into factor exposures and residual or selection effects. Brinson-style attribution separates allocation, selection, and interaction for benchmarked portfolios.

> **INVESTMENT IMPLICATION**
>
> A high Sharpe over a short sample may be luck. Adjust confidence for sample size, serial correlation, non-normality, and selection bias before treating a track record as evidence of skill.

PART VI

Systematic Research and Practice

Signals, arbitrage, regimes, machine learning, and the honest backtest. This part is where a promising in-sample result either survives costs, bias, and governance, or is retired.

22 Signal Research and Cross-Sectional Strategies

The pipeline from raw feature to portfolio, and the diagnostics that separate a signal from an unintended sector bet.

22.1 Signal pipeline

1. Define universe.
2. Calculate raw feature.
3. Lag information to availability date.
4. Winsorize or clean if justified.
5. Normalize or rank.
6. Neutralize unwanted exposures.
7. Form portfolios.
8. Apply turnover and cost assumptions.
9. Evaluate out of sample.

22.2 Information coefficient

IC is the correlation between a signal and subsequent returns, often cross-sectionally. Rank IC uses rank correlation and is less sensitive to outliers.

22.3 Quantiles

Sort securities into deciles or quintiles and compare future returns. Monotonicity across buckets is often more persuasive than a single top-minus-bottom spread.

22.4 Signal decay

Measure performance by forecast horizon. A signal that decays in one day cannot support monthly rebalancing assumptions.

22.5 Neutralization

A stock-selection signal may unintentionally be a sector, size, beta, or country bet. Regression or bucket neutralization can isolate the intended exposure.

PART VI · SYSTEMATIC RESEARCH AND PRACTICE CHAPTER 23

23 Statistical Arbitrage and Mean Reversion

Convergence trades require a defensible equilibrium, a stable hedge ratio, and an honest account of what happens when neither holds.

23.1 Mean reversion

Mean-reversion strategies assume deviations from an equilibrium or relative relationship tend to reverse. The equilibrium must be economically and statistically defensible.

23.2 Pairs trading

Simple correlation is insufficient for pairs trading. Cointegration, stable hedge ratios, spread stationarity, transaction costs, and structural breaks matter.

23.3 Z-score signal

A spread z-score standardizes deviation from an estimated mean. Entry and exit thresholds should be validated out of sample and not optimized excessively.

23.4 Half-life

Mean-reversion half-life estimates how quickly a deviation decays. It can guide holding period but is itself estimated with error.

KEY TAKEAWAY — ARBITRAGE IS RARELY RISKLESS

Funding, borrow availability, model breakdown, crowding, execution, and convergence timing can turn apparent arbitrage into substantial risk.

PART VI · SYSTEMATIC RESEARCH AND PRACTICE
CHAPTER 24

24 Regime Detection and Cross-Asset Models

Rule-based and latent-state approaches, and the sample-size problem that conditioning always creates.

24.1 Regime variables

CORE CONCEPT — THE EIGHT STATE VARIABLES

Growth trend	Credit spreads
Inflation trend	Yield curve
Volatility	Dollar
Liquidity	Commodity trend

24.2 Rule-based regimes

Rule-based systems use transparent thresholds or standardized scores. They are interpretable and easy to audit but can be sensitive to threshold selection.

24.3 Hidden Markov models

HMMs infer latent states from observed data and transition probabilities. They can capture probabilistic regimes but require careful specification and are vulnerable to unstable state interpretation.

24.4 Clustering

K-means or hierarchical clustering can group observations into similar macro or market states without predefined labels. Scaling and feature choice dominate results.

24.5 Regime-conditional correlations

Estimate returns, volatility, correlations, and factor premia conditional on regime. Sample sizes shrink quickly, so uncertainty must be explicit.

25 Machine Learning for Investment Research

Where flexible models genuinely help, and the leakage and tuning failures that make them look like they do when they do not.

25.1 When machine learning helps

Machine learning is useful when relationships are nonlinear, interactions are numerous, or prediction matters more than coefficient interpretation. It does not solve weak data, leakage, nonstationarity, or poor economic design.

25.2 Regularization

TABLE 25.1 — REGULARIZATION METHODS

METHOD	EFFECT
Ridge / L2	Shrinks coefficients; useful with correlated predictors
Lasso / L1	Can set coefficients to zero; performs feature selection
Elastic Net	Combines L1 and L2

Source: BIIRG Quantitative Research Division.

25.3 Trees and ensembles

Decision trees capture nonlinear splits and interactions. Random forests reduce variance through bagging; boosting sequentially fits errors. Finance datasets require conservative tuning because signal-to-noise is low.

25.4 Cross-validation

Random k-fold cross-validation can leak future information in time series. Use walk-forward, expanding-window, rolling-window, or purged and embargoed methods appropriate to the prediction horizon.

25.5 Principal component analysis

PCA reduces dimensionality by finding orthogonal variance directions. Components maximize variance, not economic meaning.

25.6 Explainability

Feature importance and SHAP-like methods can help diagnose models, but explanation of a fitted model is not proof of causality.

PART VI · SYSTEMATIC RESEARCH AND PRACTICE CHAPTER 26

26 Backtesting and Research Bias

The five ways a historical simulation flatters a strategy, and the testing discipline that removes each one.

26.1 Look-ahead bias

Using information before it was actually available contaminates a backtest. Lag fundamentals to publication dates, not fiscal-period end dates.

26.2 Survivorship bias

Testing only current constituents excludes delisted and failed securities and can materially overstate performance.

26.3 Data snooping

Repeatedly modifying a strategy after viewing results makes the final backtest an in-sample product even if each individual test appears reasonable.

26.4 Overfitting

A strategy with many parameters can fit historical noise. Prefer parsimonious rules, broad parameter stability, and genuine holdout tests.

26.5 Walk-forward testing

Estimate using only information available up to each date, then test on future observations. This approximates the researcher's real-time information set.

BIIRG STANDARD — EVERY MODEL HAS A BASELINE

Equal weight, buy-and-hold, historical mean, a simple moving average, or an established factor model. Complexity must earn its place against a stated benchmark, not against nothing.

27 Transaction Costs, Liquidity, and Execution

Where gross Sharpe goes to die. Turnover, impact, and capacity decide whether a signal is a strategy.

27.1 Cost components

- Bid-ask spread.
- Commissions and fees.
- Market impact.
- Slippage.
- Borrow cost.
- Funding cost.
- Taxes where applicable.

27.2 Turnover

Turnover is a major determinant of implementability. A high gross Sharpe can disappear after realistic costs.

27.3 Market impact

Impact increases with order size relative to market liquidity and urgency. Backtests should constrain participation and avoid assuming unlimited execution at closing prices.

27.4 Capacity

Capacity is the capital a strategy can deploy before market impact materially erodes returns. It depends on liquidity, turnover, concentration, and signal decay.

27.5 Rebalancing

Calendar, threshold, and optimization-based rebalancing trade off signal freshness against costs. Model the rule explicitly.

28 Model Validation, Monitoring, and Governance

Six validation layers, the monitoring set that distinguishes drawdown from failure, and kill criteria written before deployment.

28.1 Validation layers

TABLE 28.1 — THE SIX VALIDATION LAYERS

LAYER	QUESTION
Conceptual	Does the model make economic and statistical sense?
Data	Are inputs accurate, timely, and unbiased?
Implementation	Does the code match the methodology?
Statistical	Is performance robust and significant?
Economic	Does it survive costs and constraints?
Stability	Does it work across periods and regimes?

Source: BIIRG Quantitative Research Division.

28.2 Monitoring

CORE CONCEPT — THE LIVE MONITORING SET

Forecast error	Residual behavior	Data quality
Signal IC	Parameter drift	Capacity
Turnover	Drawdown	Exposure drift

28.3 Model decay

A model can decay because the economic relationship changes, competitors arbitrage it, market structure evolves, or input data change. Monitoring should distinguish temporary drawdown from structural failure.

28.4 Governance

Maintain model owner, purpose, version, inputs, assumptions, limitations, validation record, approval status, and change log. Research should be reproducible by another analyst.

BIIRG STANDARD — KILL CRITERIA, REQUIRED BEFORE DEPLOYMENT

Define the conditions for reducing or retiring a model: persistent out-of-sample failure, data break, exposure drift, economic mechanism invalidation, excessive turnover or cost, or unacceptable tail behavior.

29 BIIRG Quant Workflow and Practical Case

The division workflow, the research-note structure, and two projects every analyst completes before independent work.

29.1 Division workflow

1. Translate an investment question into a falsifiable hypothesis.
2. Write a data dictionary before modeling.
3. Create a clean point-in-time dataset.
4. Perform exploratory statistics and plots.
5. Specify benchmark and model.
6. Run in-sample estimation only for development.
7. Reserve validation data.
8. Backtest with costs, lags, and constraints.
9. Run robustness and regime tests.
10. Document results, failures, and limitations.
11. Translate output into portfolio and risk use.
12. Monitor live performance without rewriting history.

29.2 Research note structure

CORE CONCEPT — THE THIRTEEN SECTIONS

Research question	Results	Risk
Economic rationale	Robustness	Portfolio application
Data and universe	Backtest	Limitations
Methodology	Costs and capacity	Monitoring plan
Model specification		

29.3 Practical case: macro regime allocation model

Build a hypothetical monthly model allocating across U.S. equities, gold, broad commodities, sukuk and fixed income, and cash using standardized growth, inflation, liquidity, and volatility indicators. The purpose is to test whether regime-conditioned allocation improves risk-adjusted performance versus a static benchmark.

29.4 Required tasks

1. Define each indicator and publication lag.
2. Standardize using only historical information.
3. Create transparent regime rules.
4. Estimate regime-conditioned asset returns and covariance.
5. Set portfolio constraints and a static benchmark.
6. Run a walk-forward backtest.
7. Include transaction costs and monthly rebalancing.
8. Report CAGR, volatility, Sharpe, Sortino, max drawdown, Calmar, turnover, and tracking error.
9. Test alternative thresholds and subperiods.
10. Run a placebo or randomized signal test.
11. Explain whether improvement comes from return timing, volatility reduction, or both.
12. State conditions under which the model should be rejected.

29.5 Second case: volatility forecast

Compare 20-day rolling volatility, EWMA, and GARCH(1,1) forecasts for a liquid index. Evaluate out-of-sample forecast error and whether improved forecasts create economic value in volatility targeting.

> **KEY TAKEAWAY — THE FINAL STANDARD**
>
> The purpose of quantitative research is not to make investing look mathematical. It is to impose measurement, falsifiability, consistency, and risk discipline on investment decisions.

REFERENCE MATERIAL

Final Analyst Checklist

Run before any quantitative result is circulated as research.

01	What is the economic hypothesis?
02	Is the dataset point-in-time?
03	Are returns and frequencies calculated correctly?
04	Is the model appropriate for the data-generating process?
05	Are inference assumptions valid?
06	Did I test out of sample?
07	Did I account for multiple testing?
08	Are transaction costs realistic?
09	Is performance robust across subperiods and regimes?
10	Are portfolio exposures understood?
11	What is the model's failure mode?
12	Can another analyst reproduce the result?

APPENDIX A

Core Formula Sheet

CONCEPT	FORMULA	NOTE
Simple return	$(P_t - P_{t-1} + D_t) / P_{t-1}$	Aggregates cross-sectionally
Log return	$\ln(P_t / P_{t-1})$	Adds across time
Cumulative return	$\Pi(1 + R_t) - 1$	Not the arithmetic mean
Variance	$E[(X - E[X])^2]$	Squared dispersion
Covariance	$E[(X - E[X])(Y - E[Y])]$	Unbounded, scale-dependent
Correlation	$\mathrm{Cov}(X,Y) / (\sigma_X \sigma_Y)$	Linear dependence only
Portfolio return	$w'\mu$	Estimated inputs
Portfolio variance	$w'\Sigma w$	Σ must be positive semidefinite
OLS	$(X'X)^{-1}X'y$	Requires full column rank
R-squared	$1 - SSE/SST$	In-sample fit only
Beta	$\mathrm{Cov}(R_i,R_m) / \mathrm{Var}(R_m)$	State window and proxy
CAPM	$E[R_i] - R_f = \beta(E[R_m] - R_f)$	Single-factor model
Sharpe	$(R_p - R_f) / \sigma_p$	Total-risk basis
Information ratio	Active return / tracking error	Benchmark-relative
Value at Risk	Loss quantile at horizon and confidence	Method-dependent
Expected shortfall	Mean loss beyond VaR	Tail severity
Max drawdown	min(Value / previous peak – 1)	Path-dependent
Annualized volatility	Periodic SD × √(periods per year)	Assumes no autocorrelation
Z-score	$(x - \mu) / \sigma$	Inherits estimated moments

Definitions vary by vendor and source. Apply the source methodology and the standard BIIRG has adopted, not a formula recalled from this sheet.

APPENDIX B

Glossary

Alpha	Return unexplained by the included factor model.
Autocorrelation	Correlation of a series with its own lag.
Backtest	Historical simulation of a strategy using specified rules.
Beta	Sensitivity to a factor, commonly the market.
Cointegration	Stationary linear combination of nonstationary series.
Convex optimization	Optimization class with favorable global-solution properties.
Cross-validation	Procedure for evaluating out-of-sample generalization.
Drawdown	Decline from prior portfolio peak.
Endogeneity	Regressor correlated with the model error.
Expected shortfall	Average loss conditional on exceeding VaR.
Factor	Common systematic return driver.
GARCH	Conditional volatility model.
Heteroskedasticity	Nonconstant residual variance.
Information coefficient	Correlation between signal and subsequent return.
Look-ahead bias	Use of future information in a historical test.
Multicollinearity	Strong linear dependence among regressors.
OLS	Least-squares linear regression estimator.
Overfitting	Fitting noise rather than generalizable signal.
PCA	Orthogonal dimensionality-reduction method.
Regularization	Penalty used to constrain model complexity.
Residual	Observed minus fitted value.
Stationarity	Stable time-series properties under a specified definition.
Survivorship bias	Excluding failed or delisted members from a historical universe.
Tracking error	Volatility of active returns.
VaR	Loss threshold at a specified probability and horizon.
Walk-forward	Sequential historical training and future testing.
White noise	Uncorrelated random process under the basic definition.

APPENDIX C

Technical Stack, Study Sequence, and References

Recommended technical stack

Excel	Transparent audit models, data checks, portfolio calculations, IC-facing outputs.
Python	pandas, NumPy, SciPy, statsmodels, scikit-learn, matplotlib, and optimization libraries as appropriate.
SQL	Querying structured research databases.
Version control	Git for reproducible code and change history.
Point-in-time data	Market and fundamental databases with as-reported history; required for serious backtesting.

Recommended study sequence

New analysts: Chapters 1 to 11 first, then 12 to 17, then 18 to 21. Strategy researchers should continue through 22 to 28 before independent project work. Every analyst should complete the regime-allocation case and one factor or volatility project.

Selected references

CFA Institute curriculum: quantitative methods, portfolio management, factor models, derivatives, and performance measurement.

Fama, E. F. and French, K. R.: factor-model research and data library.

Markowitz, H.: *Portfolio Selection* and mean-variance portfolio theory.

Engle, R.: ARCH volatility modeling; Bollerslev, T.: GARCH.

Box, Jenkins et al.: time-series analysis and ARIMA methodology.

Hamilton, J.: *Time Series Analysis* and regime-switching methods.

Newey, W. and West, K.: heteroskedasticity and autocorrelation-consistent covariance estimation.

Ledoit, O. and Wolf, M.: covariance shrinkage.

Black, F. and Litterman, R.: global portfolio optimization and equilibrium views.

Bailey, Borwein, López de Prado, Zhu: research on backtest overfitting and selection bias.

Related volumes

Volume I for the equity fundamentals behind cross-sectional signals, Volume II for fixed income and sukuk instruments, Volume III for the appraisal-based data problems referenced in Chapters 7 and 17, and Volume IV for the macro state variables used in the regime work of Chapter 24.

BIIRG

BARUCH ISLAMIC INVESTMENT RESEARCH GROUP

Analyst Training Series · Volume V of V
Quantitative Research · 2026 Edition

EDUCATIONAL USE ONLY