Risk Analytics

Machine Learning and Optimization
for Data-Driven Decision Making

Fernando S. Oliveira

Draft version — April 14, 2026

Chapter 5

Tail Risk Optimization with Discrete Distributions

Chapter 5

Tail Risk Optimization with Discrete Distributions

5.1 Introduction

Chapter 4 introduced the newsvendor problem as a simple and powerful framework for analysing decisions under uncertainty. In that setting, the decision maker chooses an order quantity Q before uncertain demand D is realized. Once demand becomes known, the mismatch between Q and D generates a loss function $L(Q, D)$ [2, 3, 4, 1].

The present chapter examines how tail-risk measures behave when demand is discrete rather than continuous. This matters because in many operational settings demand is naturally discrete: the number of emergency requests, equipment failures, service calls, contamination incidents, or spare-part needs. When demand is discrete, the induced loss distribution is also discrete, and that changes how downside risk must be interpreted and measured.

As in Chapter 4, the two central tail-sensitive risk measures are Value at Risk (VaR) and Conditional Value at Risk (CVaR). VaR identifies where the adverse tail of the loss distribution begins. CVaR measures how severe that tail is on average. In a continuous distribution, this distinction is conceptually straightforward because the cumulative distribution function is smooth. In a discrete distribution, however, the cumulative distribution function jumps. The VaR threshold may therefore coincide with a support point carrying positive probability mass. This means that the event

$$L(Q, D) \geq \mathrm{VaR}_\alpha(L(Q, D))$$

may contain more than the intended upper-tail probability $1 - \alpha$. That is why the discrete case requires more care than the continuous one [9, 5, 7].

The chapter studies two discrete probability models that are especially useful in risk analytics. The first is the binomial distribution, which is appropriate when demand is bounded by a fixed number of opportunities.

The second is the Poisson distribution, which is appropriate when demand arises through rare arrivals over time [8, 10, 6]. These two models therefore represent two different structures of uncertainty: bounded discrete demand and unbounded rare-arrival demand.

To preserve the Chapter 4 logic while using original applications, the two worked examples are both newsvendor-style settings. The first concerns emergency medical drone kits that must be pre-positioned before a high-risk holiday weekend in a mountain region. Demand is bounded and is modeled with a binomial distribution. The second concerns backup portable generators that must be stocked for municipal emergency depots during a storm season. Demand arises through outage-related requests over time and is modeled with a Poisson distribution. In both cases, the decision variable is the stock level Q, and the purpose is to understand how tail-sensitive criteria alter the recommended policy.

The chapter proceeds as follows. Section 5.2 develops VaR, TCE, and exact discrete CVaR for the induced loss $L(Q, D)$. Section 5.3 applies the framework to the binomial-demand problem and gives explicit formulas for VaR, TCE, and CVaR in that setting. Section 5.4 does the same for the Poisson case. Section 5.5 compares the two models. Section 5.6 summarizes the main lessons.

5.2 VaR, TCE, and CVaR for Discrete Losses

As in Chapter 4, the decision maker chooses an order quantity Q before demand D is observed. Once demand becomes known, the mismatch between the prepared quantity and the required quantity generates a loss $L(Q, D)$. In the present chapter, D is discrete, so the induced loss distribution is also discrete.

For a confidence level $\alpha \in (0, 1)$, Value at Risk is defined as

$$\mathrm{VaR}_\alpha(L(Q, D)) = \inf\{\ell \in \mathbb{R} : P(L(Q, D) \le \ell) \ge \alpha\}.$$

Thus, $\mathrm{VaR}_\alpha(L(Q, D))$ is the smallest threshold such that losses do not exceed that value with probability at least α. It identifies where the adverse tail begins.

A useful first tail measure is the tail conditional expectation,

$$\mathrm{TCE}_\alpha(L(Q, D)) = E[L(Q, D) \mid L(Q, D) \ge \mathrm{VaR}_\alpha(L(Q, D))] .$$

TCE answers the intuitive question: if the realized loss is at least as large as the VaR threshold, what is the average loss?

© 2026 Fernando S. Oliveira

In continuous settings, TCE and CVaR often coincide under standard conditions. In discrete settings, however, they need not. The reason is that the event

$$L(Q,D) \geq \mathrm{VaR}_\alpha(L(Q,D))$$

may contain more probability mass than the intended upper tail $1-\alpha$, because the VaR point itself may carry positive probability mass. TCE then averages over a tail set that is too large.

Exact discrete CVaR corrects for this by averaging over the worst exactly $1-\alpha$ probability mass rather than over the whole event $L(Q,D) \geq \mathrm{VaR}_\alpha(L(Q,D))$. Let

$$v_\alpha(Q) = \mathrm{VaR}_\alpha(L(Q,D)), \qquad F_{L(Q,D)}(\ell) = P(L(Q,D) \leq \ell).$$

Then exact discrete CVaR can be written as

$$\mathrm{CVaR}_\alpha(L(Q,D)) = \frac{E\big[L(Q,D)\mathbf{1}_{\{L(Q,D)>v_\alpha(Q)\}}\big] + v_\alpha(Q)\big(F_{L(Q,D)}(v_\alpha(Q)) - \alpha\big)}{1-\alpha}.$$

This expression has a simple interpretation. All losses strictly above the VaR threshold are included in full. At the VaR point itself, only the fraction of probability mass needed to complete the exact worst $1-\alpha$ tail is included. In other words, $F_{L(Q,D)}(v_\alpha(Q)) - \alpha$ is precisely the fraction of probability mass at the VaR point that must be retained in order to complete the exact tail. Exact discrete CVaR therefore differs from TCE because TCE averages over all realizations at or above the VaR threshold, whereas exact discrete CVaR averages over the worst exactly $1-\alpha$ probability mass [9, 5, 7].

A short toy example makes this distinction concrete. Suppose the loss distribution is

$$P(L=0) = 0.80, \qquad P(L=10) = 0.15, \qquad P(L=50) = 0.05.$$

At confidence level $\alpha = 0.90$, the VaR is

$$\mathrm{VaR}_{0.90}(L) = 10,$$

because cumulative probability first reaches 0.90 at the loss level 10. The tail conditional expectation is then

$$\mathrm{TCE}_{0.90}(L) = E[L \mid L \geq 10] = \frac{10(0.15) + 50(0.05)}{0.20} = 20.$$

But exact discrete CVaR averages over the worst exactly 10% of the distribution. The worst 5% is already given by the loss 50, and only half of the probability mass at $L = 10$ is needed to complete the tail. Thus

$$\mathrm{CVaR}_{0.90}(L) = \frac{50(0.05) + 10(0.05)}{0.10} = 30.$$

So TCE and exact CVaR differ because the atom at the VaR point is larger than the probability mass needed to complete the exact upper tail.

As in Chapter 4, these measures are most useful when they are embedded into a decision rule. A downside-sensitive newsvendor solves

$$Q^*_{\lambda,\alpha} \in \arg\min_{Q\in\mathcal{Q}} \left\{E[L(Q,D)] + \lambda\, \mathrm{CVaR}_\alpha(L(Q,D))\right\}, \qquad \lambda \geq 0.$$

When $\lambda = 0$, the decision reduces to expected-loss minimization. As λ increases, more weight is placed on severe upper-tail losses. The parameter α determines how deep into the tail the analysis goes.

The general definitions above apply to any discrete loss distribution. In the next two sections, they are specialized to the binomial and Poisson newsvendor settings.

5.3 Binomial Losses

5.3.1 When is the binomial model appropriate

The binomial distribution is appropriate when demand is bounded by a fixed number of opportunities. This occurs when each of a known number of locations, service zones, or potential users may or may not generate one unit of demand during the planning horizon. Demand above that upper bound is impossible.

5.3.2 Model setup

Consider a regional emergency-response authority that must pre-position compact medical drone kits before a high-risk holiday weekend. Let D denote the number of qualifying emergency requests over the weekend, and suppose

$$D \sim \mathrm{Bin}(n,p), \qquad n = 24, \quad p = 0.18.$$

The authority chooses stock level Q before demand is realized.

© 2026 Fernando S. Oliveira

Unused kits generate preparation, transport, and recertification costs. Unmet requests generate much larger shortage costs because substitute response is slower and more costly. Let

$$C_o = 40, \qquad C_u = 220.$$

Then the induced loss function is

$$L(Q, D) = C_o(Q - D)^+ + C_u(D - Q)^+.$$

For the binomial-demand case, define the loss attached to demand realization $k \in \{0, 1, \ldots, n\}$ as

$$\ell_k(Q) = C_o(Q - k)^+ + C_u(k - Q)^+.$$

The corresponding probability is

$$\pi_k = \binom{n}{k} p^k (1 - p)^{n-k}.$$

5.3.3 Specific VaR, TCE, and CVaR formulas for the binomial case

Because the newsvendor loss is piecewise linear and not globally monotone in k, the VaR is most naturally written in terms of the induced loss support. The α-level VaR is

$$\mathrm{VaR}_\alpha(L(Q, D)) = \inf \left\{ \ell \in \mathbb{R} : \sum_{k:\, \ell_k(Q) \le \ell} \binom{n}{k} p^k (1 - p)^{n-k} \ge \alpha \right\}.$$

This formula says: sort the loss realizations induced by the binomial demand distribution, accumulate their probabilities, and identify the first loss level at which cumulative probability reaches α.

The corresponding tail conditional expectation is

$$\mathrm{TCE}_\alpha(L(Q, D)) = \frac{\sum_{k:\, \ell_k(Q) \ge v_\alpha(Q)} \ell_k(Q) \binom{n}{k} p^k (1 - p)^{n-k}}{\sum_{k:\, \ell_k(Q) \ge v_\alpha(Q)} \binom{n}{k} p^k (1 - p)^{n-k}}, \qquad v_\alpha(Q) = \mathrm{VaR}_\alpha(L(Q, D)).$$

The exact discrete CVaR is

$$\mathrm{CVaR}_\alpha(L(Q, D)) = \frac{1}{1 - \alpha} \left[\sum_{k:\, \ell_k(Q) > v_\alpha(Q)} \ell_k(Q) \binom{n}{k} p^k (1 - p)^{n-k} + v_\alpha(Q) \left(\sum_{k:\, \ell_k(Q) \le v_\alpha(Q)} \binom{n}{k} p^k (1 - p)^{n-k} - \alpha \right) \right]. \tag{5.1}$$

These expressions are specific to the binomial model because the weights are binomial probabilities attached to a bounded support. The crucial structural feature is that demand cannot exceed n, so the induced loss distribution also has a finite upper tail. TCE and exact CVaR are still distinct when the VaR point carries positive probability mass, but the extreme tail is capped by construction.

5.3.4 A downside-sensitive decision model

The authority chooses Q by solving

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

Worked Example: Pre-positioning emergency medical drone kits

A regional emergency-response authority must decide how many compact medical drone kits to pre-position before a three-day holiday weekend in a mountain tourism region. Historical operational data suggest that the total number of qualifying requests can be approximated by a binomial demand model with 24 high-risk excursion zones and a zone-level request probability of 0.18.
The authority therefore models

$$D \sim \mathrm{Bin}(24, 0.18), \qquad L(Q,D) = 40(Q-D)^+ + 220(D-Q)^+.$$

Table 5.1 reports expected loss, VaR, CVaR, and the mean–CVaR objective for three candidate stock levels at the baseline pair

$$(\lambda, \alpha) = (0.6, 0.95).$$

The pattern is intuitive. A low stock level performs poorly in the upper tail because shortages are severe. A very high stock level performs poorly because of overage losses. The recommended policy therefore lies in an economically meaningful interior region.

5.3.5 Sensitivity analysis

Sensitivity analysis should follow the same style as Chapter 4. Four figures are especially useful. In the figures below, the objective is evaluated at the baseline pair

$$(\lambda, \alpha) = (0.6, 0.95),$$

Table 5.1: Illustrative performance measures for selected stock levels in the binomial-demand emergency-response example, with $\lambda = 0.6$ and $\alpha = 0.95$.

Q	$E[L(Q,D)]$	$\mathrm{VaR}_{0.95}$	$\mathrm{CVaR}_{0.95}$	$E[L] + \lambda\,\mathrm{CVaR}$
5	146.79	660.00	779.19	614.30
7	128.03	240.00	358.37	343.05
9	189.27	320.00	332.57	388.81

while the policy functions vary one parameter at a time.

The first shows the induced loss distribution for a representative stock level $Q = 7$, together with $\mathrm{VaR}_{0.95}$ and $\mathrm{CVaR}_{0.95}$. The horizontal axis is truncated at 1000 for readability.

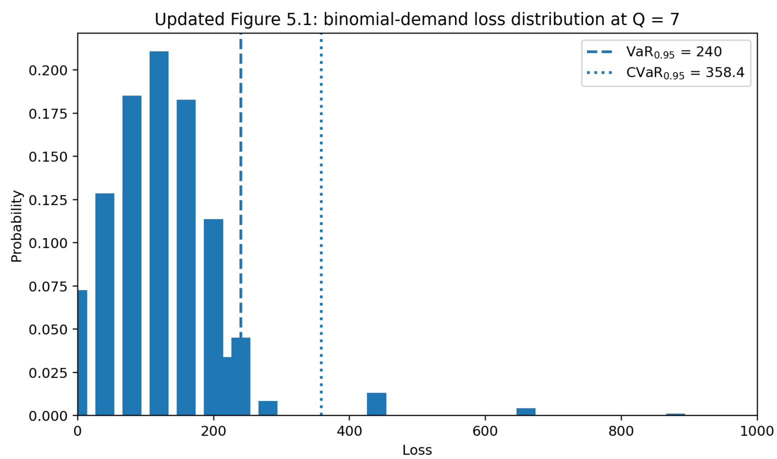

Figure 5.1: Loss distribution for the binomial-demand emergency-response example at $Q = 7$. The horizontal axis is truncated at 1000 for readability. The dashed line marks $\mathrm{VaR}_{0.95}(L(Q,D))$, while the dotted line marks $\mathrm{CVaR}_{0.95}(L(Q,D))$.

The second figure shows the mean–CVaR objective as a function of Q.

The third figure shows the optimal stock level as a function of λ, holding $\alpha = 0.95$ fixed.

The fourth figure shows the optimal stock level as a function of α, holding $\lambda = 0.6$ fixed.

Figures 5.1– 5.4 show that the binomial case involves a severe but finite tail. The main managerial problem is to protect against a concentrated but bounded surge in demand.

© 2026 Fernando S. Oliveira

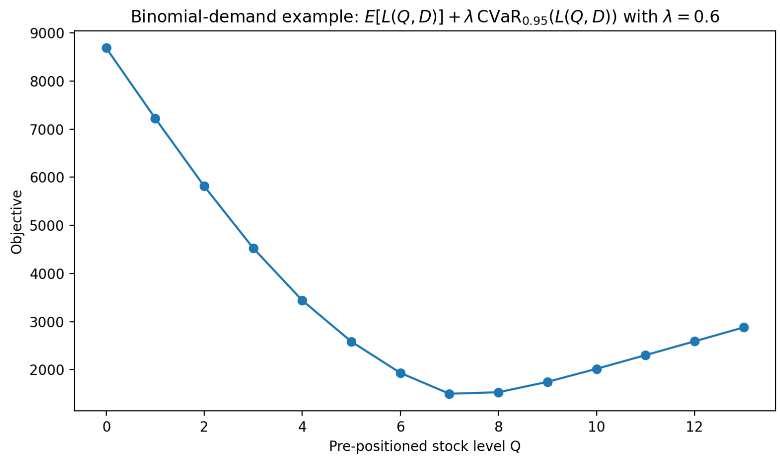

Figure 5.2: Mean–CVaR objective as a function of Q in the binomial-demand emergency-response example, evaluated at $(\lambda, \alpha) = (0.6, 0.95)$.

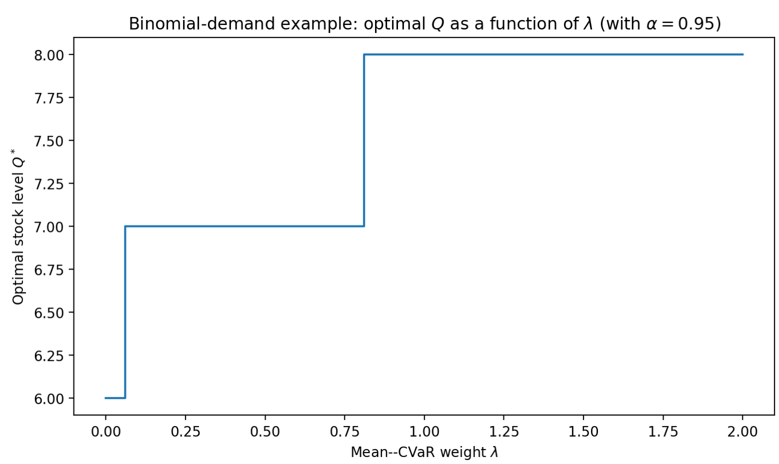

Figure 5.3: Optimal stock level Q^* as a function of λ in the binomial-demand emergency-response example, holding $\alpha = 0.95$ fixed.

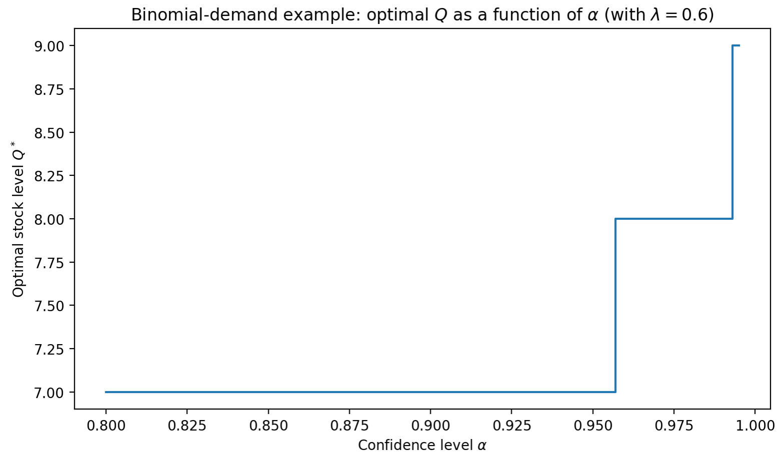

Figure 5.4: Optimal stock level Q^* as a function of α in the binomial-demand emergency-response example, holding $\lambda = 0.6$ fixed.

5.3.6 Managerial interpretation

The binomial model is appropriate when demand is bounded by a fixed number of opportunities. In the present example, demand cannot exceed the number of monitored high-risk zones. The practical implication is that the analyst is protecting against a surge that may be severe but is ultimately capped by the structure of the problem.

5.4 Poisson Losses

5.4.1 When is the Poisson model appropriate

The Poisson distribution is appropriate when demand arises through rare arrivals over time and there is no natural hard upper bound within the planning horizon. This occurs in emergency-service logistics, spare-parts planning, and rare-failure replacement systems. Unlike the binomial case, demand is not capped by a fixed number of opportunities.

Because the Poisson model is intended to represent a different structure of risk from the binomial case, it is helpful to use an application in which the expected count is not smaller than the binomial benchmark. Here the focus is a seasonal stock of backup portable generators for municipal emergency depots, where multiple outage-related requests may arise across the season.

5.4.2 Model setup

Consider a network of municipal emergency depots that must stock backup portable generators for outage-related incidents during a storm season. Let D denote the number of urgent generator requests during the planning horizon, and suppose

$$D \sim \text{Pois}(\mu), \qquad \mu = 5.2.$$

The authority chooses stock level Q before demand is realized.

Unused generators generate storage and maintenance costs, while shortages generate much larger response losses because emergency deployment delays disrupt critical local services. Let

$$C_o = 90, \qquad C_u = 520.$$

Then the induced loss function is

$$L(Q, D) = C_o(Q - D)^+ + C_u(D - Q)^+.$$

For the Poisson-demand case, define the loss attached to demand realization $k \in \{0, 1, 2, \dots\}$ as

$$\ell_k(Q) = C_o(Q - k)^+ + C_u(k - Q)^+.$$

The corresponding probability is

$$\pi_k = e^{-\mu}\frac{\mu^k}{k!}.$$

5.4.3 Specific VaR, TCE, and CVaR formulas for the Poisson case

Because the loss is again a non-monotone function of k, the VaR is most naturally written on the induced loss support. The α-level VaR is

$$\text{VaR}_\alpha(L(Q, D)) = \inf\left\{\ell \in \mathbb{R} : \sum_{k:\, \ell_k(Q) \le \ell} e^{-\mu}\frac{\mu^k}{k!} \ge \alpha\right\}.$$

The corresponding tail conditional expectation is

$$\text{TCE}_\alpha(L(Q, D)) = \frac{\sum_{k:\, \ell_k(Q) \ge v_\alpha(Q)} \ell_k(Q) e^{-\mu}\frac{\mu^k}{k!}}{\sum_{k:\, \ell_k(Q) \ge v_\alpha(Q)} e^{-\mu}\frac{\mu^k}{k!}}, \qquad v_\alpha(Q) = \text{VaR}_\alpha(L(Q, D)).$$

© 2026 Fernando S. Oliveira

The exact discrete CVaR is

$$\mathrm{CVaR}_\alpha(L(Q,D)) = \frac{1}{1-\alpha}\left[\sum_{k:\,\ell_k(Q)>v_\alpha(Q)} \ell_k(Q)e^{-\mu}\frac{\mu^k}{k!} + v_\alpha(Q)\left(\sum_{k:\,\ell_k(Q)\le v_\alpha(Q)} e^{-\mu}\frac{\mu^k}{k!} - \alpha\right)\right]. \tag{5.2}$$

These expressions are specific to the Poisson model because the weights are Poisson probabilities attached to an unbounded support. This is the main structural contrast with the binomial case. Even when the expected demand is moderate, the probability model still permits increasingly large realizations, and that is exactly why the upper tail can remain influential far above the mean.

5.4.4 A downside-sensitive decision model

The authority chooses Q by solving

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

Worked Example: Stocking backup portable generators

A network of municipal emergency depots must stock backup portable generators for outage-related incidents during a storm season. Historical records suggest that the number of urgent generator requests over the season can be approximated by

$$D \sim \mathrm{Pois}(5.2).$$

The choice of $\mu = 5.2$ is deliberate. It makes the Poisson example clearly different from the binomial case by ensuring that the expected count is not smaller than the binomial benchmark $24 \times 0.18 = 4.32$. More importantly, it fits the interpretation of the Poisson model: requests arrive across the season as outage-related incidents occur over time, rather than being limited by a fixed number of potential sources. Unused generators generate storage and maintenance costs, while shortages generate much larger response losses because delays can disrupt critical local services. The mismatch loss is

$$L(Q,D) = 90(Q-D)^+ + 520(D-Q)^+.$$

© 2026 Fernando S. Oliveira

Table 5.2 reports expected loss, VaR, CVaR, and the mean–CVaR objective for three candidate stock levels at the baseline pair

$$(\lambda, \alpha) = (0.5, 0.95).$$

The selected values $Q = 7, 8, 9$ are especially informative because they lie around the most relevant part of the objective function and reveal the local trade-off between average loss and tail protection more clearly than a table centered farther away from the baseline optimum.

Table 5.2: Illustrative performance measures for selected stock levels in the Poisson-demand emergency-generator example, with $\lambda = 0.5$ and $\alpha = 0.95$.

Q	$E[L(Q,D)]$	$\text{VaR}_{0.95}$	$\text{CVaR}_{0.95}$	$E[L] + \lambda\,\text{CVaR}$
7	348.66	1040.00	1757.52	1227.42
8	344.07	630.00	1270.17	979.15
9	384.09	720.00	964.04	866.11

5.4.5 Sensitivity analysis

Sensitivity analysis again follows the Chapter 4 logic, but the Poisson case should now produce a clearly different pattern because the support of demand is unbounded and shortage losses remain influential farther into the tail. In the figures below, the objective is evaluated at the baseline pair

$$(\lambda, \alpha) = (0.5, 0.95),$$

while the policy functions vary one parameter at a time.

The first figure shows the induced loss distribution for a representative stock level $Q = 8$, with the horizontal axis truncated for readability.

The second figure shows the mean–CVaR objective as a function of Q.

The third figure shows the optimal stock level as a function of λ, holding $\alpha = 0.95$ fixed.

The fourth figure shows the optimal stock level as a function of α, holding $\lambda = 0.5$ fixed.

Figures 5.5– 5.8 show a more visibly different pattern from the binomial case. The recommended stock level moves more strongly with both λ and α, reflecting stronger sensitivity to open-ended upper-tail exposure.

© 2026 Fernando S. Oliveira

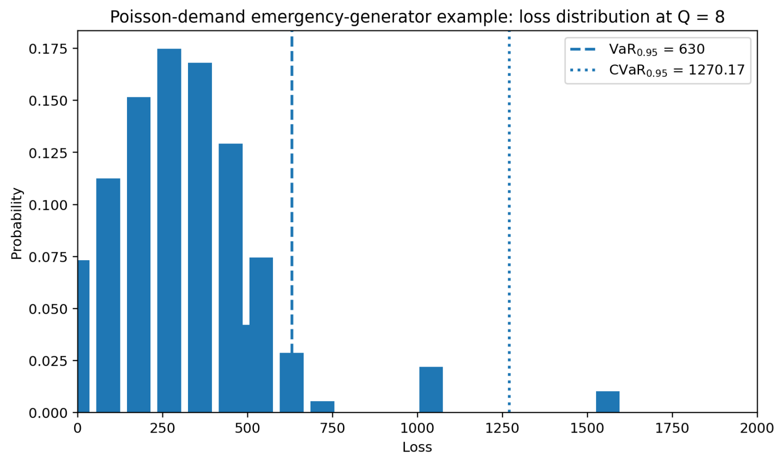

Figure 5.5: Loss distribution for the Poisson-demand emergency-generator example at $Q = 8$. The horizontal axis is truncated at 2000 for readability. The dashed line marks $\mathrm{VaR}_{0.95}(L(Q, D))$, while the dotted line marks $\mathrm{CVaR}_{0.95}(L(Q, D))$.

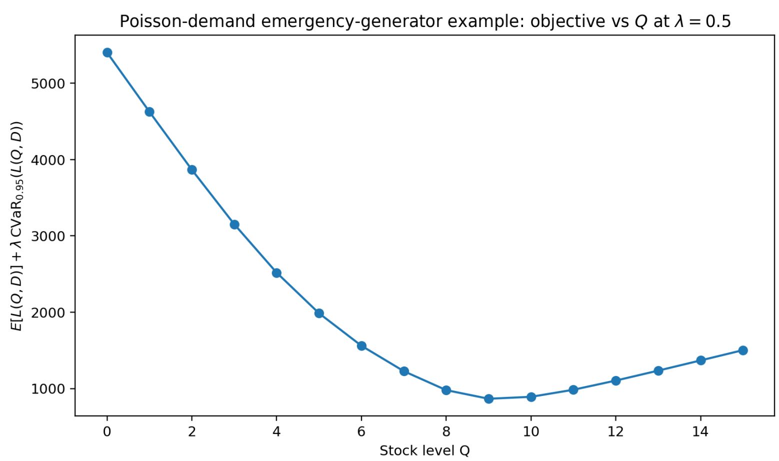

Figure 5.6: Mean–CVaR objective as a function of Q in the Poisson-demand emergency-generator example, evaluated at $(\lambda, \alpha) = (0.5, 0.95)$.

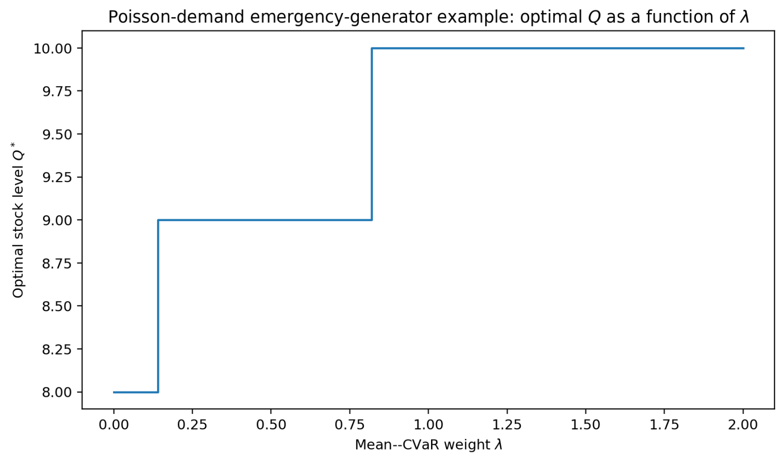

Figure 5.7: Optimal stock level Q^* as a function of λ in the Poisson-demand emergency-generator example, holding $\alpha = 0.95$ fixed.

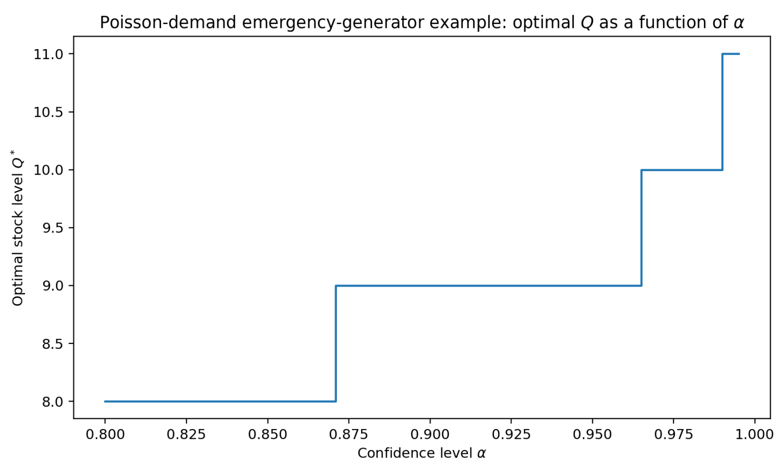

Figure 5.8: Optimal stock level Q^* as a function of α in the Poisson-demand emergency-generator example, holding $\lambda = 0.5$ fixed.

© 2026 Fernando S. Oliveira

5.4.6 Managerial interpretation

The Poisson model is appropriate when demand arises through rare arrivals over time rather than through a bounded set of opportunities. The practical implication is that the analyst is protecting against rare shortage episodes that may escalate much farther than in the bounded binomial case.

5.5 Comparing Binomial and Poisson Tail Risk

The binomial and Poisson models are both discrete demand models, but they capture different structures of uncertainty. In the binomial case, the support is bounded, so the extreme upper tail is finite. In the Poisson case, the support is unbounded, so the tail remains open-ended.

Table 5.3: Structural comparison of the two discrete-demand settings.

Feature	Binomial case	Poisson case
Demand support	Bounded	Unbounded
Tail type	Severe but finite	Open-ended
Typical interpretation	Fixed opportunities	Rare arrivals over time
Effect on CVaR	Tail capped by support	Tail can remain influential far above the mean

This difference has direct implications for downside-sensitive decision making. In the binomial case, increasing λ or α may shift the optimal policy, but the range of extreme outcomes is capped by construction. In the Poisson case, the same increase often has a stronger effect because the upper tail remains exposed to increasingly large demand realizations.

In practice, this means that binomial demand typically calls for protection against concentrated but capped surges, whereas Poisson demand calls for protection against rare but open-ended shortage episodes. This is the central managerial contrast between the two models.

5.6 Summary and Key Takeaways

This chapter has extended the logic of Chapter 4 from continuous or smoothly approximated losses to discrete loss distributions. Its central message is that VaR and CVaR remain meaningful and useful in discrete settings, but the

© 2026 Fernando S. Oliveira

discrete case requires additional care because the loss distribution places positive probability mass on individual support points.

The chapter first reintroduced the Chapter 4 newsvendor framework and then developed VaR, TCE, and exact discrete CVaR for the induced loss $L(Q, D)$. It explained why the discrete case is more delicate than the continuous one: the exact upper tail may require only a fraction of the probability mass at the VaR point, so TCE and exact CVaR need not coincide.

The chapter then applied the framework to two original examples. The binomial application modeled bounded weekend demand for emergency medical drone kits. The Poisson application modeled seasonal backup-generator demand at municipal emergency depots under outage-related conditions. In both cases, the decision variable Q directly shaped the loss distribution, so VaR and CVaR became part of an endogenous decision problem rather than merely descriptive statistics.

The main practical lesson is the same as in Chapter 4, but now in a discrete setting: expected loss alone is often insufficient. VaR identifies where the adverse tail begins, but CVaR provides the richer picture when the decision maker cares about the average severity of the worst outcomes. In this sense, Chapter 5 does not replace Chapter 4's analysis of downside risk, but extends it from continuous approximations to discrete operational environments in which the structure of probability mass itself matters for decision making.

Key Takeaways

- $\text{VaR}_\alpha(L(Q, D))$ identifies where the adverse tail begins.
- $\text{TCE}_\alpha(L(Q, D))$ averages losses at or beyond the VaR threshold, but need not coincide with exact CVaR in discrete models.
- $\text{CVaR}_\alpha(L(Q, D))$ measures the average loss in the exact worst $1 - \alpha$ probability mass.
- A binomial demand model is appropriate when demand is bounded by a fixed number of opportunities.
- A Poisson demand model is appropriate when demand arises through rare arrivals over time.
- In both cases, VaR and CVaR are most useful when embedded into a decision problem over Q.

© 2026 Fernando S. Oliveira

- Sensitivity to λ and α is a central managerial insight, not merely a numerical detail.

Bibliography

[1] Edward J. Anderson. *Business Risk Management: Models and Analysis*. John Wiley & Sons, Ltd, Chichester, West Sussex, UK, 2014.

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

[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] Rainer Göb. Estimating value at risk and conditional value at risk for count variables. *Quality and Reliability Engineering International*, 27(5):659–672, 2011.

[6] Matthew J. Hassett, Donald G. Stewart, and Jelena Milovanovic. *Probability for Risk Management*. ACTEX Learning, Greenland, NH, 3 edition, 2021.

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

[8] Emanuel Parzen. *Modern Probability Theory and Its Applications*. John Wiley & Sons, Inc., New York, 1960.

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

[10] Sheldon M. Ross. *Introduction to Probability Models*. Academic Press, Oxford, 11 edition, 2014.