On service level measures in stochastic inventory control

Dr. Roberto Rossi

The University of Edinburgh Business School, The University of Edinburgh, UK
roberto.rossi@ed.ac.uk

Friday, June the 21th, 2013

Decision making under uncertainty

Key elements

A problem of **decision making under uncertainty** features:

- Decision variables
- Random variables
- A set of decision/observation stages
- A set of stochastic constraints, e.g. chance constraints
- Objective function

We assume that randomized decisions at a given stage — i.e. **randomized policies** — are **forbidden**.

Decision making under uncertainty

Decision trees

To gain insights into the issues we are going to discuss next, we adopt a **graphical depiction** of the problem: “decision trees”.

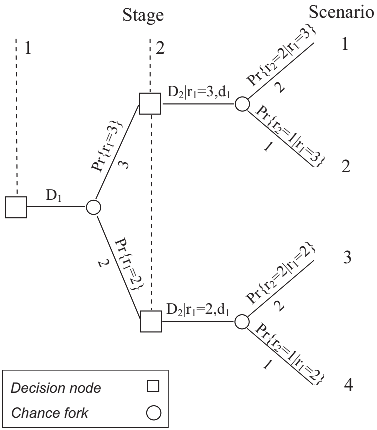

Decision making under uncertainty

Service level measures

In this work we are concerned with **service level measures** in decision making under uncertainty.

Decision making under uncertainty

Service level measures

In this work we are concerned with **service level measures** in decision making under uncertainty.

How do we **measure** a service level?

Decision making under uncertainty

Service level measures

In this work we are concerned with **service level measures** in decision making under uncertainty.

How do we **measure** a service level?

A common service level measure in stochastic decision making takes the form of a **chance constraint**.

A chance constraint is a particular type of stochastic constraint that must be satisfied according to a **prescribed probability**.

A motivating example

Consider the **stochastic constraint**

$$D_1 - r_1 + D_2 - r_2 = 0$$

r_1 and r_2 are independent and the two values in their support are equally likely.

Under a chance constraint

$$\Pr\{D_1 - r_1 + D_2 - r_2 \geq 0\} \geq 0.25$$

the assignment
$D_1 = 2$,
$(D_2|r_1 = 3, D_1) = 2$,
$(D_2|r_1 = 2, D_1) = 0$
is feasible.

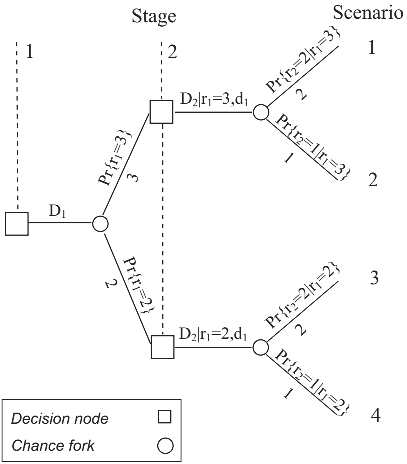

A motivating example

In general, we can **condition** a probability to a random variable.

Let E be an event, we may write

$$\Pr\{E|r_1\}$$

this is a function which takes value $\Pr\{E|r_1 = i\}$ when $r_1 = i$.

Enforcing a chance constraint such as

$$\Pr\{E|r_1\} \geq \alpha$$

means making sure that this function does not take a value less than α **for each value in the support** of r_1.

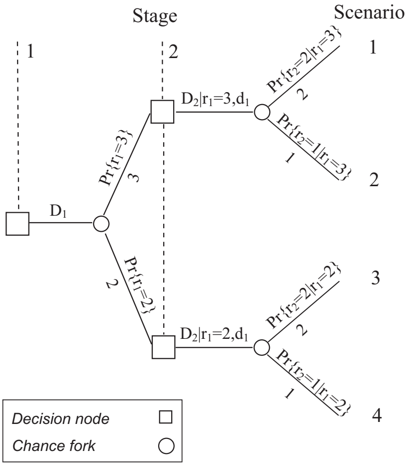

A motivating example

In practical situations chance constraints such as

$$\Pr\{D_1 - r_1 + D_2 - r_2 \geq 0 | r_1\} \geq \alpha$$

may become relevant.

Under this chance constraint, the previous assignment
$D_1 = 2$,
$(D_2 | r_1 = 3, D_1) = 2$,
$(D_2 | r_1 = 2, D_1) = 0$
would be infeasible.

In what follows we will discuss why these service level measures become particularly relevant in stochastic inventory control.

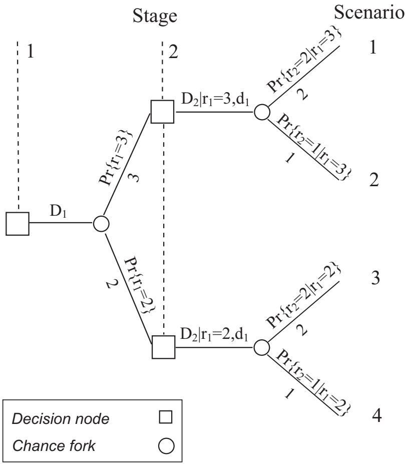

The periodic review production/inventory problem

In a periodic review system **inventory is reviewed only at discrete points in time**. We review inventory only at the beginning and at the end of a period.

Orders can be placed only at the beginning of a period.

Demand is a random variable d with known distribution.

The **delivery lead-time** is constant and equal to L periods.

Unmet demand is **backordered** and fulfilled as soon as a replenishment arrives.

There is a **service level constraint** enforcing a specified probability α of no-stockout per period — α service level.

A **holding cost** of $\$h$ per period is paid for each unit carried in stock to the next period.

The periodic review production/inventory problem

Numerical example

We consider a **planning horizon** comprising 4 periods.

In each period we observe a **random demand** that follows a **discrete distribution**. The **probability mass functions** of the demand in each period are the following.

$$\begin{aligned}
\text{pmf}(d_1) &= \{18(0.5), 26(0.5)\}\\
\text{pmf}(d_2) &= \{52(0.5), 6(0.5)\}\\
\text{pmf}(d_3) &= \{9(0.5), 43(0.5)\}\\
\text{pmf}(d_4) &= \{20(0.5), 11(0.5)\}
\end{aligned}$$

The **delivery lead time** is set to 0.

The **holding cost** h is set to \$10

The prescribed **no stockout probability**α is 0.85.

The complete set of **scenarios** is presented below

Scenario	d_1	d_2	d_3	d_4	Probability
1	18	52	9	20	0.0625
2	18	52	9	11	0.0625
3	18	52	43	20	0.0625
4	18	52	43	11	0.0625
5	18	6	9	20	0.0625
6	18	6	9	11	0.0625
7	18	6	43	20	0.0625
8	18	6	43	11	0.0625
9	26	52	9	20	0.0625
10	26	52	9	11	0.0625
11	26	52	43	20	0.0625
12	26	52	43	11	0.0625
13	26	6	9	20	0.0625
14	26	6	9	11	0.0625
15	26	6	43	20	0.0625
16	26	6	43	11	0.0625

The periodic review production/inventory problem

Numerical example

Let I_t denote the **inventory level** — i.e. on hand stock minus backorders — at the end of period t.

If the **event of interest** is a no stockout in period t, i.e.

$$I_t \geq 0 \leftrightarrow I_{t-1} + Q_t - d_t \geq 0$$

then we may enforce either the constraint

$$\Pr\{I_{t-1} + Q_t \geq d_t\} \geq \alpha$$

or the constraint

$$\Pr\{I_{t-1} + Q_t \geq d_t | I_{t-1}\} \geq \alpha$$

The **optimal solution**, obtained by a trivial **scenario based MILP model**, can be represented by means of a **decision tree**.

The periodic review production/inventory problem

Numerical example

Optimal solution under service measure $\Pr\{I_t \geq 0\} \geq \alpha$

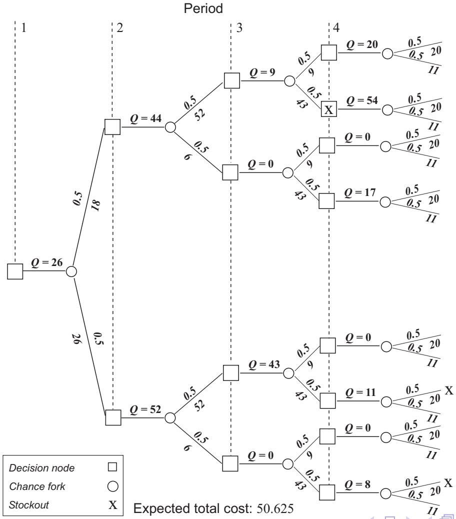

The periodic review production/inventory problem

Numerical example

Optimal solution under service measure $\Pr\{I_t \geq 0 | I_{t-1}\} \geq \alpha$

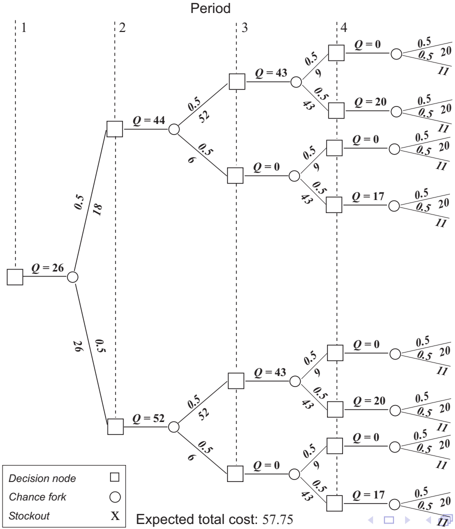

Expected total cost: 57.75

The periodic review production/inventory problem

Numerical example

Optimal solution under service measure $\Pr\{I_t \geq 0 | I_{t-1}\} \geq \alpha$ takes the form of a "base stock" policy with base stock levels 26, 52, 43, 20 in period 1, 2, 3, and 4, respectively.

In a **base stock policy**, an order is placed as soon as the inventory position drops below the "base stock level". The **inventory position** comprises items in stocks minus backorders plus incoming orders not yet received.

In general, this example shows that **the optimal policy** under the service measure $\Pr\{I_t \geq 0\} \geq \alpha$ is **not a base stock policy.**

The **cost** of an optimal solution under service measure $\Pr\{I_t \geq 0 | I_{t-1}\} \geq \alpha$ is **higher**.

Optimal solution under service measure $\Pr\{I_t \geq 0 | I_{t-1}\} \geq \alpha$ takes the form of a "base stock" policy with base stock levels 26, 52, 43, 20 in period 1, 2, 3, and 4, respectively.

In a **base stock policy**, an order is placed as soon as the inventory position drops below the "base stock level". The **inventory position** comprises items in stocks minus backorders plus incoming orders not yet received.

In general, this example shows that **the optimal policy** under the service measure $\Pr\{I_t \geq 0\} \geq \alpha$ is **not a base stock policy.**

The **cost** of an optimal solution under service measure $\Pr\{I_t \geq 0 | I_{t-1}\} \geq \alpha$ is **higher**. *However, in certain cases this measure may better reflect contractual requirements and industrial practices.*

The continuous review production/inventory problem

In a continuous review production/inventory problem inventory is **monitored continuously** and orders can be placed at each time instant.

Demand — measured in units per period — is a random variable d with known distribution.

The **delivery lead-time** is constant and equal to L periods.

Unmet demand is **backordered** and fulfilled as soon as a replenishment arrives.

We consider a **service level constraint** enforcing a specified probability α of no-stockout over the replenishment lead time — α service level.

A **holding cost** of $\$h$ per period is paid for each unit carried in stock.

The α service level is generally **defined informally** as the

Continuous review:
"no-stockout probability over the replenishment lead time"; or

Periodic review:
"no-stockout probability per period".

We now aim to to **formalize** the service level definition for the continuous review case in **mathematical terms**.

The continuous review production/inventory problem

α service level

We denote the **inventory position** at time t as I_t^p.

Since demand is stationary we express the **service level** at time t as

$$\Pr\{I_{t+L} \geq 0\} \geq \alpha$$

where I_t is a **random variable** representing the inventory level, i.e. items in stocks minus backorders, at time t.

Let d_L denote the **demand over the lead-time**.

By exploiting the fact that the inventory position **tracks incoming orders**, we can rewrite our service level constraint as

$$\Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha \tag{1}$$

Alternatively, one may adopt the following chance constraint to express the service level constraint

$$\Pr\{I_t^p - d_L + Q_t \geq 0 | I_t\} \geq \alpha \tag{2}$$

The continuous review production/inventory problem

α service level: $\Pr\{I_t^p - d_L + Q_t \geq 0 | I_t\} \geq \alpha$

When the service level is formulated as

$$\Pr\{I_t^p - d_L + Q_t \geq 0 | I_t\} \geq \alpha$$

since we condition the event $I_t^p - d_L + Q_t \geq 0$ on I_t, I_t^p **becomes a scalar value**.

The optimal order quantity can be immediately obtained by simply **inverting the cumulative distribution function** of the demand over lead-time

$$Q_t = \min\{Q | I_t^p + Q \geq \mathsf{cdf}_{d_L}^{-1}(\alpha)\}$$

where $\mathsf{cdf}_{d_L}^{-1}(\alpha)$ denotes the inverse cumulative distribution of d_L.

This shows that, under the cost structure discussed above and this service level measure, a **base stock policy** with base stock level

$$S = \min\{s | s \geq \mathsf{cdf}_{d_L}^{-1}(\alpha)\} \tag{3}$$

becomes optimal regardless of the nature of the demand distribution, i.e. continuous or discrete, as long as randomized policies are forbidden.

The continuous review production/inventory problem

However, if the service level is

$$\Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$$

a base stock policy is not optimal, in general, when demand follows a discrete distribution.

Furthermore, if demand follows a continuous distribution, a base stock policy is **only optimal if specific conditions are met.**

The continuous review production/inventory problem

Discrete distributions ($\Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$)

When **demand distribution is discrete**, it may not be possible to find a base stock level that guarantees exactly a service level α.

Example We consider a Poisson demand with rate $\lambda = 3$ units/period.

The lead time for an order is $L = 2$ periods. Hence demand over lead-time follows a Poisson distribution with rate λL.

Holding cost is $h = \$4$ per unit per period.

We enforce an α service level with $\alpha = 0.7$.

The optimal base stock level is $S = 7$, since this is the minimum base stock level for which the cumulative distribution of a poisson with rate λL exceeds $\alpha = 0.7$.

However, the **actual service level** associated with this base stock level is **0.743**. The expected total cost per period of the optimal base stock policy is $ 9.14.

The continuous review production/inventory problem

Discrete distributions ($\Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$)

It is possible to devise **a better control policy** by allowing the base stock level to **depend on the demand realizations in previous periods.**

Example (cont.) For instance, we may adopt a policy which "**orders up to 7** if demand in the past L periods has been less or equal to 7, and that **orders up to 6** otherwise."

This policy ensures a no-stockout probability per replenishment cycle of **0.708**, the expected total cost per period is \$ 8.82 $<$ \$ 9.14.

In fact, also **this policy is not optimal** and there may be a better "past demand over lead time dependent policy" which guarantees a service level of exactly $\alpha = 0.7$ at minimum cost. However, **finding such a policy is NP-hard.**

Theorem (1)

Under a discrete demand distribution finding a policy that guarantees exactly a given α service level is NP-hard.

Proof.

Reduction from 0-1 Knapsack (omitted). □

The continuous review production/inventory problem

Continuous distributions ($\Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$)

If demand follows a **continuous distribution**...

The continuous review production/inventory problem

Continuous distributions ($\Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$)

If demand follows a **continuous distribution**...a base stock policy is optimal **only if specific conditions are met.**

Theorem (2)

A base stock policy is optimal for the production/inventory problem with continuously distributed stationary stochastic demand under α *service level constraints formulated as*

$$\Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$$

only if

$$\alpha \frac{d\ cdf_{d_L}^{-1}(\alpha)}{d\alpha}$$

is increasing in α*.*

Proof.

next slide...

Proof.

(sketch) Consider the following function

$$b(S) = \int_0^S (S - i)\mathsf{pdf}_{d_L}(i)\mathsf{d}i$$

which computes the **expected buffer stock associated with a base stock level** S.

To prove that a base stock policy is optimal we must show that a demand or stock level dependent policy cannot produce a better cost.

This can be shown by considering the following function

$$h(\alpha) = \int_0^{\mathsf{cdf}_{d_L}^{-1}(\alpha)} (\mathsf{cdf}_{d_L}^{-1}(\alpha) - i)\mathsf{pdf}_{d_L}(i)\mathsf{d}i$$

that represents the **expected buffer stock as a function of the service level** α; and by showing that such a function is **convex** in α. □

Theorem (3)

If demand over leadtime follows an exponential distribution with parameter λ, $h(\alpha)$ is convex.

Proof.

In the paper... □

The continuous review production/inventory problem

Continuous distributions ($\Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$)

We now provide an example in which $h(\alpha)$ is **not convex** and it is possible to find a policy that **beats a base stock policy** for a given service level.

Example 4 We consider a continuous review system in which demand follows a **Beta distribution** with parameter $\alpha = \beta = 0.2$ and thus expected value $\alpha/(\alpha + \beta) = 0.5$ units/period.

The **lead time** for an order is $L = 1$ periods. **Holding cost** is $h = \$4$ per unit per period. We enforce an α **service level** with $\alpha = 0.7$.

It immediately follows that the **optimal base stock level** is $S = 0.94$, since this is the base stock level for which the cumulative distribution of the demand over lead time is **exactly equal** to $\alpha = 0.7$. The **expected total cost** per period of the optimal base stock policy is \$ 4.70 per period.

However, we now consider a policy that

"orders up to $S_1 = 0.9999$ if demand over lead time has been lower than 0.5 and that orders up to $S_2 = 0.4999$ otherwise."

This policy ensures a **service level of 0.7082**, that is slightly higher than the prescribed one. However, the **expected total cost** per period of this policy is only \$ 4.59.

Observation

Under a demand that follows a **continuous distribution** and that **satisfies Theorem 2**, the service level measures

$$\Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$$

and

$$\Pr\{I_t^p - d_L + Q_t \geq 0 | I_t\} \geq \alpha$$

provide the **same cost performance.**

Final considerations

When should we apply one or the other measure?

Stationary demand

- demand truly stationary, e.g. low-cost consumables (so called type "C" items), and accurately estimated over a long time span

Final considerations

When should we apply one or the other measure?

Stationary demand

- demand truly stationary, e.g. low-cost consumables (so called type "C" items), and accurately estimated over a long time span
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$

Final considerations

When should we apply one or the other measure?

Stationary demand

- demand truly stationary, e.g. low-cost consumables (so called type "C" items), and accurately estimated over a long time span
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$
- inventory systems controlled under a "rolling horizon" strategy with forecast updates

Final considerations

When should we apply one or the other measure?

Stationary demand

- demand truly stationary, e.g. low-cost consumables (so called type "C" items), and accurately estimated over a long time span
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$
- inventory systems controlled under a "rolling horizon" strategy with forecast updates
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0 | I_t\} \geq \alpha$

Final considerations

When should we apply one or the other measure?

Stationary demand

- demand truly stationary, e.g. low-cost consumables (so called type "C" items), and accurately estimated over a long time span
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$
- inventory systems controlled under a "rolling horizon" strategy with forecast updates
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0 | I_t\} \geq \alpha$

Nonstationary demand

- frequent and substantial forecast updates, e.g. launch of a new product

Final considerations

When should we apply one or the other measure?

Stationary demand

- demand truly stationary, e.g. low-cost consumables (so called type "C" items), and accurately estimated over a long time span
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$
- inventory systems controlled under a "rolling horizon" strategy with forecast updates
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0 | I_t\} \geq \alpha$

Nonstationary demand

- frequent and substantial forecast updates, e.g. launch of a new product
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0 | I_t\} \geq \alpha$

Final considerations

When should we apply one or the other measure?

Stationary demand

- demand truly stationary, e.g. low-cost consumables (so called type "C" items), and accurately estimated over a long time span
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$
- inventory systems controlled under a "rolling horizon" strategy with forecast updates
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0 | I_t\} \geq \alpha$

Nonstationary demand

- frequent and substantial forecast updates, e.g. launch of a new product
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0 | I_t\} \geq \alpha$
- demand process subject to little or no forecast updates

Final considerations

When should we apply one or the other measure?

Stationary demand

- demand truly stationary, e.g. low-cost consumables (so called type "C" items), and accurately estimated over a long time span
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$
- inventory systems controlled under a "rolling horizon" strategy with forecast updates
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0 | I_t\} \geq \alpha$

Nonstationary demand

- frequent and substantial forecast updates, e.g. launch of a new product
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0 | I_t\} \geq \alpha$
- demand process subject to little or no forecast updates
  $\rightarrow \Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha$

Questions