Risk Analytics

Machine Learning and Optimization
for Data-Driven Decision Making

Fernando S. Oliveira

Draft version — May 27, 2026

Chapter 10

Time Series Foundations for Risk

Chapter 10

Time Series Foundations for Risk

10.1 Introduction: Risk as an Evolving Management Problem

Risk management becomes more difficult when the source of risk is not a single uncertain outcome, but a variable that changes month after month. A procurement exposure may look acceptable today and become material after several adverse price movements. A supplier may appear reliable until delivery delays begin to accumulate. A market may look calm until volatility rises and the same exposure suddenly produces a much wider range of possible outcomes. In these situations, the relevant question is not only "What is the risk now?" but also "How is the risk evolving, how persistent is it, and what future paths should the organization be prepared to face?"

This chapter introduces the time-series foundations needed to answer those questions. It connects the statistical language of time series to the managerial language of monitoring, early warning, volatility, persistence, and contingency planning. The running case is copper procurement risk. A manufacturer buys copper every month and is exposed to changes in input prices. The firm does not need a complete commodity-trading model to ask useful questions. It needs to know whether copper prices are high relative to recent history, whether monthly price changes are becoming more volatile, whether adverse changes tend to persist, and what several adverse months would imply for procurement planning.

The chapter develops a practical workflow. The analyst first distinguishes the price level from the monthly price change. The price level describes the economic exposure; the price change describes the short-term risk movement. Moving averages summarize recent direction. Rolling standard deviations summarize current volatility. Non-stationarity diagnostics explain why price levels should not automatically be treated as stable around an old average.

Autocorrelation shows whether monthly changes tend to continue, reverse, or fade. Rolling VaR and CVaR then summarize the recent adverse tail of monthly price changes. Finally, simulation and stress paths convert one-period indicators into multi-period scenarios.

The chapter contributes to the architecture of the book in three ways. First, it extends VaR and CVaR from static loss distributions to rolling, time-updated risk indicators. Second, it prepares the reader for scenario trees, event trees, and sequential decision models by showing how future paths can be generated from time-series information. Third, it clarifies why non-stationarity, persistence, and changing volatility matter for management: historical data are useful only when they remain informative about the current decision environment.

10.2 Running Case: Copper Procurement Risk

The running case uses monthly commodity price data from the World Bank Commodity Price Data, commonly known as the Pink Sheet [18]. The workbook used for this chapter contains monthly prices in nominal US dollars from 1960M01 through 2025M12. The empirical illustrations focus on 2015–2025 so that the analysis reflects a recent period with visible commodity-price movements, including the post-pandemic commodity cycle.

The managerial problem is a procurement decision under evolving price risk. Consider a manufacturer that buys copper as a critical input. The firm sells finished products under contracts that adjust only gradually. A sudden increase in copper prices therefore affects margins before the firm can fully pass the cost on to customers. If the price increase is short-lived, the firm may absorb it. If the increase persists for several months, the problem becomes strategic: the firm may need to hedge, adjust inventory, renegotiate contracts, revise budgets, or escalate the exposure to senior management.

This is why the case is framed as a time-series risk problem rather than a one-period loss problem. The procurement director is not only interested in the latest copper price. The monitoring system must show whether the current price is high relative to the recent baseline, whether monthly changes are becoming more volatile, whether adverse changes tend to persist or reverse, and what would happen if adverse changes continued for several months.

The statistical analysis focuses on the price level and the monthly price change. The one-month horizon is used for rolling VaR and CVaR because the firm purchases monthly and needs a current short-horizon indicator of adverse price movements. The six-month horizon is used for scenario analysis because procurement, hedging, inventory, and contract responses usually require decisions over several months. The same structure could be repeated for three, twelve, or twenty-four months if the decision horizon required it.

Figure 10.1 plots the copper price and its 12-month moving average. The figure introduces the central managerial tension of the chapter. The price level determines the size of the procurement bill, but the path of prices determines whether the organization is facing a temporary disturbance, a persistent cost squeeze, or a possible change in market regime.

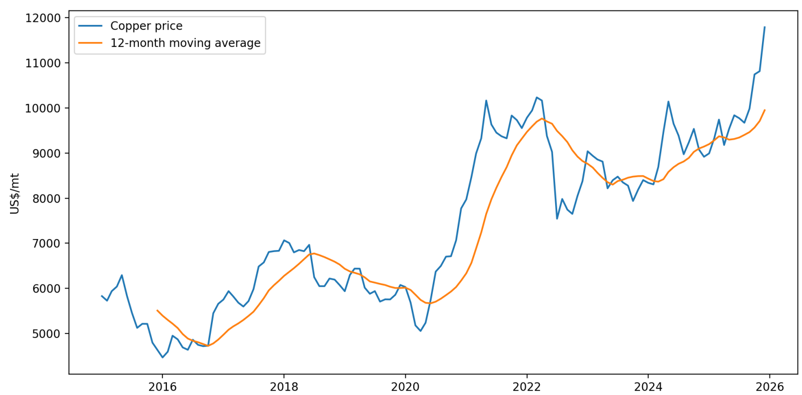

Figure 10.1: Copper price and 12-month moving average, World Bank Pink Sheet, 2015–2025.

The moving average is not a final forecast. It is a smoothing device that helps the manager separate short-term noise from the recent baseline. If the current price is far above its recent average, the procurement discussion changes. The relevant question becomes whether the deviation is likely to fade, persist, or mark a new environment.

10.3 Defining the Risk Variable: Prices, Changes, and Exposure

The first analytical task is to define the risk variable. In the copper case, there are two related but distinct objects. The price level, denoted P_t, determines the current cost per metric ton. The monthly price change, denoted R_t, determines the short-term cost movement from one month to the next:

$$R_t = \frac{P_t - P_{t-1}}{P_{t-1}}. \tag{10.1}$$

This distinction is essential. The price level answers the exposure question: "How expensive is copper now?" The monthly price change answers the risk-dynamics question: "How much has the exposure moved, and how volatile are those movements?" For a copper buyer, positive values of R_t are adverse because they increase input costs. Negative values are favourable because the input becomes cheaper. The relevant adverse tail is therefore the upper tail of monthly price changes.

The chapter focuses on prices and monthly price changes directly. This keeps the example general and avoids building the analysis around an artificial loss function. If a manager later wants a monetary number, a percentage price movement can be multiplied by the relevant purchase quantity and current price. The statistical object, however, remains the monthly price change.

The use of monthly price changes also improves the statistical interpretation. Commodity price levels often move over time and may not fluctuate around a stable long-run mean. By contrast, monthly changes are usually more suitable for estimating recent volatility, autocorrelation, and empirical tail measures.

For multi-period scenarios, the chapter converts simulated or stress-tested price changes back into a price path. This is done recursively:

$$P_{t+h} = P_{t+h-1}(1 + R_{t+h}), \qquad h = 1, \ldots, H. \tag{10.2}$$

The recursion is important because dynamic risk is path-dependent. A single 8% monthly price increase may be manageable. Several consecutive increases may create a strategic problem. The price after six months depends on the sequence of monthly changes, not only on the average change.

The same logic applies beyond copper. A credit-risk analyst may monitor monthly default rates; an operations manager may monitor delivery-delay rates; an energy analyst may monitor load, price, or outage rates; and a macro-risk analyst may monitor inflation, exchange rates, or production indicators. The general principle is unchanged: the analyst must define the time-indexed variable, transform it into a decision-relevant risk measure, and then ask how its distribution changes through time.

10.4 Smoothing Risk Signals: Moving Averages and Rolling Volatility

Raw time-series observations are noisy. A single monthly price change may reflect temporary news, reporting noise, inventory adjustment, or a genuine shift in market conditions. Managers therefore need indicators that are stable enough to avoid overreaction but responsive enough to detect deterioration. Smoothing methods provide this bridge between raw data and management interpretation [13, 11].

The simplest smoothing device is the moving average. For a window of w months, the rolling mean of monthly price changes is:

$$\bar{R}_{t,w} = \frac{1}{w}\sum_{j=0}^{w-1} R_{t-j}. \tag{10.3}$$

The notation $\bar{R}_{t,w}$ means the average value of R at time t, computed using the most recent w observations. For example, $\bar{R}_{t,12}$ is the average monthly price change over the most recent twelve months. Operationally, the calculation is simple: each month the analyst drops the oldest observation, adds the newest observation, and recomputes the average over the updated window. The moving average is not a forecast by itself. It is a recent baseline against which the latest observation can be compared.

The same window can be used to estimate recent volatility:

$$s_{t,w} = \sqrt{\frac{1}{w-1}\sum_{j=0}^{w-1}\left(R_{t-j} - \bar{R}_{t,w}\right)^2}. \tag{10.4}$$

The statistic $s_{t,w}$ is the rolling standard deviation. In the copper case, it measures the recent dispersion of monthly input-cost changes around their rolling mean. A rising $s_{t,w}$ means that the current range of plausible monthly cost movements is widening. This is often more important for risk management than the average itself, because a stable mean can coexist with increasing uncertainty.

The window length w is a modelling and governance choice. A short window reacts quickly to new market conditions but can be dominated by a few observations. A long window is more stable but may react slowly after a regime change. The choice should therefore be linked to the management decision. A 12-month window is useful for summarizing recent annual behavior and communicating the current baseline. A 36-month window gives more observations for empirical tail measures such as VaR and CVaR. Neither choice is universally correct; each reflects a trade-off between responsiveness and stability.

Figure 10.2 applies these moving-window indicators to copper. The monthly price changes show the raw signal. The 12-month moving average shows recent average cost pressure. The rolling standard deviation shows the current volatility of monthly price changes.

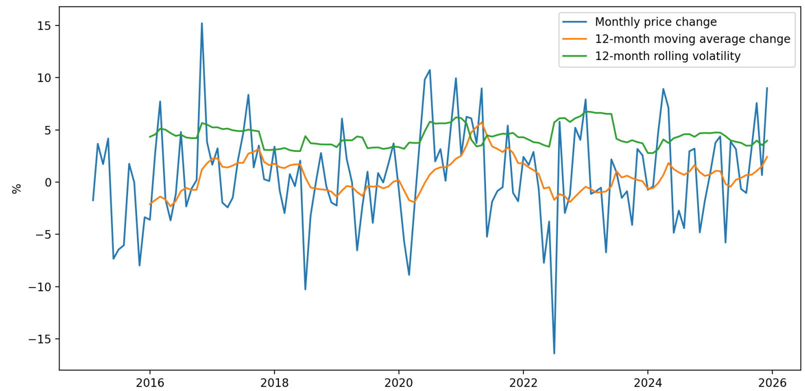

Figure 10.2: Copper monthly price changes, 12-month moving average change, and rolling volatility.

Copper example. At the end of the sample, copper has a latest price of US$11,785.25 per metric ton and a 12-month moving average price of US$9,947.31 per metric ton. The latest monthly price change is 9.0%, the 12-month moving average monthly change is 2.4%, and the 12-month rolling monthly volatility is 4.0%. The firm is therefore facing both a high price level and recent upward cost pressure. The rolling volatility indicates that future monthly movements remain materially uncertain.

Table 10.1: Smoothed copper indicators at the end of the sample

Indicator	Value
Latest copper price	US$11,785.25/mt
12-month moving average price	US$9,947.31/mt
Latest monthly price change	9.0%
12-month moving average monthly change	2.4%
12-month rolling monthly volatility	4.0%
36-month rolling mean monthly change	1.0%
36-month rolling monthly volatility	4.2%

The table should not be read as a mechanical trading or procurement rule. It is a monitoring summary. The latest monthly increase is well above the 12-month average change, so the current month deserves attention. The price level is also above the 12-month moving average price, so the exposure is already expensive. The rolling volatility is material, so the next monthly movement should not be treated as predictable simply because the recent mean is positive.

Smoothing has limitations. It can hide turning points, especially when the window is long. It can also create a false sense of stability if the underlying process has changed. For that reason, smoothing should be combined with the non-stationarity diagnostics in the next section, the persistence analysis in Section 10.6, and the rolling tail measures in Section 10.7. A good risk dashboard does not ask whether the moving average is right. It asks what the moving average, volatility, persistence, and tail measures jointly imply for action.

10.5 Non-Stationarity, Differencing, and Risk Interpretation

Sections 10.3 and 10.4 made two practical choices. First, the chapter distinguished between the copper price level, which determines the current procurement bill, and the monthly price change, which measures the short-term movement in that bill. Second, it used moving averages and rolling volatility to summarize recent behaviour. This section explains why those choices are not merely cosmetic. They are necessary because price levels and price changes often have different time-series properties.

A stationary time series fluctuates around a stable statistical structure. In simple terms, its mean, variance, and dependence pattern do not drift systematically through time. A non-stationary time series does not have such a stable reference structure. Its level may drift, shocks may accumulate, or structural breaks may shift the process to a new environment. This distinction matters for risk analysis because many risk summaries are interpreted relative to a baseline. If the baseline itself is moving, then a historical average can become a poor guide to current exposure [1, 2, 10, 3]. Non-stationarity is also one reason why simple correlations and regressions based on time-series levels can become misleading if the analyst ignores the underlying time-series structure [9, 5, 3].

Commodity price levels are often non-stationary over economically relevant samples. Copper prices can move to a new level because of changes in global demand, mine supply, inventories, exchange rates, energy costs, regulation, financing conditions, or expectations about future industrial activity. In a risk-management setting, this means that a high price should not automatically be interpreted as a temporary deviation from an old long-run average. It may instead be the new starting point from which future procurement exposure is measured.

A simple representation of a non-stationary price level is a random walk with drift:

$$p_t = p_{t-1} + g + u_t, \qquad p_t = \log(P_t). \tag{10.5}$$

Here, p_t is the logarithm of the copper price, g is an average monthly drift, and u_t is a new shock. Equation (10.5) is non-stationary because shocks accumulate in the level. A positive shock raises p_t, and the higher value becomes the starting point for the next period. There is no fixed long-run mean to which the price level must immediately return.

Taking first differences gives:

$$\Delta p_t = p_t - p_{t-1} = g + u_t. \tag{10.6}$$

For small changes, Δp_t is close to the percentage price change R_t defined in Equation (10.1). This explains the practical strategy used in the chapter. The price level P_t remains economically meaningful because it determines the procurement bill. The monthly price change R_t, however, is usually more suitable for estimating recent volatility, autocorrelation, VaR, CVaR, and scenario inputs.

The copper data illustrate the distinction. Figure 10.3 shows the autocorrelation of log copper prices over 2015–2025, omitting lag 0 because lag 0 is always equal to one by definition and does not provide information about persistence. The autocorrelations are very high and decay slowly. This pattern is typical of a persistent, non-stationary, or near-non-stationary price level. If a manager treated the price level as a stable distribution around a fixed historical mean, the resulting risk assessment could understate the possibility that the firm is already operating in a new cost environment.

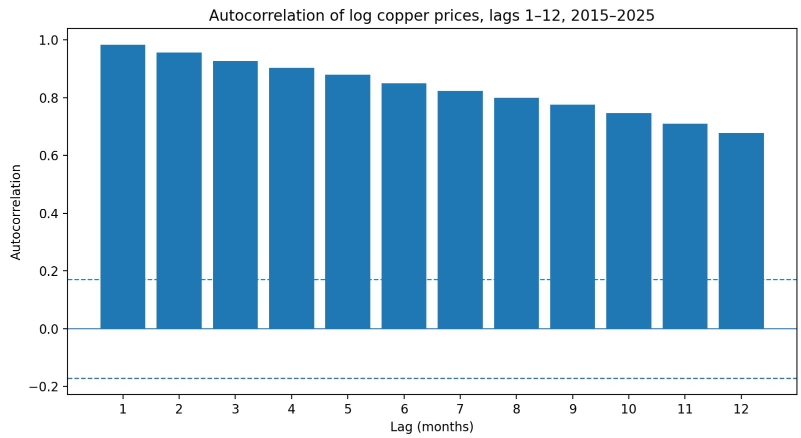

Figure 10.3: Autocorrelation of log copper prices, lags 1–12, 2015–2025. Lag 0 is omitted because it is always equal to one by definition.

Table 10.2 compares the price-level diagnostic with the price-change diagnostic. The table has been designed as a risk-management diagnostic rather than as a compact econometric output. It answers three questions: what was estimated, what the diagnostic implies statistically, and what the result means for the procurement decision.

Table 10.2: Copper price-level and price-change diagnostics, 2015–2025

Variable	Lag coefficient	ACF(1)	Risk-management interpretation
Log price level p_t	0.996	0.983	Highly persistent level. Do not treat the long-run average price as a stable procurement baseline. Current exposure should be measured from the current price level.
Monthly change R_t	0.254	0.251	Moderate short-run persistence. Use monthly changes for rolling volatility, VaR, CVaR, persistence checks, and short-horizon scenario inputs.

The purpose of Table 10.2 is to support a disciplined modelling choice. The price level is the exposure variable, but it is highly persistent and can shift to a new range. The monthly change is the short-horizon risk variable because it measures the percentage movement that changes the next procurement bill. This is why the chapter estimates rolling volatility, VaR, CVaR, and scenario inputs using monthly changes rather than price levels.

This distinction connects directly to the smoothing methods in Section 10.4. A moving average of the price level is useful as a visual baseline, but it should not be interpreted as a stable equilibrium price. A moving average of monthly price changes, by contrast, summarizes recent cost pressure. A rolling standard deviation of monthly price changes summarizes recent uncertainty. The analyst is therefore smoothing the variable that is closer to being stable over the short horizon, while still reporting the current price level because that is where the money is at risk.

The risk-management implication is strongest for scenario design. Multi-month scenarios should start from the current price and evolve recursively through simulated or stressed monthly changes, rather than assuming an immediate return to an old average. The practical message is simple: use prices to understand the level of exposure, and use price changes to estimate recent volatility, persistence, VaR, CVaR, and short-horizon scenario inputs. This distinction avoids treating a non-stationary price level as if it were a stable distribution while preserving the direct managerial interpretation of the procurement bill.

10.6 Autocorrelation, Persistence, and Mean Reversion

Sections 10.3–10.5 established the basic modelling sequence used in this chapter. Price levels determine the economic exposure, but monthly price changes are the more appropriate object for short-horizon risk measurement. Moving averages and rolling standard deviations then summarize recent mean pressure and volatility. The next step is to ask whether monthly changes are independent disturbances or whether they contain persistence. This matters because procurement risk is not only about the size of the next monthly increase. It is also about whether an adverse increase is likely to be followed by further adverse increases.

Autocorrelation measures the relationship between a time series and its lagged values [1, 2, 10]. The first-order autocorrelation, denoted ACF(1), is the correlation between R_t and R_{t-1}. A positive ACF(1) means that high monthly changes tend to be followed by high monthly changes and low changes tend to be followed by low changes. A negative ACF(1) means that changes tend to reverse. A value close to zero means weak linear dependence from one month to the next.

For risk management, autocorrelation is a bridge between the one-period indicators in Sections 10.4 and 10.5 and the multi-period scenarios developed later in the chapter. If changes are close to independent, a one-month VaR or CVaR is mainly a short-term monitoring statistic. If changes are positively autocorrelated, a one-month adverse movement may be the beginning of a more persistent cost episode. If changes are negatively autocorrelated, a single adverse movement is more likely to be partly corrected. Evidence on mean reversion in financial and commodity-related series is mixed and horizon-dependent, which is why the chapter uses autocorrelation as a diagnostic rather than as a trading rule [14, 7].

A simple autoregressive model of order one, or AR(1), represents persistence as:

$$R_t = a + bR_{t-1} + \varepsilon_t. \tag{10.7}$$

The term "AR(1)" means that the current value of the series is modelled using one lag of itself. In this case, the current monthly price change depends on the previous monthly price change plus a new shock. The coefficient b measures persistence. If $b > 0$, positive changes tend to be followed by positive changes. If $b < 0$, changes tend to reverse. If b is close to zero, the previous month contains little linear information about the current month.

If $|b| < 1$, the AR(1) model has a stable mean. Taking expectations in Equation (10.7) and using $E[R_t] = E[R_{t-1}] = \mu$ gives:

$$\mu = a + b\mu.$$

Solving for μ gives:

$$\mu = \frac{a}{1-b}. \tag{10.8}$$

Substituting $a = (1-b)\mu$ back into Equation (10.7) gives:

$$R_t = (1-b)\mu + bR_{t-1} + \varepsilon_t.$$

Subtracting μ from both sides gives the mean-reverting form:

$$R_t - \mu = b(R_{t-1} - \mu) + \varepsilon_t. \tag{10.9}$$

This equation should be interpreted carefully. It does not say that the copper price level must return to an old historical average. Section 10.5 showed that copper price levels are highly persistent and may be non-stationary. Equation (10.9) applies to monthly price changes. It says that, if the monthly-change process is stable and $|b| < 1$, deviations of monthly changes from their recent mean tend to shrink over time.

For a stable AR(1) process, the magnitude of b provides a compact measure of persistence. If b is close to zero, deviations from the recent mean fade quickly. If b is positive and large, adverse price changes are more likely to carry into the following months. If $b < 0$, price changes tend to reverse. This interpretation is useful for management because it links a statistical coefficient to the planning question of whether a price movement is likely to remain a one-month disturbance or become a multi-month exposure.

Figure 10.4 shows the autocorrelation of copper monthly price changes over 2015–2025. Lag 0 is omitted because it is always equal to one by definition and does not provide information about persistence. The first-order autocorrelation is approximately 0.25. This is not strong, but it is large enough to caution against treating monthly changes as completely independent.

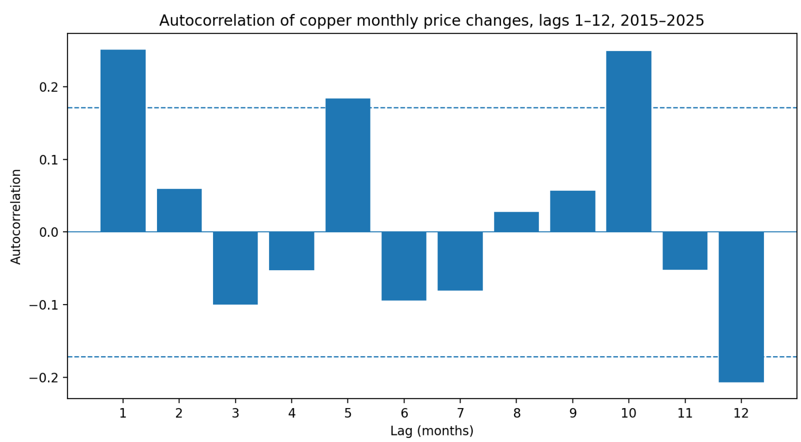

Figure 10.4: Autocorrelation of copper monthly price changes, lags 1–12, 2015–2025. Lag 0 is omitted because it is always equal to one by definition.

Copper example. The estimated AR(1) coefficient for monthly copper price changes is approximately $b = 0.254$, close to the sample ACF(1) of 0.251 reported in Table 10.2. The managerial interpretation is moderate persistence, not strong momentum. One adverse month should not be mechanically extrapolated into a long trend, but it should not be ignored when the price level, rolling mean, volatility, and tail indicators are also deteriorating.

This interpretation connects directly to risk management. For a monthly buyer, persistence determines whether a price shock is mainly an accounting problem for the current purchase or an operational problem for the procurement plan. A low-persistence shock may be handled through monitoring and short-term budget adjustment. A persistent shock may require escalation, supplier renegotiation, hedging, inventory changes, or customer price pass-through. A structural change is more serious still, because the old historical baseline may no longer be the right benchmark.

Table 10.3: Persistence, mean reversion, and managerial interpretation

Pattern	Risk interpretation	Managerial implication
Low autocorrelation	Shocks are mostly isolated	Monitor; avoid overreaction
Positive autocorrelation	Shocks may continue across purchase cycles	Consider hedging, inventory, or escalation
Negative autocorrelation	Shocks tend to reverse	Avoid extrapolating one bad month too far
High persistence	Adverse states decay slowly	Treat risk as multi-period exposure
Structural break	Old baseline may not apply	Recalibrate indicators and add stress scenarios

The section therefore adds two management questions to the workflow. First, within the current regime, do shocks tend to fade, reverse, or continue? Second, is the current regime still comparable to the historical data used for estimation? The first question is about persistence. The second is about non-stationarity and structural change. A good time-series risk report should answer both before converting rolling VaR and CVaR into scenario paths.

10.7 Rolling VaR and CVaR for Time-Series Risk

The previous sections established the modelling sequence used in this chapter. Section 10.3 defined the monthly copper price change, R_t, as the short-term risk variable. Section 10.4 summarized recent mean behaviour and volatility. Section 10.5 explained why the non-stationary price level should not be treated as a stable distribution. Section 10.6 then used autocorrelation to assess whether monthly changes show persistence. Rolling VaR and CVaR build directly on this workflow by translating the recent distribution of R_t

into a current downside-risk indicator.

VaR means Value at Risk in this chapter; variance is written explicitly as $\mathrm{Var}(\cdot)$. For the copper buyer, the adverse variable is the monthly price change R_t. Since price increases are adverse, the relevant tail is the upper tail. Negative values of R_t are not ignored; they are favourable from the buyer's perspective because copper has become cheaper. If the decision problem were about a copper producer, the adverse tail would instead be the lower tail because price decreases reduce revenue.

Let the rolling window at time t be:

$$W_{t,w} = \{R_{t-w+1}, \dots, R_t\}. \tag{10.10}$$

The empirical rolling VaR at confidence level α is the empirical α-quantile of the observations in the current window:

$$\mathrm{VaR}^R_{\alpha,t} = Q_\alpha(W_{t,w}). \tag{10.11}$$

The function Q_α denotes the empirical quantile. For $\alpha = 0.95$, it is the value such that approximately 95% of observations in the rolling window are below it and 5% are above it. Because the buyer is harmed by price increases, this upper-tail quantile is the relevant Value at Risk. The superscript R means that the measure is expressed as a monthly price change.

The empirical rolling CVaR is the average of the observations in the current window that are at least as large as the VaR threshold:

$$\mathrm{CVaR}^R_{\alpha,t} = \text{average}\left\{R_j \in W_{t,w} : R_j \geq \mathrm{VaR}^R_{\alpha,t}\right\}. \tag{10.12}$$

VaR is the adverse threshold. CVaR is the average outcome beyond that threshold. In continuous distributions this corresponds to the usual expected-shortfall interpretation [12, 15, 16]. In finite empirical samples, special care may be required when there are ties at the VaR point, but the tail-average definition is sufficient for the rolling management indicator used here.

The notation avoids hats to keep the chapter readable. The meaning is that VaR and CVaR are empirical rolling estimates recalculated from the current window. There is no simple exact scalar recursion for a rolling quantile comparable to a moving average. A rolling quantile depends on the ordering of all observations in the window. Operationally, the window is updated by dropping the oldest observation and adding the newest one; VaR and CVaR are then recomputed from the updated window.

Copper example. At the end of the copper sample, using a 36-month rolling window, the empirical 95% monthly VaR is 8.2%, and the empirical 95% monthly CVaR is 9.0%. These are price-change measures, not dollar losses. The VaR says that recent adverse monthly copper-price increases crossed an 8.2% threshold. The CVaR says that the average increase in the recent tail was about 9.0%. For a buyer, both numbers describe upside price pressure; for a producer, the relevant tail would normally be price decreases.

Table 10.4: Rolling one-month copper price-change risk measures

Measure	Price change (%)	Risk-management interpretation
Empirical $\mathrm{VaR}^R_{.95}$	8.2	Adverse monthly price-increase threshold
Empirical $\mathrm{CVaR}^R_{.95}$	9.0	Average recent tail increase beyond VaR

Figure 10.5 shows the same measures through time. The figure should be read as an early-warning chart. A rising VaR line means that the recent threshold for an adverse monthly price increase is moving upward. A rising CVaR line means that the recent tail is not only crossing a higher threshold but also becoming more severe on average. A decline does not mean there is no risk; it means that the most adverse observations in the rolling window have become less severe or have dropped out of the window.

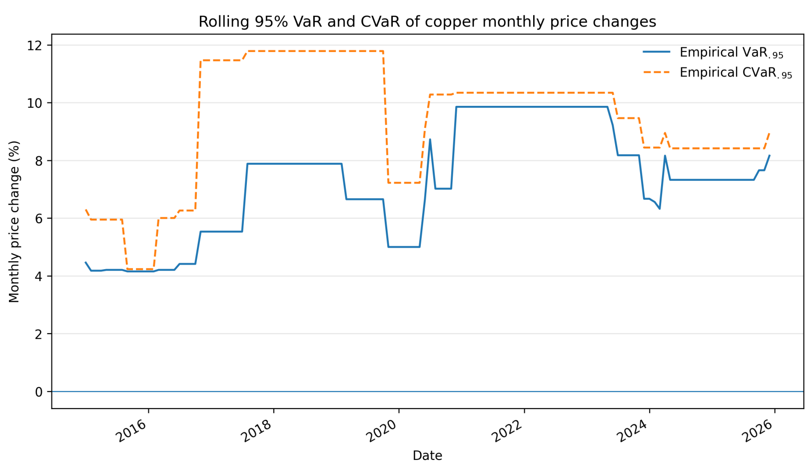

Figure 10.5: Empirical rolling one-month copper VaR and CVaR, expressed as monthly price-change percentages.

Table 10.5 extends the same idea to several commodities. The VaR and CVaR columns are monthly price-change percentages over 2015–2025. The

ACF(1) column is the first-order autocorrelation of monthly price changes, defined as the correlation between R_t and R_{t-1}. It indicates whether monthly changes tend to persist from one month to the next.

Table 10.5: Selected commodity risk indicators from World Bank Pink Sheet data, 2015–2025

Commodity	Price US$/unit	12m chg. %	Vol. %	VaR$_{.95}$ %	CVaR$_{.95}$ %	ACF(1)
Copper	11,785.25	32.2	4.6	8.6	10.4	0.25
Aluminum	2,875.53	13.2	4.4	7.1	9.3	0.32
Crude oil, avg.	60.88	-15.8	10.1	15.2	23.2	0.27
Natural gas, Europe	9.48	-31.6	17.1	36.2	46.8	0.33
Wheat, US HRW	242.80	-3.7	5.8	9.1	12.7	0.10

The table gives a concise managerial comparison. The 12-month change and the VaR/CVaR columns answer different questions. The 12-month change reports what happened over the latest year in the sample. VaR and CVaR report the upper tail of monthly price increases over 2015–2025. For this reason, European natural gas can show a negative latest 12-month change and still have a very high VaR and CVaR: prices may have fallen over the latest year, while the historical monthly series still contains large upward jumps. A common dashboard can compare commodities, but action thresholds should differ by commodity. A 10% monthly move may be extreme for one commodity and less exceptional for another.

Rolling empirical VaR and CVaR are transparent, but they have limitations. They depend on the window length. They are backward-looking. They use very few tail observations when the window is short. They may miss new structural breaks. For these reasons, rolling VaR and CVaR should be used as monitoring indicators, not as complete forecasts. The next section uses simulation and stress paths to move from one-period indicators to multi-period risk scenarios.

10.8 Simulation and Scenario Paths

Rolling VaR and CVaR summarize one-period downside exposure. Many managerial decisions, however, depend on paths. A procurement director does not only care about the next month. A sequence of adverse monthly price changes can damage margins, exhaust budgets, or force contract renegotiation. Simulation converts the rolling indicators from the previous section into a distribution of possible future price paths.

The simulation uses the same building blocks developed in Sections 10.3–10.7: monthly price changes, rolling mean, rolling volatility, autocorrelation, and the recursive price equation. The return equation for simulated path m is:

$$R_{t+h} = \mu_t + b\,(R_{t+h-1} - \mu_t) + e_{t+h}, \qquad h = 1, \ldots, H. \tag{10.13}$$

Here H is the horizon, μ_t is the recent rolling mean of monthly price changes, b is the estimated AR(1) persistence coefficient, and e_{t+h} is a shock drawn from recent empirical residuals. The coefficient b carries forward the persistence analysis from Section 10.6. If $b = 0$, the next simulated change depends only on the recent mean and the new shock. If $b > 0$, a high current price change tends to raise the next simulated price change. If $b < 0$, the simulation incorporates partial reversal. The simulation repeats this recursion many times, but the equation itself describes one simulated path at a time.

The empirical residuals are computed from the same rolling window used for the current risk indicators:

$$e_j = R_j - \left[\mu_t + b\,(R_{j-1} - \mu_t)\right]. \tag{10.14}$$

This empirical route resamples recent residuals with replacement, a residual-bootstrap idea used to generate simulated series while preserving an estimated dependence structure [6, 17]. Monte Carlo simulation then converts uncertain inputs into a distribution of possible paths rather than a single point estimate [8]. Similar time-series simulation logic is used in procurement and inventory settings when future requirements are volatile [4].

After each simulated monthly change is generated, the price is updated using the same recursion introduced in Equation (10.2). In words, a simulated path is built month by month: draw a price change, apply it to the previous price, and repeat. Monthly changes are the stochastic risk process. Prices are the accumulated economic path obtained by applying those changes recursively.

Table 10.6: Pseudo-code for empirical residual simulation of copper price paths

Block	Procedure
Inputs	Choose the horizon $H = 6$, number of simulated paths $M =$ 10,000, rolling window $w = 36$, latest price P_t, and latest monthly price change R_t.
Estimate state	Compute the recent mean μ_t and estimate the persistence coefficient b from the AR(1) regression $R_j = a + bR_{j-1} + \varepsilon_j$ over the same window.
Build shocks	Compute residuals $e_j = R_j - [\mu_t + b(R_{j-1} - \mu_t)]$ and store them as the empirical shock pool.
Repeat paths	For each simulated path, start from the latest observed price and latest observed monthly change.
Repeat months	For each future month, draw one residual from the shock pool, compute the next monthly change using the AR(1) update, and update the price with $P_{t+h} = P_{t+h-1}(1 + R_{t+h})$.
Output	After all paths are generated, report the median simulated price and the 5th–95th percentile range at each horizon.

Table 10.6 is written as pseudocode rather than software code. It defines inputs, estimates the state variables before the simulation loop, separates the loop over paths from the loop over future months, and states the outputs. This makes the simulation reproducible and easier to audit.

Figure 10.6 shows the resulting six-month copper price paths. The simulation uses 10,000 paths. The central line reports the median simulated path, and the shaded region reports the 5th–95th percentile range. The figure should not be read as a point forecast. It is a decision chart. The median path shows the central tendency under the current rolling state. The lower part of the fan shows possible price relief for the buyer. The upper part shows adverse procurement risk.

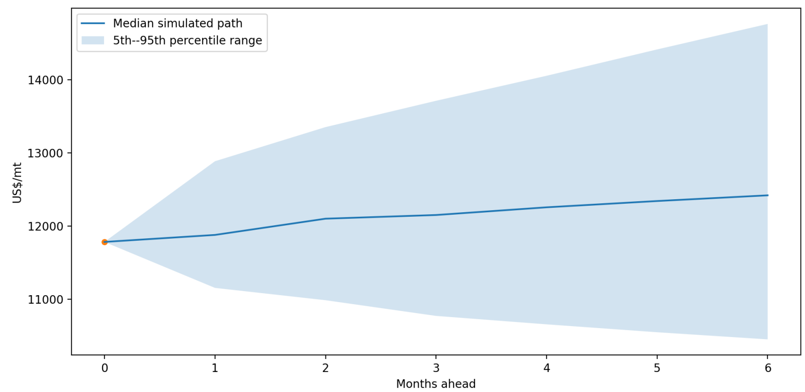

Figure 10.6: Simulated six-month copper price paths using rolling mean, AR(1) persistence, and empirical residual sampling.

The simulation also clarifies the difference between probabilistic paths and deterministic stress tests. The fan chart samples many possible residual sequences. The stress paths in Table 10.7 deliberately impose persistent adverse monthly changes. Starting from the latest copper price of US$11,785.25 per metric ton, the baseline applies the recent 36-month mean monthly change. The adverse path applies the rolling empirical 95th percentile monthly increase. The severe path applies the rolling empirical CVaR monthly increase. These scenarios are not probability forecasts; they are management tests.

Table 10.7: Illustrative six-month copper price scenarios

Scenario	Monthly change (%)	Six-month change (%)
Baseline	1.0	6.4
Adverse VaR path	8.2	60.1
Severe CVaR path	9.0	67.4

Table 10.7 is intentionally severe because it repeats adverse monthly changes for six consecutive months. The baseline is a budget reference under recent average pressure. The VaR and CVaR paths are stress tests: they ask whether the firm could absorb repeated adverse-threshold or average-tail months through hedging, inventory, price-adjustment clauses, supplier renegotiation, or escalation.

The empirical residual approach is transparent because simulated shocks come from recent observed residuals. Its limitation is sample size: with a 36-month window, the residual pool is small. Simulation should therefore be interpreted together with rolling VaR, rolling CVaR, deterministic stress paths, and managerial judgement.

10.9 Summary and Managerial Takeaways

This chapter introduced time-series foundations for risk analytics through a practical commodity procurement case. The central message is that risk changes through time. A risk variable should therefore be represented not only as a one-period random variable, but as a sequence of observations that can be monitored, smoothed, stress-tested, and translated into decisions.

The copper example showed why the distinction between price levels and price changes matters. Price levels determine the economic exposure and are essential for managerial interpretation. Monthly price changes are more useful for short-term risk measurement because they provide a more stable basis for rolling volatility, autocorrelation, VaR, CVaR, and scenario inputs. Non-stationarity explains why the current price environment may not be well represented by an old historical average.

The methods developed in the chapter are deliberately simple. Moving averages summarize recent direction. Rolling standard deviations summarize current volatility. Autocorrelation indicates whether shocks tend to persist or reverse. Rolling VaR and CVaR summarize the recent adverse tail of monthly price changes. Simulation then converts those one-period indicators into multi-period scenario paths. Together, these tools turn a price history into a dynamic risk-monitoring framework.

The managerial lesson is direct: time-series analysis is useful when it changes the quality of risk decisions. It helps managers decide whether to monitor, hedge, build inventory, renegotiate contracts, revise budgets, or escalate exposure. The purpose is not only better point forecasting. It is better communication about how risk is evolving and what future paths the organization should be prepared to manage.

Managerial takeaways.

- Use price levels to understand exposure, but use price changes for short-horizon risk measurement.
- Treat non-stationarity as a warning that old averages may no longer be valid baselines.
- Use moving averages, rolling volatility, and autocorrelation as monitoring indicators, not mechanical rules.

- Interpret rolling VaR as an adverse threshold and rolling CVaR as the average severity beyond that threshold.
- Use simulation and stress paths to translate one-month indicators into multi-month planning questions.
- Judge the analysis by whether it improves monitoring, hedging, inventory, budgeting, escalation, and risk communication.

Bibliography

[1] George E. P. Box, Gwilym M. Jenkins, Gregory C. Reinsel, and Greta M. Ljung. *Time Series Analysis: Forecasting and Control*. John Wiley & Sons, Hoboken, NJ, 5 edition, 2016.

[2] Peter J. Brockwell and Richard A. Davis. *Introduction to Time Series and Forecasting*. Springer Texts in Statistics. Springer, Cham, 3 edition, 2016.

[3] Michael P. Clements and David F. Hendry. *Forecasting Non-Stationary Economic Time Series*. MIT Press, Cambridge, MA, 1999.

[4] Robert Davies, Tim Coole, and David Osipyw. The application of time series modelling and monte carlo simulation: Forecasting volatile inventory requirements. *Applied Mathematics*, 5:1152–1168, 2014.

[5] David A. Dickey and Wayne A. Fuller. Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74(366):427–431, 1979.

[6] Bradley Efron and Robert J. Tibshirani. *An Introduction to the Bootstrap*. Chapman & Hall/CRC, New York, 1994.

[7] Eugene F. Fama and Kenneth R. French. Permanent and temporary components of stock prices. *Journal of Political Economy*, 96(2):246–273, 1988.

[8] Paul Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer, New York, 2004.

[9] Clive W. J. Granger and Paul Newbold. Spurious regressions in econometrics. *Journal of Econometrics*, 2(2):111–120, 1974.

[10] James D. Hamilton. *Time Series Analysis*. Princeton University Press, Princeton, NJ, 1994.

[11] Rob J. Hyndman and George Athanasopoulos. *Forecasting: Principles and Practice*. OTexts, Melbourne, 3 edition, 2021.

[12] Alexander J. McNeil, Rüdiger Frey, and Paul Embrechts. *Quantitative Risk Management: Concepts, Techniques, and Tools*. Princeton University Press, Princeton, NJ, 2005.

[13] Douglas C. Montgomery, Cheryl L. Jennings, and Murat Kulahci. *Introduction to Time Series Analysis and Forecasting*. Wiley Series in Probability and Statistics. John Wiley & Sons, Hoboken, NJ, 2 edition, 2015.

[14] James M. Poterba and Lawrence H. Summers. Mean reversion in stock prices: Evidence and implications. *Journal of Financial Economics*, 22(1):27–59, 1988.

[15] R. Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–41, 2000.

[16] R. Tyrrell Rockafellar and Stanislav Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7):1443–1471, 2002.

[17] Robert H. Shumway and David S. Stoffer. *Time Series Analysis and Its Applications: With R Examples*. Springer Texts in Statistics. Springer, Cham, 4 edition, 2017.

[18] World Bank. Commodity Price Data (The Pink Sheet): Monthly Prices. Excel workbook, CMO-Historical-Data-Monthly.xlsx, 2025. Monthly commodity prices in nominal U.S. dollars. Dataset used in Chapter 10 case study.

