Risk Analytics

Machine Learning and Optimization
for Data-Driven Decision Making

Fernando S. Oliveira

Draft version — April 6, 2026

Chapter 4

VaR, CVaR and Downside Risk Assessment

Chapter 4

VaR, CVaR and Downside Risk Assessment

4.1 Introduction: From Utility Theory to Downside-Risk Measures

Chapter 3 developed the normative foundations of risky choice. It explained how risk-neutral, risk-averse, and risk-seeking behavior can be represented through the curvature of the utility function, and how certainty equivalents and risk premia translate those attitudes into monetary terms. That framework remains fundamental because it explains why two decision makers may rank the same uncertain prospect differently. In practice, however, risk analytics often requires a more operational language. Once uncertainty is represented through data, forecasts, or probability models, the analyst must still decide how to summarize the adverse part of the distribution and how to use that summary in a policy decision.

This chapter develops one such bridge. Its purpose is to introduce Value at Risk (VaR) and Conditional Value at Risk (CVaR) as practical summaries of downside risk, explain their interpretation and properties, and show how they can be embedded into an optimization problem. The chapter has two layers. The first layer is conceptual: it introduces VaR and CVaR for a generic loss variable, derives their normal-distribution formulas, and clarifies how they differ from the utility-based notion of risk aversion developed in Chapter 3. The second layer is applied: it uses an original newsvendor example based on seasonal battery backup kits to show how these risk measures affect actual policy choice.

The distinction between *policy evaluation* and *policy optimization* is fundamental throughout the chapter. Policy evaluation asks what expected loss, VaR, or CVaR are associated with a given decision. Policy optimization asks which decision minimizes a chosen objective. This distinction matters because the expected-loss-optimal decision and the CVaR-sensitive decision need not coincide. A VaR-based optimization problem could also be defined, but this chapter focuses on CVaR because it is more informative about tail severity and better behaved in optimization.

4.2 VaR and CVaR as Downside-Risk Measures

4.2.1 Loss functions and tail-oriented risk measurement

To keep the chapter conceptually clean, the analysis is organized around a loss variable L. This is important because VaR and CVaR are defined on adverse outcomes. If a model is naturally written in terms of profit Π, the corresponding loss can always be written as

$$L = -\Pi,$$

or, more generally, as another economically meaningful downside-oriented transformation. In this chapter, we work directly with loss functions.

The central issue is simple: once a loss distribution has been induced by a decision, how should its adverse tail be summarized? Mean and variance often do not suffice. Two decisions may have the same expected loss and similar dispersion while differing sharply in how severe the worst outcomes are. VaR and CVaR are designed precisely to describe that part of the distribution.

4.2.2 Value at Risk

For a loss variable L, the Value at Risk at confidence level $\alpha \in (0,1)$ is defined as the smallest threshold ℓ such that

$$P(L \leq \ell) \geq \alpha.$$

Equivalently,

$$\mathrm{VaR}_{\alpha}(L) = \inf\{\ell \in \mathbb{R} : P(L \leq \ell) \geq \alpha\}. \tag{4.1}$$

This chapter uses the standard *confidence-level* convention: α denotes the confidence level, while $1-\alpha$ is the corresponding tail probability. Thus, $\alpha = 0.95$ refers to the worst 5% tail. An equivalent statement is

$$P(L > \mathrm{VaR}_{\alpha}(L)) \leq 1 - \alpha.$$

VaR should therefore be interpreted as a threshold. If $\mathrm{VaR}_{0.95}(L) = x$, then only 5% of losses exceed x.

4.2.3 Conditional Value at Risk

For continuous loss distributions, Conditional Value at Risk, also called Expected Shortfall, can be written as the average loss in the adverse tail beyond VaR:

$$\mathrm{CVaR}_\alpha(L) = E[L \mid L \geq \mathrm{VaR}_\alpha(L)]. \tag{4.2}$$

The difference between VaR and CVaR is fundamental. VaR identifies where the tail begins. CVaR measures how severe losses are on average once the decision maker is already in that tail. Thus, VaR is a threshold, while CVaR is a tail average.

4.2.4 Normal benchmark: VaR and CVaR for a normal loss variable

To see the mechanics clearly, suppose a generic loss variable Y satisfies

$$Y \sim N(\mu, \sigma^2).$$

Then

$$P(Y \leq \ell) = \Phi\left(\frac{\ell - \mu}{\sigma}\right),$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. Setting this equal to α gives

$$\Phi\left(\frac{\ell - \mu}{\sigma}\right) = \alpha.$$

Applying the inverse normal distribution function yields

$$\frac{\ell - \mu}{\sigma} = z_\alpha, \qquad z_\alpha = \Phi^{-1}(\alpha).$$

Hence

$$\mathrm{VaR}_\alpha(Y) = \mu + \sigma z_\alpha. \tag{4.3}$$

For CVaR, standardize Y as

$$Z = \frac{Y - \mu}{\sigma}, \qquad Z \sim N(0, 1),$$

so that

$$Y = \mu + \sigma Z.$$

Then

$$\mathrm{CVaR}_\alpha(Y) = E[Y \mid Y \geq \mathrm{VaR}_\alpha(Y)] = \mu + \sigma E[Z \mid Z \geq z_\alpha].$$

For the standard normal distribution,

$$E[Z \mid Z \geq z] = \frac{\phi(z)}{1 - \Phi(z)},$$

where $\phi(\cdot)$ is the standard normal density. Substituting $z = z_\alpha$ and using $\Phi(z_\alpha) = \alpha$ gives

$$\mathrm{CVaR}_\alpha(Y) = \mu + \sigma \frac{\phi(z_\alpha)}{1 - \alpha}. \tag{4.4}$$

This formula is central. The cutoff z_α identifies the relevant tail threshold in the standard normal distribution. The ratio

$$\frac{\phi(z_\alpha)}{1 - \alpha}$$

is a standardized tail-severity multiplier. Multiplying by σ converts that multiplier into the units of the loss variable, and adding μ shifts it to the correct location.

Worked Example: Normal CVaR

Using

$$Y \sim N(100, 20^2), \qquad \alpha = 0.95,$$

we have

$$z_{0.95} \approx 1.6449, \qquad \phi(z_{0.95}) \approx 0.1031.$$

So

$$\mathrm{VaR}_{0.95}(Y) = 100 + 20(1.6449) \approx 132.90,$$

and

$$\mathrm{CVaR}_{0.95}(Y) = 100 + 20\frac{0.1031}{0.05} \approx 141.24.$$

Thus, VaR gives the threshold of the worst 5% losses, while CVaR gives the average loss within that same worst 5% tail.

Figure 4.1 should be read as a direct visualization of the worked example. The dashed vertical line marks the quantile threshold that separates ordinary losses from the worst 5% of outcomes. The dotted line lies further into the tail because CVaR is not a threshold but an average over that tail region. The figure therefore makes visible the most important conceptual distinction of the chapter: VaR identifies where the tail begins, while CVaR measures how severe the tail is on average.

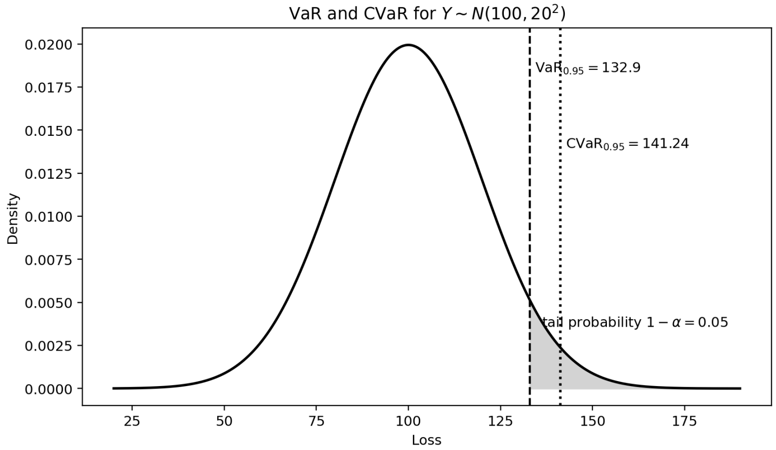

Figure 4.1: VaR and CVaR for the normal loss example $Y \sim N(100, 20^2)$ at confidence level $\alpha = 0.95$. The dashed line marks $\mathrm{VaR}_{0.95}(Y) \approx 132.90$, while the dotted line marks $\mathrm{CVaR}_{0.95}(Y) \approx 141.24$. The shaded region is the worst 5% tail of the loss distribution.

4.2.5 *Properties of VaR and CVaR*

VaR is intuitive but incomplete. VaR is easy to communicate and aligns naturally with threshold-based managerial thinking. But it says nothing about what happens beyond the threshold. Two decisions may have the same VaR and yet very different extreme-tail behavior.

CVaR is tail-sensitive. CVaR incorporates the magnitude of losses inside the tail. For this reason, it is generally more informative than VaR when rare severe outcomes matter.

CVaR is coherent; VaR may not be. CVaR is a coherent risk measure: it satisfies monotonicity, translation invariance, positive homogeneity, and subadditivity. In particular, it behaves well in optimization because it is also convex. VaR may fail subadditivity in some settings, which is one reason CVaR is often preferred in optimization and risk management [2, 9, 10, 6].

VaR and CVaR are not utility functions. This is the key conceptual bridge to Chapter 3. In Chapter 3, risk aversion was represented by the curvature of a utility function:

$$\max_x E[u(\Pi(x))].$$

VaR and CVaR are different objects. They are not utility functions and do not represent risk aversion through concavity. Instead, they are distributional summaries of downside risk. They become operational representations of downside sensitivity only when embedded in a decision rule such as

$$\min_x \left\{ E[L(x)] + \lambda \, \mathrm{CVaR}_\alpha(L(x)) \right\}.$$

4.3 Policy Evaluation versus Optimal Policy Choice

It is important to distinguish between two different tasks.

Policy evaluation treats the decision as fixed. The analyst asks: what are the expected loss, VaR, and CVaR associated with this policy?

Optimal policy choice treats the decision variable itself as endogenous. The analyst asks: which policy minimizes expected loss, VaR, CVaR, or a mean–CVaR objective?

This distinction is essential. A policy may look attractive under expected loss but unattractive under CVaR, and vice versa.

Abstractly, if a policy variable is denoted by x, a downside-sensitive decision maker may solve

$$x^* \in \arg\min_{x \in \mathcal{X}} \mathrm{CVaR}_\alpha(L(x)), \tag{4.5}$$

or, more generally,

$$x^* \in \arg\min_{x \in \mathcal{X}} \left\{ E[L(x)] + \lambda \, \mathrm{CVaR}_\alpha(L(x)) \right\}. \tag{4.6}$$

Here $\lambda \geq 0$ controls how much weight is placed on tail losses. If $\lambda = 0$, the decision rule reduces to expected-loss minimization only. If $\lambda > 0$, adverse-tail outcomes matter explicitly. Larger λ means greater downside sensitivity.

4.4 The Newsvendor Problem as an Application

The newsvendor model is one of the most important benchmark models in operations management, inventory control, and supply chain theory. Its value is not only pedagogical. It captures a very common real-world problem: a decision must be made *before* uncertainty is resolved, and ordering too much and ordering too little are both costly, but in different ways [1, 3, 4, 8].

This model is especially suitable for the present chapter because it gives a simple and economically meaningful setting in which:

- uncertainty is explicit;
- the loss function is asymmetric;
- the risk-neutral and downside-sensitive policies can differ.

Consider a retailer selling portable battery backup kits before the summer blackout season. The retailer has a single opportunity to order inventory before demand is realized. If the season is mild, some units remain unsold and must be cleared at a discount. If the season is severe and demand exceeds inventory, profitable sales are lost.

Let Q denote the order quantity and D uncertain demand. The generic newsvendor parameters are:

- selling price p,
- procurement cost c,
- salvage value v,

with

$$p > c > v.$$

For downside-risk analysis, the relevant loss function is the standard mismatch loss:

$$L(Q, D) = C_o(Q - D)^+ + C_u(D - Q)^+, \tag{4.7}$$

where

$$C_o = c - v, \qquad C_u = p - c.$$

This formulation is preferable here to simply setting $L = -\Pi$, because it separates the two economically meaningful sources of downside risk:

- overage loss when inventory remains unsold;
- underage loss when demand exceeds supply.

Running Example: Seasonal Battery Backup Kits

The retailer must order battery backup kits before the season begins. Demand is modeled as

$$D \sim N(1200, 420^2).$$

The economic parameters are

$$p = 980, \qquad c = 620, \qquad v = 390.$$

So

$$C_u = p - c = 360, \qquad C_o = c - v = 230.$$

The loss function is therefore

$$L(Q, D) = 230(Q - D)^+ + 360(D - Q)^+.$$

Thus:

- if $D < Q$, the retailer carries excess units and incurs overage loss;
- if $D > Q$, the retailer loses profitable sales and incurs underage loss.

4.5 The Critical Ratio and the Risk-Neutral Benchmark

The expected-loss benchmark is equivalent to the expected-profit benchmark of the classical newsvendor model. In this setting, minimizing expected mismatch loss is equivalent to maximizing expected profit because the two objective functions differ only by terms that do not affect the optimizer. Let Q^{EP} denote this classical expected-profit-maximizing, equivalently expected-loss-minimizing, order quantity.

The key idea is marginal. Suppose the retailer increases the order quantity by one unit. That marginal unit is beneficial only if demand turns out to exceed the original order quantity, in which case the unit avoids an underage loss C_u. But the same marginal unit is harmful if demand turns out to be below the order quantity, in which case it generates an overage loss C_o.

So the optimal policy is obtained by balancing the expected marginal gain from ordering more against the expected marginal cost of ordering more. This condition leads to the critical-ratio rule

$$F_D(Q^{EP}) = \frac{C_u}{C_u + C_o}. \tag{4.8}$$

This expression can be interpreted very directly. The left-hand side, $F_D(Q^{EP})$, is the probability that demand is at or below the order quantity.

The right-hand side is the ratio of the underage cost to the total mismatch cost. Thus the optimal order quantity is the quantile of demand that balances the two types of marginal mismatch loss.

For the battery backup kit problem,

$$\frac{C_u}{C_u + C_o} = \frac{360}{360 + 230} = 0.6102.$$

Since

$$D \sim N(1200, 420^2),$$

we obtain

$$Q^{EP} = \mu + z^*\sigma, \qquad \Phi(z^*) = 0.6102.$$

The corresponding quantile is approximately

$$z^* \approx 0.28,$$

so

$$Q^{EP} = 1200 + 0.28(420) \approx 1318. \tag{4.9}$$

This is the risk-neutral benchmark in the classical newsvendor sense because it minimizes expected loss only. It does not distinguish between moderate and severe losses once the loss distribution has been induced.

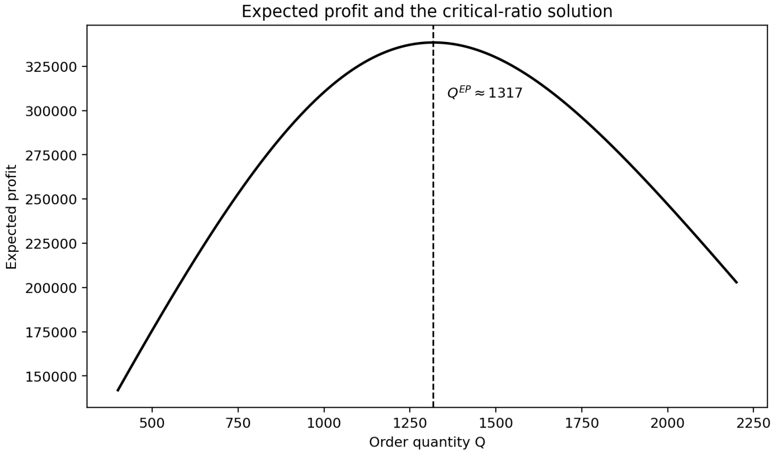

Figure 4.2: Expected profit as a function of the order quantity Q for the battery backup kit problem. The dashed line marks the expected-profit-maximizing quantity $Q^{EP} \approx 1318$.

Figure 4.2 shows the risk-neutral benchmark graphically. The curve is concave because very small order quantities sacrifice too many profitable sales, while very large order quantities generate too much leftover inventory. The critical-ratio quantity appears at the peak of this curve because it is exactly the quantity that balances the two mismatch costs in expected-value terms.

Equation (4.8) can also be read as a direct policy rule: order up to the demand quantile whose cumulative probability equals the relative importance of shortage cost. The larger C_u is relative to C_o, the further into the right tail the retailer orders. This is why the critical ratio is not just a formula but an economic summary of the trade-off embedded in the model.

4.6 Evaluating the Newsvendor Policy through VaR and CVaR

The critical-ratio solution $Q^{EP} \approx 1318$ is optimal under expected profit, but it induces a specific distribution of losses through (4.7). VaR and CVaR provide two ways of evaluating the tail properties of that policy.

Because the loss function is piecewise linear in demand, the induced loss distribution is not normal, even though demand itself is normal. This is why the newsvendor is a useful application: a simple normal demand model can still generate a nontrivial loss distribution.

4.6.1 From the loss function to the induced loss distribution

For a given order quantity Q and threshold $\ell \geq 0$, the event

$$L(Q, D) \leq \ell$$

means that the realized mismatch loss does not exceed ℓ. Because the loss function is piecewise linear, this condition is easiest to analyze in two cases.

If demand is below the order quantity, $D \leq Q$, then the loss is purely an overage loss:

$$L(Q, D) = C_o(Q - D).$$

So the requirement $L(Q, D) \leq \ell$ implies

$$C_o(Q - D) \leq \ell, \qquad D \geq Q - \frac{\ell}{C_o}.$$

If demand is above the order quantity, $D \geq Q$, then the loss is purely an underage loss:

$$L(Q, D) = C_u(D - Q).$$

So the requirement $L(Q, D) \leq \ell$ implies

$$C_u(D - Q) \leq \ell, \qquad D \leq Q + \frac{\ell}{C_u}.$$

Combining the two cases gives

$$L(Q, D) \leq \ell \quad \Longleftrightarrow \quad Q - \frac{\ell}{C_o} \leq D \leq Q + \frac{\ell}{C_u}. \tag{4.10}$$

This is the key step. It says that the event "loss does not exceed ℓ" is exactly the same as the event that demand lies inside a particular interval around Q.

Since $D \sim N(\mu, \sigma^2)$, the cumulative distribution function of the induced loss is therefore

$$F_{L_Q}(\ell) = \Phi\left(\frac{Q + \ell/C_u - \mu}{\sigma}\right) - \Phi\left(\frac{Q - \ell/C_o - \mu}{\sigma}\right), \qquad \ell \geq 0. \tag{4.11}$$

This equation gives the full CDF of the induced loss for any order quantity Q.

4.6.2 Computing the Value at Risk of the induced loss

For a fixed order quantity Q, let $v_\alpha(Q)$ denote the α-level Value at Risk of the induced loss. By definition, it must satisfy

$$F_{L_Q}(v_\alpha(Q)) = \alpha.$$

Substituting (4.11) gives

$$\Phi\left(\frac{Q + v_\alpha(Q)/C_u - \mu}{\sigma}\right) - \Phi\left(\frac{Q - v_\alpha(Q)/C_o - \mu}{\sigma}\right) = \alpha. \tag{4.12}$$

This is the VaR equation for the newsvendor loss. It usually has no simple closed-form solution, but it is a one-dimensional equation in the unknown $v_\alpha(Q)$, so it can be solved numerically very easily.

4.6.3 Computing the Conditional Value at Risk of the induced loss

Once the VaR threshold $v_\alpha(Q)$ is known, the CVaR is the average loss conditional on being in the adverse tail:

$$\mathrm{CVaR}_\alpha(L_Q) = E[L(Q,D) \mid L(Q,D) \geq v_\alpha(Q)].$$

For the newsvendor loss, that tail event comes from two disjoint demand regions.

On the low-demand side, severe losses come from excess inventory:

$$D \leq a(Q) = Q - \frac{v_\alpha(Q)}{C_o}.$$

On the high-demand side, severe losses come from lost sales:

$$D \geq b(Q) = Q + \frac{v_\alpha(Q)}{C_u}.$$

So the adverse tail of the induced loss does not come from one single interval. It comes from:

- a lower-demand overage tail;
- an upper-demand underage tail.

Define the standardized cutoffs

$$z_a = \frac{a(Q) - \mu}{\sigma}, \qquad z_b = \frac{b(Q) - \mu}{\sigma}. \tag{4.13}$$

These standardized quantities are introduced because demand is normal. They allow the tail contributions to be written using the standard normal density $\phi(\cdot)$ and distribution function $\Phi(\cdot)$.

Because demand is continuous, the induced loss distribution is also continuous, so the tail event $L_Q \geq \mathrm{VaR}_\alpha(L_Q)$ has probability exactly $1 - \alpha$. Using the standard formulas for truncated normal expectations, the CVaR of the newsvendor loss is

$$\mathrm{CVaR}_\alpha(L_Q) = \frac{C_o\left[(Q-\mu)\Phi(z_a) + \sigma\phi(z_a)\right] + C_u\left[\sigma\phi(z_b) + (\mu - Q)(1 - \Phi(z_b))\right]}{1-\alpha}. \tag{4.14}$$

This expression can be read term by term.

The first bracket,

$$C_o \left[(Q-\mu)\Phi(z_a) + \sigma\phi(z_a) \right],$$

is the expected overage loss contributed by the low-demand tail.

The second bracket,

$$C_u \left[\sigma\phi(z_b) + (\mu - Q)(1-\Phi(z_b)) \right],$$

is the expected underage loss contributed by the high-demand tail.

The denominator $1-\alpha$ is the total probability mass of the tail. So the expression is the average loss conditional on being in the tail.

Equation (4.14) is explicit once the VaR threshold $v_\alpha(Q)$ is known. However, it is not a fully closed-form expression in Q alone, because $v_\alpha(Q)$ is determined implicitly by equation (4.12). The relationship between the two equations is direct: equation (4.12) determines the tail boundary, while equation (4.14) computes the average loss over the tail defined by that boundary.

Using the standardized cutoffs z_a and z_b, equation (4.12) can be written as

$$\Phi(z_b) - \Phi(z_a) = \alpha,$$

so that the total tail probability is

$$\Phi(z_a) + 1 - \Phi(z_b) = 1 - \alpha.$$

Equation (4.14) then evaluates the conditional expectation of the loss over exactly those two tail regions.

Figure 4.3 is a policy-evaluation figure. For each possible order quantity, it asks:

- what is the threshold of the worst 5% losses?
- what is the average loss within that worst 5% tail?

The U-shaped curves reflect the basic newsvendor trade-off. If Q is too low, shortage losses dominate. If Q is too high, overage losses dominate. The gap between the VaR and CVaR curves measures how much worse the average tail outcome is than the threshold that defines the tail.

Figure 4.4 complements the previous one by showing how the realized loss varies with demand for selected policies. Each curve is V-shaped. On the left side of the vertex, overage losses dominate; on the right side, underage losses dominate. This figure is useful because it shows structurally how different order quantities induce different loss distributions.

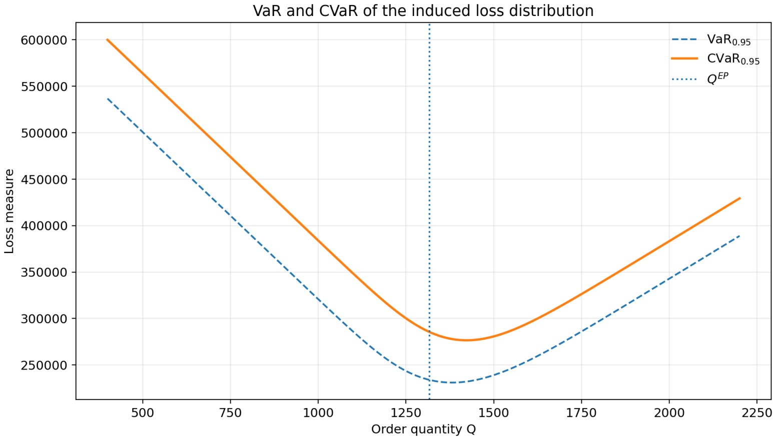

Figure 4.3: VaR and CVaR of the induced loss distribution as functions of the order quantity Q for the battery backup kit problem, using $\alpha = 0.95$. The dotted vertical line marks the expected-profit benchmark $Q^{EP} \approx 1318$.

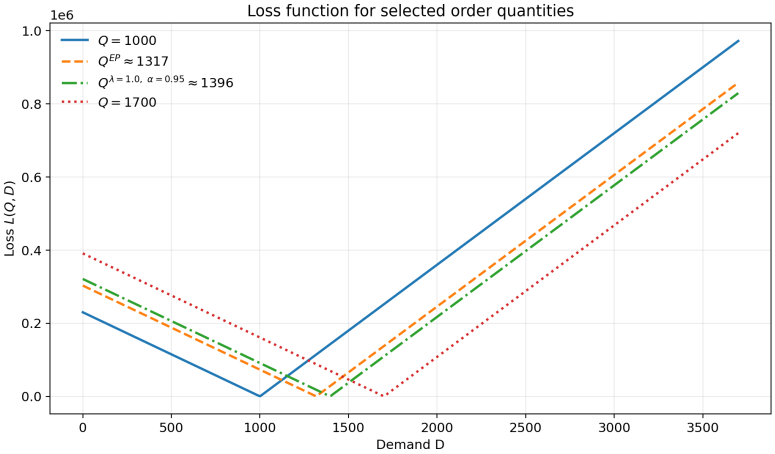

Figure 4.4: Loss function $L(Q, D)$ plotted against realized demand D for selected order quantities in the battery backup kit problem.

4.7 Optimal Policy Choice under Mean–CVaR

The previous section evaluated VaR and CVaR for a given order quantity. We now move to policy optimization. The mean–CVaR formulation studied here belongs to a broader family of risk-aware optimization models. In related work on robust and distributionally robust decision-making, the analyst does not rely on a single fully specified probability distribution, but instead optimizes against a set of plausible distributions or uncertainty realizations. This broader perspective has been developed extensively in the optimization literature and is especially relevant for inventory and supply-chain problems in which tail behavior is difficult to estimate precisely [11, 5, 7].

A downside-sensitive retailer may solve

$$Q^{CVaR} \in \arg\min_{Q} \mathrm{CVaR}_\alpha(L_Q), \tag{4.15}$$

or, more generally,

$$Q^{\lambda,\alpha} \in \arg\min_{Q} \left\{ E[L(Q,D)] + \lambda \, \mathrm{CVaR}_\alpha(L_Q) \right\}. \tag{4.16}$$

This objective combines average loss and tail severity. The parameter $\lambda \geq 0$ controls how much weight is placed on the CVaR term. When $\lambda = 0$, the problem reduces to minimizing expected loss only. That is the risk-neutral benchmark in the classical newsvendor sense. When $\lambda > 0$, the decision maker becomes increasingly sensitive to severe tail outcomes.

Let

$$z = \frac{Q-\mu}{\sigma}.$$

Then the expected mismatch loss has the closed form

$$E[L(Q,D)] = C_o\big[(Q-\mu)\Phi(z) + \sigma\phi(z)\big] + C_u\big[\sigma\phi(z) + (\mu - Q)(1-\Phi(z))\big]. \tag{4.17}$$

Differentiating (4.17) gives

$$(C_o + C_u)\Phi\left(\frac{Q-\mu}{\sigma}\right) - C_u = 0,$$

which reproduces the critical-ratio rule. So the classical newsvendor solution is the optimizer of expected loss, but not necessarily of CVaR or mean–CVaR.

Equation (4.16) should therefore be interpreted as a generalization of the risk-neutral problem. The first term $E[L(Q,D)]$ measures average mismatch cost. The second term $\lambda \, \mathrm{CVaR}_\alpha(L_Q)$ penalizes the severity of the adverse tail. The role of λ is to determine how much the decision maker is willing to trade expected performance for tail protection.

There is generally no comparably simple closed-form "CVaR critical ratio". The reason is that the CVaR term depends on the threshold $v_\alpha(Q)$, and that threshold itself depends implicitly on Q through equation (4.12). So the optimization problem is explicit but not closed form. In practice, it is solved numerically.

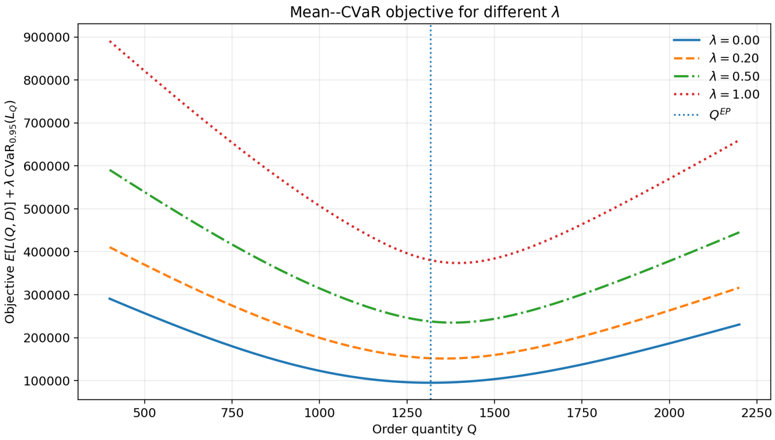

Figure 4.5: Mean–CVaR objective $E[L(Q,D)]+\lambda\,\mathrm{CVaR}_{0.95}(L_Q)$ as a function of the order quantity Q for different values of λ.

Figure 4.5 is the key optimization figure. When $\lambda = 0$, the curve is simply expected loss, and the minimum occurs at the critical-ratio solution Q^{EP}. As λ increases, the objective tilts because more weight is placed on adverse-tail outcomes. The minimizing quantity shifts. This shift is the formal expression of downside-sensitive choice.

4.8 Sensitivity Analysis in λlambda and αalpha

The parameters λ and α are not merely computational settings. They are operational levers controlling downside sensitivity.

In the objective

$$E[L(Q,D)] + \lambda\,\mathrm{CVaR}_\alpha(L_Q),$$

λ controls how much weight is placed on CVaR relative to expected loss. If $\lambda = 0$, the decision rule reduces to the expected-loss benchmark. Larger values of λ place more emphasis on tail severity.

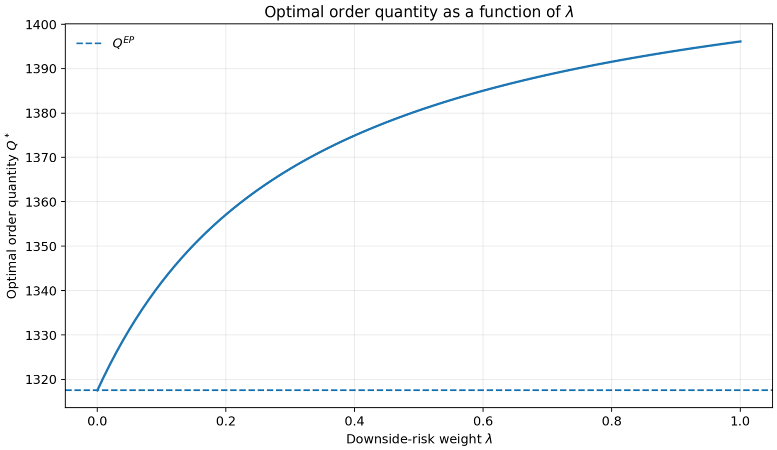

Figure 4.6: Optimal order quantity Q^* as a function of the downside-risk weight λ, using $\alpha = 0.95$.

Figure 4.6 shows how the optimal policy changes as λ changes. Near $\lambda = 0$, the optimal quantity is close to the risk-neutral benchmark. As λ rises, the optimal quantity shifts. The stepwise appearance is due to numerical grid search; the key insight is the policy movement, not the steps themselves.

The role of α is subtler. In this chapter, α is the confidence level, so the corresponding tail probability is $1 - \alpha$. Thus:

- $\alpha = 0.95$ means the worst 5% tail;
- $\alpha = 0.99$ means the worst 1% tail.

Changing α changes which part of the tail the decision maker cares about. A smaller α focuses on a broader and less extreme tail. A larger α focuses on a deeper and more extreme part of the loss distribution.

Figure 4.7 shows how the shape of the optimization problem itself changes when the confidence level changes. Different definitions of the relevant adverse tail lead to different objective functions and, therefore,ß potentially different recommended policies.

Figure 4.8 translates the previous figure into a direct policy statement. It shows that the confidence level is not a harmless technical parameter. Different choices of α imply different optimal order quantities because they emphasize different parts of the adverse tail.

Taken together, Figures 4.6, 4.7, and 4.8 show why sensitivity analysis is a central part of downside-risk modeling. The recommended policy depends not

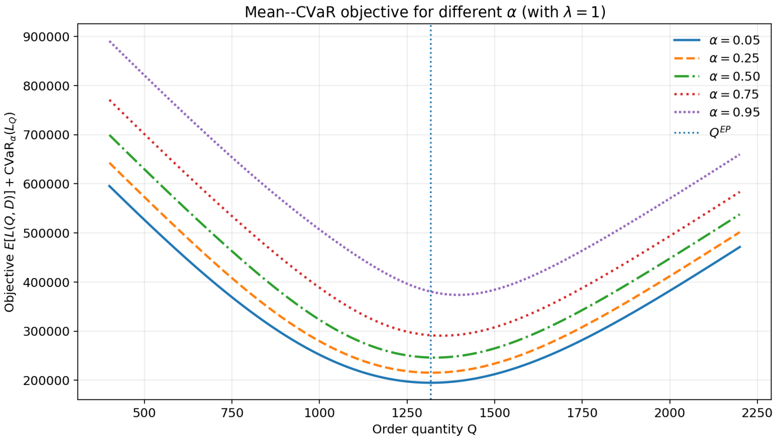

Figure 4.7: Mean–CVaR objective $E[L(Q,D)] + \text{CVaR}_\alpha(L_Q)$ as a function of the order quantity Q for different values of α, using $\lambda = 1$.

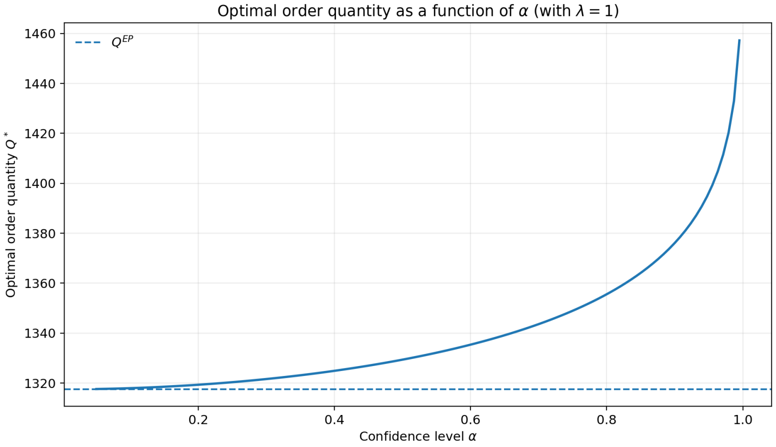

Figure 4.8: Optimal order quantity Q^* as a function of the confidence level α, using $\lambda = 1$.

only on the data and the economic environment, but also on how downside sensitivity is represented mathematically.

4.9 Summary and Key Takeaways

This chapter has shown how VaR and CVaR can be used to operationalize downside-sensitive choice in a one-period stochastic decision problem.

The chapter first introduced VaR and CVaR abstractly, then derived their closed-form expressions for a normally distributed loss variable. It clarified their interpretation and their properties, and explained how they differ from the utility-based notion of risk aversion developed in Chapter 3. It then distinguished policy evaluation from policy optimization. Finally, it applied the methodology to an original newsvendor example based on seasonal battery backup kits, derived the critical-ratio benchmark, evaluated that benchmark through VaR and CVaR, and showed how the optimal quantity changes once tail losses are penalized directly.

Key Takeaways

- VaR is a quantile-based tail threshold; CVaR is the average loss in the tail beyond that threshold.
- For a normal loss variable $Y \sim N(\mu, \sigma^2)$,

$$\mathrm{VaR}_\alpha(Y) = \mu + \sigma z_\alpha, \qquad \mathrm{CVaR}_\alpha(Y) = \mu + \sigma \frac{\phi(z_\alpha)}{1-\alpha}.$$

- VaR and CVaR are risk measures, not utility functions.
- The critical-ratio rule defines the risk-neutral benchmark in the newsvendor model.
- There is no equally simple closed-form "CVaR critical ratio" for the standard newsvendor problem; the downside-sensitive optimum must generally be solved numerically.
- Sensitivity to λ and α is a central managerial insight, not merely a numerical detail.

Bibliography

[1] Kenneth J. Arrow, Theodore Harris, and Jacob Marschak. Optimal inventory policy. *Econometrica*, 19(3):250–272, 1951.

[2] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.

[3] Gérard P. Cachon and Christian Terwiesch. *Matching Supply with Demand: An Introduction to Operations Management*. McGraw-Hill Education, New York, 3 edition, 2013.

[4] Sunil Chopra and Peter Meindl. *Supply Chain Management: Strategy, Planning, and Operation*. Pearson, Boston, 6 edition, 2016.

[5] Joel Goh and Melvyn Sim. Distributionally robust optimization and its tractable approximations. *Operations Research*, 58(4-part-1):902–917, 2010.

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

[7] Fernando S. Oliveira. Procurement risk management in a petroleum refinery. *Decision Sciences*, 54(3):277–296, 2023.

[8] Nicholas C. Petruzzi and Maqbool Dada. Pricing and the newsvendor problem: A review with extensions. *Operations Research*, 47(2):183–194, 1999.

[9] R. Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–42, 2000.

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

[11] Chuen-Teck See and Melvyn Sim. Robust approximation to multi-period inventory management. *Operations Research*, 58(3):583–594, 2010.