Piecewise linear approximations of the standard normal first order loss function and an application to stochastic inventory control

Dr. Roberto Rossi

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

IFORS, July, 2014

Introduction

Papers

This presentation illustrates results covered in the following works:

Roberto Rossi, S. Armagan Tarim, Brahim Hnich, and Steven D. Prestwich. Piecewise linear approximations of the standard normal first order loss function.
Applied Mathematics and Computation, arXiv:1307.1708, 2013

Roberto Rossi, Onur A. Kilic, and S. Armagan Tarim. Piecewise linear approximations for the static-dynamic uncertainty strategy in stochastic lot-sizing.
Submitted to Omega, arXiv:1307.5942, 2013

Other related works:

Roberto Rossi, S. Armagan Tarim, and Onur A. Kilic. A mixed integer linear programming heuristic for computing nonstationary (s,s) policy parameters.
In Hendrix E.M.T. and L. Liberti, editors, *Proceedings of the Malaga XII Global Optimization workshop*, pages xxx–xxx, 2014

Roberto Rossi and E.M.T. Hendrix. Computing linearisation parameters of arbitrarily distributed first order loss functions.
In Hendrix E.M.T. and L. Liberti, editors, *Proceedings of the Malaga XII Global Optimization workshop*, pages xxx–xxx, 2014

Introduction

Research questions

Our investigation tries to answer the following two key questions:

How can we produce “effective” **piecewise linearisations of the first order loss function**?

How can we employ these linearizations to **model in a seamless way** a number of variants of **the stochastic lot-sizing problem** under a **static-dynamic uncertainty** control policy, thus avoiding ad-hoc solutions?

The newsboy problem

Random demand

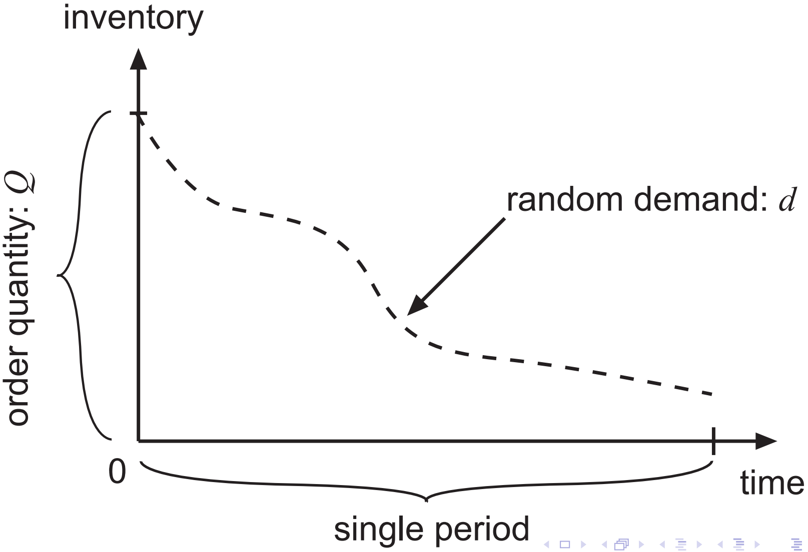

The newsboy problem

Random demand

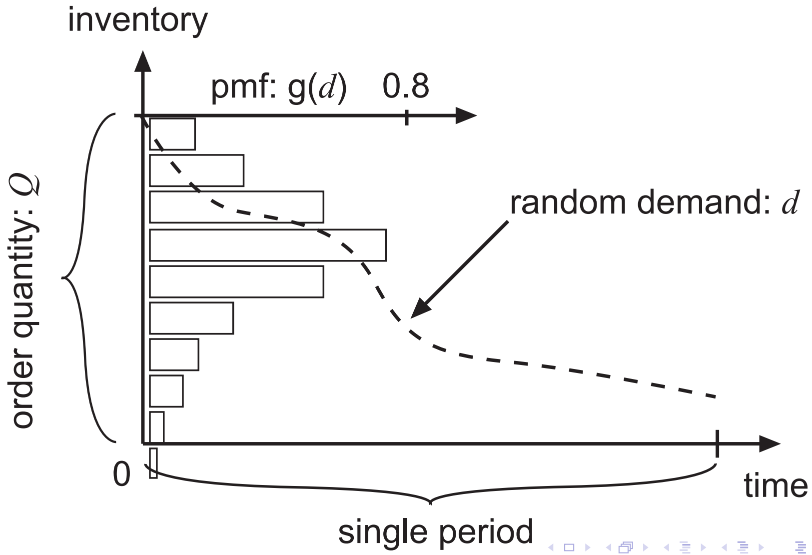

The newsboy problem

Random demand

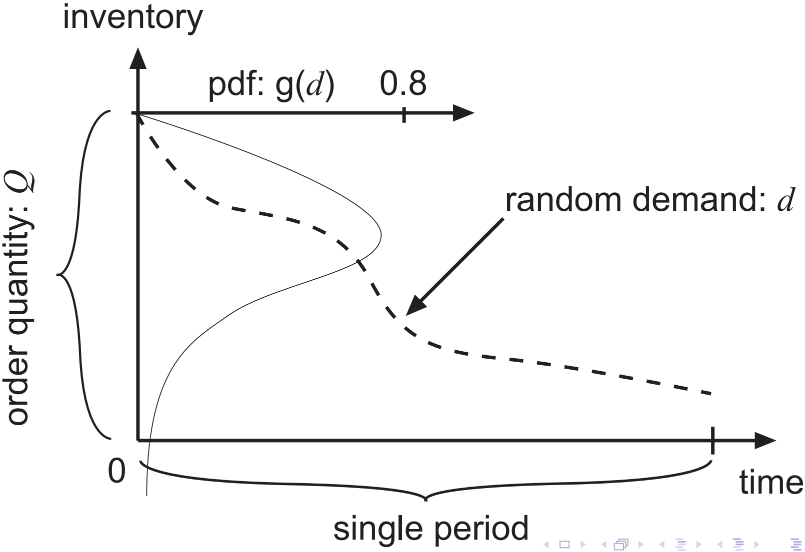

The newsboy problem

Cost structure under random demand

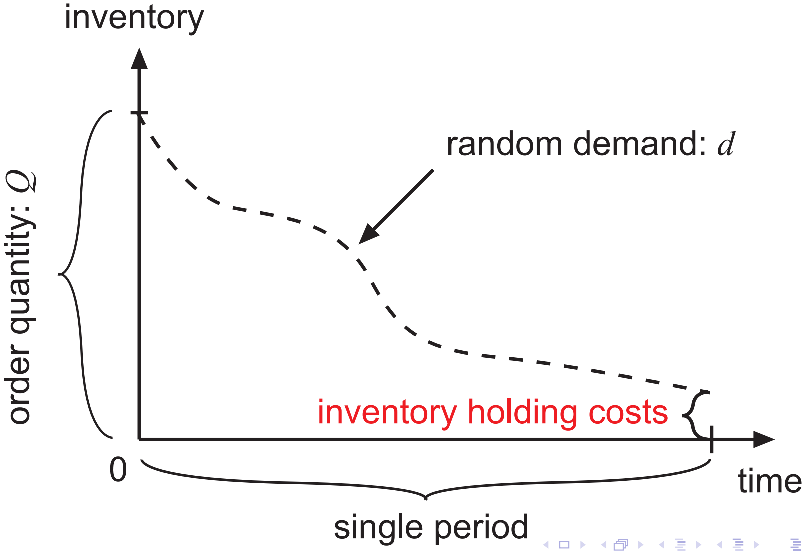

The newsboy problem

Cost structure under random demand

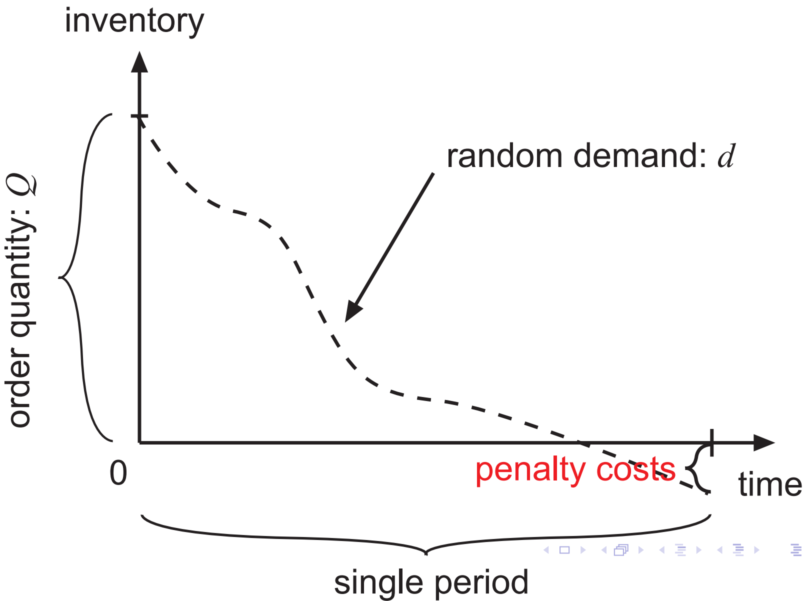

The newsboy problem

Cost structure under random demand

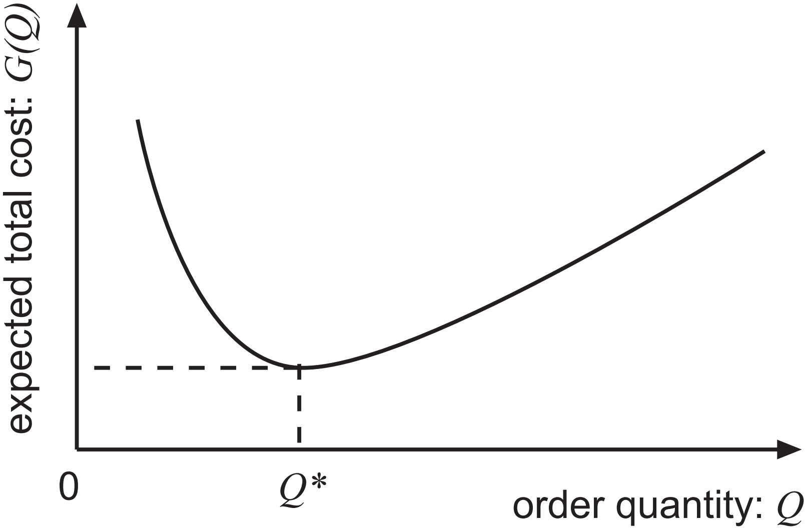

The newsboy problem

Mathematical formulation

Consider

- d: a **one-period** random demand that follows a **probability distribution** $f(d)$
- h: unit **holding cost**
- p: unit **penalty cost**

Let I be the end of period inventory and

$$g(I) = hI^{+} + pI^{-},$$

where $I^{+} = \max(I, 0)$ and $I^{-} = -\min(I, 0)$.

The **expected total cost** is $G(Q) = E[g(Q - d)]$, where $E[\cdot]$ denotes the expected value.

The newsboy problem

Mathematical formulation

Define:

$$E[I^+] = E[\max(Q-d,0)]: \quad \text{complementary first order loss function}$$
$$E[I^-] = E[\max(d-Q,0)]: \quad \text{first order loss function}$$

The **expected total cost** comprises two separable components

$$G(Q) = E[g(Q-d)] = hE[I^+] + pE[I^-]$$

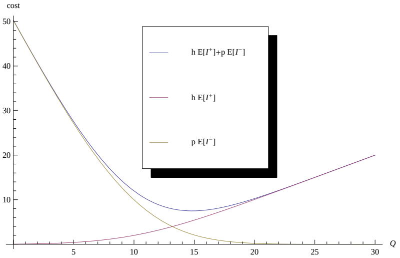

$d = \text{Normal}(10, 5)$
$h = \$1$
$p = \$5$

The first order loss function

A graphical outlook

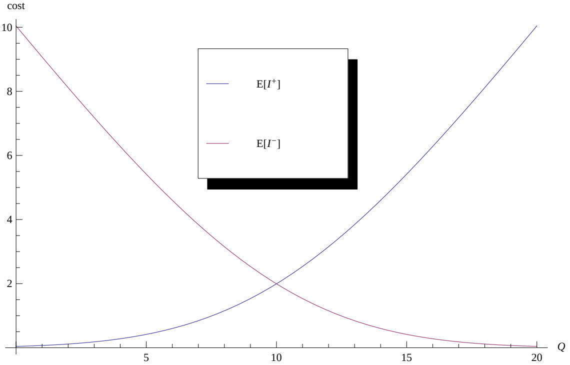

$E[I^+]$: complementary first order loss function
$E[I^-]$: first order loss function

The first order loss function

Properties

Consider a continuous random variable ω with support over $\mathbb{R}$, probability density function $g_\omega(x) : \mathbb{R} \to (0,1)$ and cumulative distribution function $G_\omega(x) : \mathbb{R} \to (0,1)$.

The **first order loss function** can be rewritten as

$$\mathcal{L}(x,\omega) = \int_{-\infty}^{\infty} \max(t-x,0) g_\omega(t)\, dt = \int_{x}^{\infty} (t-x) g_\omega(t)\, dt. \quad (1)$$

The **complementary first order loss function** can be rewritten as

$$\widehat{\mathcal{L}}(x,\omega) = \int_{-\infty}^{\infty} \max(x-t,0) g_\omega(t)\, dt = \int_{-\infty}^{x} (x-t) g_\omega(t)\, dt. \quad (2)$$

Lemma

$\mathcal{L}(x,\omega)$ and $\widehat{\mathcal{L}}(x,\omega)$ are convex in x.

The first order loss function

Properties

There is a close relationship between the first order loss function and the complementary first order loss function.

Lemma

The first order loss function $\mathcal{L}(x, \omega)$ *can also be expressed as*

$$\mathcal{L}(x, \omega) = \widehat{\mathcal{L}}(x, \omega) - (x - \tilde{\omega}) \qquad (3)$$

where $\tilde{\omega} = \mathsf{E}[\omega]$.

The first order loss function

Properties

Lemma

The first order loss function $\mathcal{L}(x,\omega)$ *can also be expressed as*

$$\mathcal{L}(x,\omega) = \int_x^\infty (1 - G_\omega(t))\ dt \tag{4}$$

Lemma

The complementary first order loss function $\widehat{\mathcal{L}}(x,\omega)$ *can also be expressed as*

$$\widehat{\mathcal{L}}(x,\omega) = \int_{-\infty}^x G_\omega(t)\, dt. \tag{5}$$

These two results are not easily derived from each other!

The first order loss function

Properties for symmetric distributions

Lemma

If the probability density function of ω is symmetric about a mean value $\tilde{\omega}$, then

$$\mathcal{L}(x,\omega) = \widehat{\mathcal{L}}(2\tilde{\omega} - x, \omega).$$

Lemma

If the probability density function of ω is symmetric about a mean value $\tilde{\omega}$, then

$$\widehat{\mathcal{L}}(x,\omega) = \widehat{\mathcal{L}}(2\tilde{\omega} - x, \omega) + (x - \tilde{\omega})$$

and

$$\mathcal{L}(x,\omega) = \mathcal{L}(2\tilde{\omega} - x, \omega) - (x - \tilde{\omega}).$$

The first order loss function

Properties for normal distribution

Let ζ be a normally distributed random variable with mean μ and standard deviation σ.

Lemma

The complementary first order loss function of ζ can be expressed in terms of the standard Normal cumulative distribution function as

$$\widehat{\mathcal{L}}(x,\zeta) = \sigma \int_{-\infty}^{\frac{x-\mu}{\sigma}} \Phi(t)\,dt = \sigma \widehat{\mathcal{L}}\left(\frac{x-\mu}{\sigma}, Z\right), \tag{6}$$

where Z is a standard Normal random variable.

Unfortunately, no closed form expression exists for $\widehat{\mathcal{L}}(x,Z)$.

The first order loss function

Non-linear approximations

Approximation to $\mathcal{L}(x, \zeta)$ have been recently discussed in

Steven K. De Schrijver, El-Houssaine Aghezzaf, and Hendrik Vanmaele. Double precision rational approximation algorithm for the inverse standard normal first order loss function.
Applied Mathematics and Computation, 219(3):1375–1382, October 2012

Steven K. De Schrijver, El-Houssaine Aghezzaf, and Hendrik Vanmaele. Double precision rational approximation algorithms for the standard normal first and second order loss functions.
Applied Mathematics and Computation, 219(4):2320–2330, November 2012

The first order loss function

Non-linear approximations

Drawbacks

Existing approximations are non-linear and cannot be easily embedded in MILP models — ad-hoc strategies are needed.

Existing approximations do not provide upper and lower bounds for $\mathcal{L}(x, \zeta)$ — it is hard to estimate the goodness of the solutions obtained and to obtain "linearisation gaps."

The first order loss function

Piecewise linear approximations

We introduce a well-known inequality from stochastic programming

Peter Kall and Stein W. Wallace. *Stochastic Programming (Wiley Interscience Series in Systems and Optimization)*.
John Wiley & Sons, August 1994, p. 167.

Theorem (Jensen's inequality)

Consider a random variable ω with support Ω and a function $f(x,s)$, which for a fixed x is convex for all $s \in \Omega$, then

$$E[f(x,\omega)] \geq f(x, E[\omega]).$$

The first order loss function

The newsboy problem & Jensen's inequality

For a fixed Q, the **total cost** is convex for all values in the support of d.

$$g_Q(d) = g(Q-d) = h\max(Q-d,0) + p\max(d-Q,0))$$

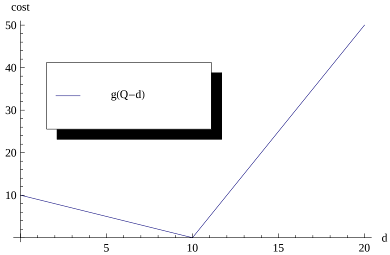

$Q = 10$
$h =$\$1
$p =$\$5

The first order loss function

The newsboy problem & Jensen's inequality

Define:

$$\begin{aligned} E[I^+] = E[\max(Q-d,0)]: &\quad \text{complementary first order loss function} \\ E[I^-] = E[\max(d-Q,0)]: &\quad \text{first order loss function} \end{aligned}$$

The **expected total cost** can be bounded from below as follows.

$$hE[I^+]+pE[I^-] \geq h\max(Q-E[d],0)+p\max(E[d]-Q,0) = g(Q-\mathsf{E}[d])$$

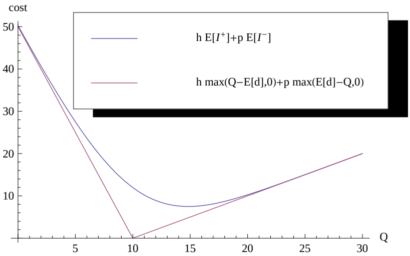

$d = \text{Normal}(10, 5)$
$\mathsf{E}[d] = 10$
$h =$ \$1
$p =$ \$5

The first order loss function

Bounding techniques

Define:

$$E[I^+] = E[\text{max}(Q - d, 0)]: \quad \text{complementary first order loss function}$$

The **complementary first order loss function** can be bounded from below as follows.

$$E[I^+] \geq h \max(Q - E[d], 0)$$

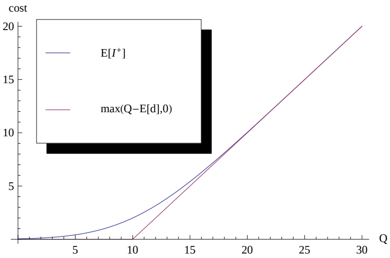

$d =$ Normal$(10, 5)$
E$[d] = 10$
$h =$\$1
$p =$\$5

The first order loss function

Bounding techniques

Define:

$$E[I^-] = E[\mathsf{max}(d - Q, 0)]: \quad \text{first order loss function}$$

The **first order loss function** can be bounded from below as follows.

$$E[I^-] \geq \max(E[d] - Q, 0)$$

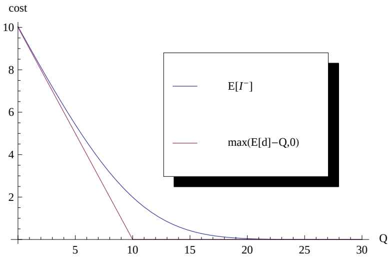

$d = \mathsf{Normal}(10, 5)$
$\mathsf{E}[d] = 10$
$h =$ \$1
$p =$ \$5

The first order loss function

Bounding techniques

Let $g_\omega(\cdot)$ denote the probability density function of ω and consider a partition of the support Ω of ω into N disjoint compact subregions $\Omega_1, \ldots, \Omega_N$. We define, for all $i = 1, \ldots, N$

$$p_i = \Pr\{\omega \in \Omega_i\} = \int_{\Omega_i} g_\omega(t)\, dt$$

$$\mathsf{E}[\omega|\Omega_i] = \frac{1}{p_i} \int_{\Omega_i} t g_\omega(t)\, dt$$

Theorem

$$E[f(x, \omega)] \geq \sum_{i=1}^{N} p_i f(x, E[\omega|\Omega_i])$$

The first order loss function

Bounding techniques

$$p_i = \Pr\{\omega \in \Omega_i\} = \int_{\Omega_i} g_\omega(t)\,dt$$

$$\mathsf{E}[\omega|\Omega_i] = \frac{1}{p_i}\int_{\Omega_i} t g_\omega(t)\,dt$$

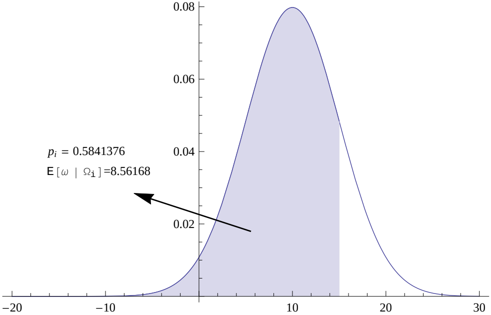

The first order loss function

Bounding techniques

$$p_i = \Pr\{\omega \in \Omega_i\} = \int_{\Omega_i} g_\omega(t)\, dt$$

$$\mathsf{E}[\omega|\Omega_i] = \frac{1}{p_i} \int_{\Omega_i} t g_\omega(t)\, dt$$

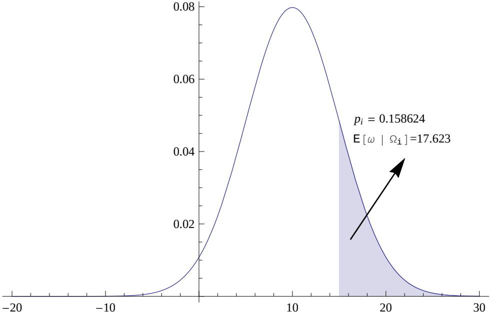

The first order loss function

Bounding techniques

For the (complementary) first order loss function ($\widehat{\mathcal{L}}_{lb}(x,\omega)$) $\mathcal{L}_{lb}(x,\omega)$ the lower bound

$$\mathsf{E}[f(x,\omega)] \geq \sum_{i=1}^{N} p_i f(x, \mathsf{E}[\omega|\Omega_i])$$

is a piecewise linear function with $N+1$ segments.

Consider the bound presented above and let $f(x,\omega) = \max(x-\omega, 0)$,

$$\widehat{\mathcal{L}}_{lb}(x,\omega) = \sum_{i=1}^{N} p_i \max(x - \mathsf{E}[\omega|\Omega_i], 0)$$

this function is equivalent to

$$\widehat{\mathcal{L}}_{lb}(x,\omega) = \begin{cases} 0 & -\infty \leq x \leq \mathsf{E}[\omega|\Omega_1] \\ p_1 x - p_1 \mathsf{E}[\omega|\Omega_1] & \mathsf{E}[\omega|\Omega_1] \leq x \leq \mathsf{E}[\omega|\Omega_2] \\ (p_1+p_2)x - (p_1\mathsf{E}[\omega|\Omega_1] + p_2\mathsf{E}[\omega|\Omega_2]) & \mathsf{E}[\omega|\Omega_2] \leq x \leq \mathsf{E}[\omega|\Omega_3] \\ \vdots & \vdots \\ (p_1+p_2+\ldots+p_N)x - (p_1\mathsf{E}[\omega|\Omega_1] + \ldots + p_N\mathsf{E}[\omega|\Omega_N]) & \mathsf{E}[\omega|\Omega_{N-1}] \leq x \leq \mathsf{E}[\omega|\Omega_N] \end{cases}$$

which is piecewise linear in x with breakpoints at $\mathsf{E}[\omega|\Omega_1], \mathsf{E}[\omega|\Omega_2], \ldots, \mathsf{E}[\omega|\Omega_N]$.

The first order loss function

Bounding techniques

$$p_i = \Pr\{\omega \in \Omega_i\} = \int_{\Omega_i} g_\omega(t)\,dt$$

$$\mathsf{E}[\omega|\Omega_i] = \frac{1}{p_i}\int_{\Omega_i} t g_\omega(t)\,dt$$

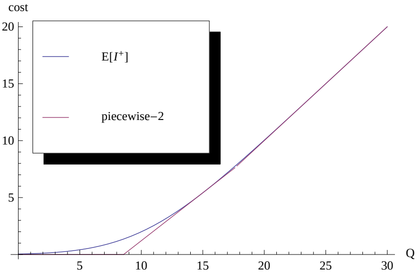

The first order loss function

Bounding techniques

$$p_i = \Pr\{\omega \in \Omega_i\} = \int_{\Omega_i} g_\omega(t)\, dt$$

$$\mathsf{E}[\omega|\Omega_i] = \frac{1}{p_i} \int_{\Omega_i} t g_\omega(t)\, dt$$

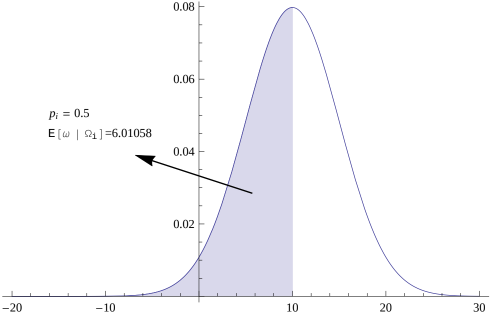

The first order loss function

Bounding techniques

$$p_i = \Pr\{\omega \in \Omega_i\} = \int_{\Omega_i} g_\omega(t)\, dt$$

$$\mathsf{E}[\omega|\Omega_i] = \frac{1}{p_i} \int_{\Omega_i} t g_\omega(t)\, dt$$

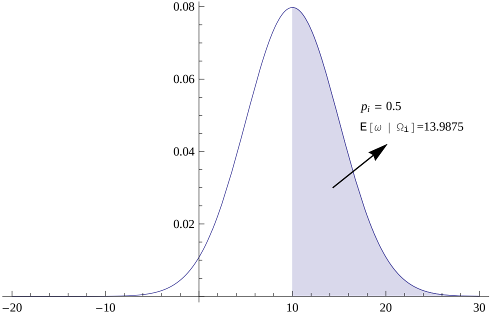

The first order loss function

Bounding techniques

$$p_i = \Pr\{\omega \in \Omega_i\} = \int_{\Omega_i} g_\omega(t)\, dt$$

$$\mathsf{E}[\omega|\Omega_i] = \frac{1}{p_i}\int_{\Omega_i} t g_\omega(t)\, dt$$

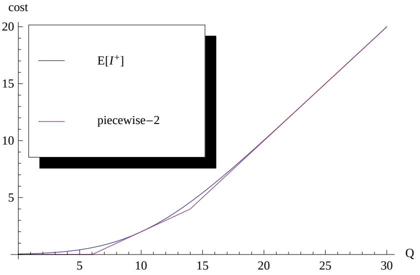

The first order loss function

Minimax optimal linearisation parameters for a standard normal random variable

Segments	Error		Piecewise linear approximation parameters									
		i	1	2	3	4	5	6	7	8	9	10
2	0.398942	b_i	∞									
		p_i	1									
		$\mathsf{E}[\omega\|\Omega_i]$	0									
3	0.120656	b_i	0	∞								
		p_i	0.5	0.5								
		$\mathsf{E}[\omega\|\Omega_i]$	-0.797885	0.797885								
4	0.0578441	b_i	-0.559725	0.559725	∞							
		p_i	0.287833	0.424333	0.287833							
		$\mathsf{E}[\omega\|\Omega_i]$	-1.18505	0	1.18505							
5	0.0339052	b_i	-0.886942	0	0.886942	∞						
		p_i	0.187555	0.312445	0.312445	0.187555						
		$\mathsf{E}[\omega\|\Omega_i]$	-1.43535	-0.415223	0.415223	1.43535						
6	0.0222709	b_i	-1.11507	-0.33895	0.33895	1.11507	∞					
		p_i	0.132411	0.234913	0.265353	0.234913	0.132411					
		$\mathsf{E}[\omega\|\Omega_i]$	-1.61805	-0.691424	0	0.691424	1.61805					
7	0.0157461	b_i	-1.28855	-0.579834	0	0.579834	1.28855	∞				
		p_i	0.0987769	0.182236	0.218987	0.218987	0.182236	0.0987769				
		$\mathsf{E}[\omega\|\Omega_i]$	-1.7608	-0.896011	-0.281889	0.281889	0.896011	1.7608				
8	0.0117218	b_i	-1.42763	-0.765185	-0.244223	0.244223	0.765185	1.42763	∞			
		p_i	0.0766989	0.145382	0.181448	0.192942	0.181448	0.145382	0.0766989			
		$\mathsf{E}[\omega\|\Omega_i]$	-1.87735	-1.05723	-0.493405	0	0.493405	1.05723	1.87735			
9	0.00906529	b_i	-1.54317	-0.914924	-0.433939	0	0.433939	0.914924	1.54317	∞		
		p_i	0.0613946	0.118721	0.152051	0.167834	0.167834	0.152051	0.118721	0.0613946		
		$\mathsf{E}[\omega\|\Omega_i]$	-1.97547	-1.18953	-0.661552	-0.213587	0.213587	0.661552	1.18953	1.97547		
10	0.00721992	b_i	-1.64166	-1.03998	-0.58826	-0.19112	0.19112	0.58826	1.03998	1.64166	∞	
		p_i	0.0503306	0.0988444	0.129004	0.146037	0.151568	0.146037	0.129004	0.0988444	0.0503306	
		$\mathsf{E}[\omega\|\Omega_i]$	-2.05996	-1.30127	-0.8004	-0.384597	0.	0.384597	0.8004	1.30127	2.05996	
11	0.00588597	b_i	-1.72725	-1.14697	-0.717801	-0.347462	0.	0.347462	0.717801	1.14697	1.72725	∞
		p_i	0.0420611	0.0836356	0.110743	0.127682	0.135878	0.135878	0.127682	0.110743	0.0836356	0.0420611
		$\mathsf{E}[\omega\|\Omega_i]$	-2.13399	-1.39768	-0.9182	-0.526575	-0.17199	0.17199	0.526575	0.9182	1.39768	2.13399

The first order loss function

Approximation error of $\widehat{\mathcal{L}}_{lb}(x, Z)$ with up to eleven segments

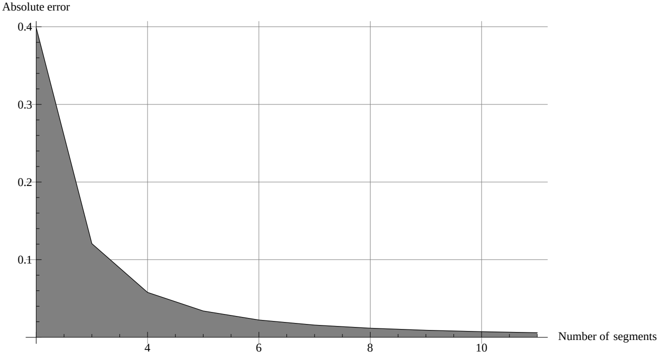

The first order loss function

Five-segment piecewise Jensen's bound for $\widehat{\mathcal{L}}(x,\zeta)$, where $\mu = 0$ and $\sigma = 1$

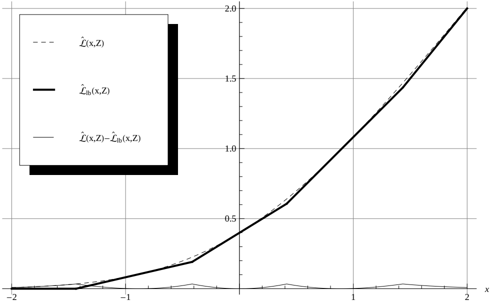

Five-segment piecewise Jensen's bound for $\widehat{\mathcal{L}}(x, Z)$, where Z is a standard normally distributed random variable. The maximum error is 0.0339052 and it is observed at $x \in \{\pm 1.43535, \pm 0.415223\}$.

The first order loss function

Five-segment piecewise Jensen's bound for $\widehat{\mathcal{L}}(x,\zeta)$, where $\mu = 20$ and $\sigma = 5$

Below we exploit the fact that the complementary first order loss function of ζ can be expressed in terms of the standard Normal cumulative distribution function as

$$\widehat{\mathcal{L}}(x,\zeta) = \sigma \int_{-\infty}^{\frac{x-\mu}{\sigma}} \Phi(t)\,dt = \sigma \widehat{\mathcal{L}}\left(\frac{x-\mu}{\sigma}, Z\right),$$

where Z is a standard Normal random variable.

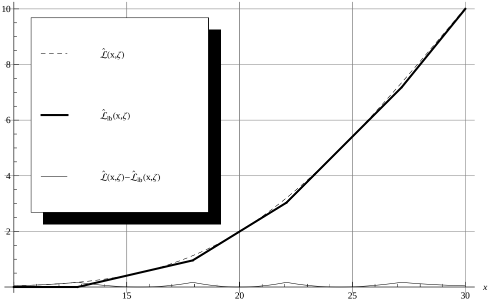

Five-segment piecewise Jensen's bound for $\widehat{\mathcal{L}}(x,\zeta)$, where ζ is a normally distributed random variable with mean $\mu = 20$ and standard deviation $\sigma = 5$. The maximum error is $\sigma 0.0339052$ and it is observed at $x \in \{\sigma(\pm 1.43535) + \mu, \sigma(\pm 0.415223) + \mu\}$.

The first order loss function

Five-segment piecewise linear upper bound for $\widehat{\mathcal{L}}(x, \zeta)$, where $\mu = 20$ and $\sigma = 5$

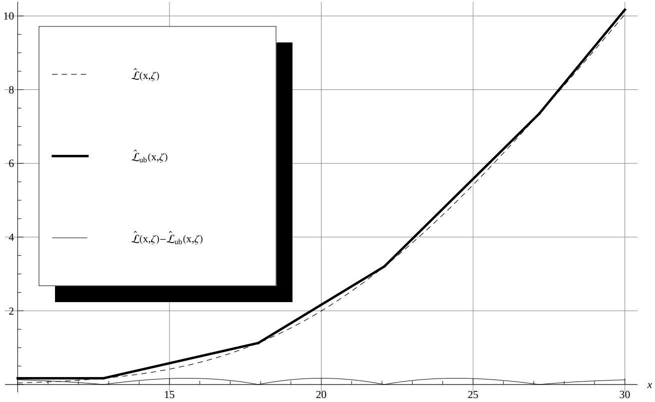

Five-segment piecewise linear upper bound for $\widehat{\mathcal{L}}(x, \zeta)$, where ζ is a normally distributed random variable with mean $\mu = 20$ and standard deviation $\sigma = 5$. The maximum error is $\sigma 0.0339052$ and it is observed at $x \in \{\pm\infty, \sigma(\pm 0.886942) + \mu, \mu\}$.

Stochastic lot-sizing

General framework

$$\min \mathsf{E}[\mathsf{TC}] = \int_{d_1} \int_{d_2} \dots \int_{d_N} \sum_{t=1}^{N} (a\delta_t + h \max(I_t, 0) + vQ_t) \times$$
$$g_1(d_1) g_2(d_2) \dots g_N(d_N)\, d(d_1) d(d_2) \dots d(d_N)$$

subject to, for $t = 1, \dots N$

$$I_t = I_0 + \sum_{i=1}^{t} (Q_i - d_i)$$

$$\delta_t = \begin{cases} 1 & \text{if } Q_t > 0, \\ 0 & \text{otherwise} \end{cases}$$

$$Q_i \geq 0,\ \ \delta_t \in \{0, 1\}$$

Stochastic lot-sizing

α service level

$$\min \mathsf{E}[\mathsf{TC}] = \int_{d_1}\int_{d_2}\ldots\int_{d_N}\sum_{t=1}^{N}(a\delta_t + h\max(I_t,0) + vQ_t)\times$$
$$g_1(d_1)g_2(d_2)\ldots g_N(d_N)\,d(d_1)d(d_2)\ldots d(d_N)$$

subject to, for $t = 1,\ldots N$

$$I_t = I_0 + \sum_{i=1}^{t}(Q_i - d_i)$$

$$\delta_t = \begin{cases} 1 & \text{if } Q_t > 0, \\ 0 & \text{otherwise} \end{cases}$$

$$\Pr\{I_t \geq 0\} \geq \alpha$$

$$Q_i \geq 0,\ \ \delta_t \in \{0,1\}$$

Stochastic lot-sizing

Penalty cost

$$\min \mathsf{E}[\mathsf{TC}] = \int_{d_1}\int_{d_2}\dots\int_{d_N}\sum_{t=1}^{N} (a\delta_t + h\max(I_t,0) + p\max(-I_t,0) + vQ_t)\times g_1(d_1)g_2(d_2)\dots g_N(d_N)\,d(d_1)d(d_2)\dots d(d_N)$$

subject to, for $t = 1, \dots N$

$$I_t = I_0 + \sum_{i=1}^{t}(Q_i - d_i)$$

$$\delta_t = \begin{cases} 1 & \text{if } Q_t > 0, \\ 0 & \text{otherwise} \end{cases}$$

$$Q_i \geq 0,\ \ \delta_t \in \{0,1\}$$

Stochastic lot-sizing

β^{cyc} service level

H. Tempelmeier. On the stochastic uncapacitated dynamic single-item lotsizing problem with service level constraints.
European Journal of Operational Research, 181(1):184–194, August 2007

$$\min \mathsf{E}[\mathsf{TC}] = \int_{d_1}\int_{d_2}\dots\int_{d_N}\sum_{t=1}^{N}(a\delta_t + h\max(I_t,0) + vQ_t)\times$$
$$g_1(d_1)g_2(d_2)\dots g_N(d_N)\,d(d_1)d(d_2)\dots d(d_N)$$

subject to, for $t = 1,\dots N$

$$I_t = I_0 + \sum_{i=1}^{t}(Q_i - d_i)$$

$$\delta_t = \begin{cases} 1 & \text{if } Q_t > 0,\\ 0 & \text{otherwise}\end{cases}$$

$$1 - \max_{i=1,\dots,m}\left[\mathsf{E}\left\{\frac{\text{Total backorders in replenishment cycle } i}{\text{Total demand in replenishment cycle } i}\right\}\right] \geq \beta^{\text{cyc}}$$

$$Q_i \geq 0,\ \ \delta_t \in \{0,1\}$$

Stochastic lot-sizing

β service level

$$\min \mathsf{E}[\mathsf{TC}] = \int_{d_1}\int_{d_2}\ldots\int_{d_N}\sum_{t=1}^{N}(a\delta_t + h\max(I_t,0) + vQ_t)\times$$
$$g_1(d_1)g_2(d_2)\ldots g_N(d_N)\,d(d_1)d(d_2)\ldots d(d_N)$$

subject to, for $t = 1,\ldots N$

$$I_t = I_0 + \sum_{i=1}^{t}(Q_i - d_i)$$

$$\delta_t = \begin{cases} 1 & \text{if } Q_t > 0, \\ 0 & \text{otherwise} \end{cases}$$

$$1 - \mathsf{E}\left\{\frac{\text{Total backorders within the planning horizon}}{\text{Total demand within the planning horizon}}\right\} \geq \beta$$

$$Q_i \geq 0,\;\; \delta_t \in \{0,1\}$$

Problem parameters

Normally distributed demand with constant coefficient of variation

$$c = \frac{\sigma_t}{\mu_t}$$

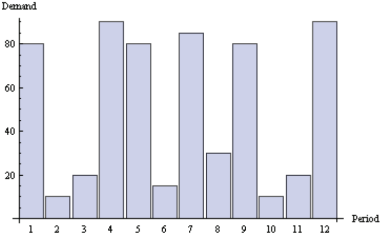

Static-dynamic uncertainty

J. H. Bookbinder and J. Y. Tan. Strategies for the probabilistic lot-sizing problem with service-level constraints.
Management Science, 34:1096–1108, 1988

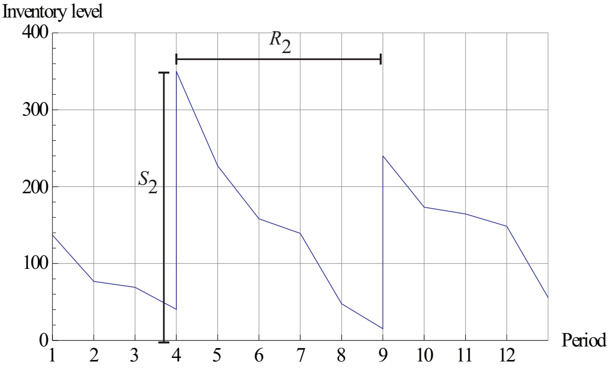

Static-dynamic uncertainty

J. H. Bookbinder and J. Y. Tan. Strategies for the probabilistic lot-sizing problem with service-level constraints.
Management Science, 34:1096–1108, 1988

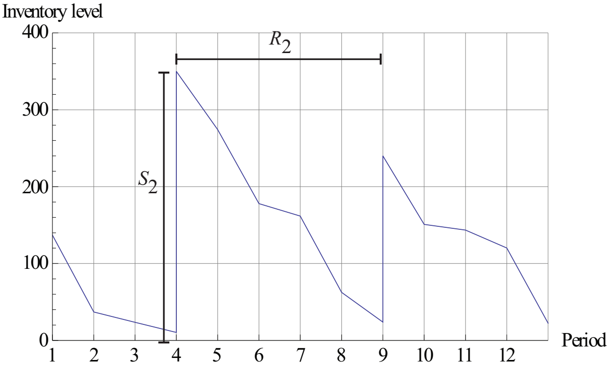

Static-dynamic uncertainty

J. H. Bookbinder and J. Y. Tan. Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science*, 34:1096–1108, 1988

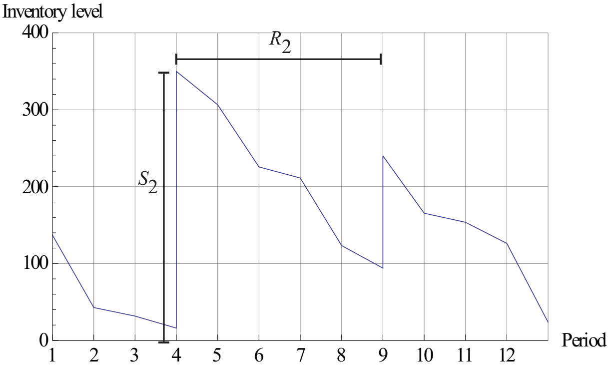

Static-dynamic uncertainty

J. H. Bookbinder and J. Y. Tan. Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science*, 34:1096–1108, 1988

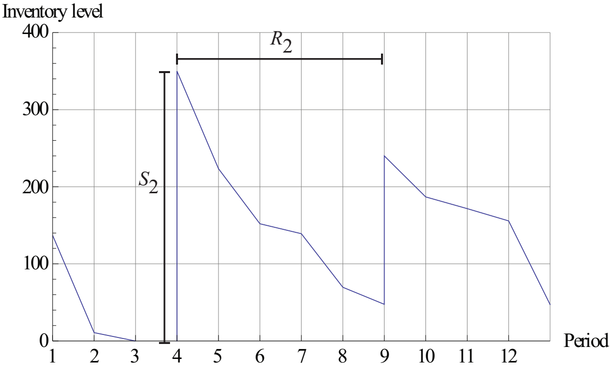

Static-dynamic uncertainty

MILP model under α service level

S. A. Tarim and Brian G. Kingsman. The stochastic dynamic production/inventory lot-sizing problem with service-level constraints.
International Journal of Production Economics, 88(1):105–119, March 2004

$$\mathrm{E[TC]} = -vI_0 + v\sum_{t=1}^{N}\tilde{d}_t + \min\sum_{t=1}^{N}(a\delta_t + h\tilde{I}_t) + v\tilde{I}_N$$

subject to, for $t = 1, \dots N$

$$\tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} \geq 0$$

$$\tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} \leq \delta_t M_t$$

$$\tilde{I}_t \geq \sum_{j=1}^{t}\left(G^{-1}_{d_{j\dots t}}(\alpha) - \sum_{k=j}^{t}\tilde{d}_k\right)P_{jt}$$

$$\sum_{j=1}^{t} P_{jt} = 1$$

$$P_{jt} \geq \delta_j - \sum_{k=j+1}^{t}\delta_k \qquad j = 1, \dots, t$$

$$P_{jt} \in \{0,1\} \qquad j = 1, \dots, t$$

$$\delta_t \in \{0,1\}$$

Static-dynamic uncertainty

Enhanced MILP model under α service level

We introduce two new sets of decision variables: $\tilde{I}_t^{lb}$ and $\tilde{I}_t^{ub}$ for $t = 1, \ldots, N$, which represent, respectively, a lower and an upper bound to the true value of $\mathsf{E}[\max(I_t, 0)]$. We introduce the following constraints in the model

$$\tilde{I}_t^{lb} \geq \tilde{I}_t \sum_{k=1}^{i} p_k - \sum_{j=1}^{t} \left(\sum_{k=1}^{i} p_k \mathsf{E}\{Z|\Omega_i\} \right) P_{jt} \sigma_{d_{j\ldots t}} \quad t = 1, \ldots, N; \; i = 1, \ldots, W$$

where $\sigma_{d_{j\ldots t}}$ denotes the standard deviation of $d_j + \ldots + d_t$ and $\tilde{I}_t^{lb} \geq 0$.

$$\tilde{I}_t^{ub} \geq \tilde{I}_t \sum_{k=1}^{i} p_k - \sum_{j=1}^{t} \left(\sum_{k=1}^{i} p_k \mathsf{E}\{Z|\Omega_i\} \right) P_{jt} \sigma_{d_{j\ldots t}} + \sum_{j=1}^{t} e^W P_{jt} \sigma_{d_{j\ldots t}} \quad \begin{array}{l} t = 1, \ldots, N, \\ i = 1, \ldots, W; \end{array}$$

where $\tilde{I}_t^{ub} \geq e^W$ and e^W denotes the maximum approximation error associated with a partition comprising W regions. Finally, the objective function becomes

$$\mathsf{E}[\mathsf{TC}] = -vI_0 + v \sum_{t=1}^{N} \tilde{d}_t + \min \sum_{t=1}^{N} (a\delta_t + h\tilde{I}_t^{lb}) + v\tilde{I}_N$$

if our aim is to compute a lower bound for the cost of an optimal plan, or

$$\mathsf{E}[\mathsf{TC}] = -vI_0 + v \sum_{t=1}^{N} \tilde{d}_t + \min \sum_{t=1}^{N} (a\delta_t + h\tilde{I}_t^{ub}) + v\tilde{I}_N$$

if our aim is to compute an upper bound for the cost of an optimal plan.

Static-dynamic uncertainty

Numerical example under α service level

We demonstrate our approach on an instance originally discussed in S. A. Tarim and Brian G. Kingsman. The stochastic dynamic production/inventory lot-sizing problem with service-level constraints. *International Journal of Production Economics*, 88(1):105–119, March 2004.

The instance comprises $N = 10$ periods in the planning horizon. Demand d_t in period t is normally distributed with mean μ_t and standard deviation σ_t.

Piecewise linear approximation (2 seg.) - E[TC]$\in [9989.07, 10314.00]$

t	1	2	3	4	5	6	7	8	9	10
δ_t	1	0	0	0	0	1	0	0	0	0
S_t	1000.46	-	-	-	-	867.35	-	-	-	-

Piecewise linear approximation (11 seg.) - E[TC]$\in [9993.66, 9998.46]$

t	1	2	3	4	5	6	7	8	9	10
δ_t	1	0	0	0	0	1	0	0	0	0
S_t	1000.46	-	-	-	-	867.35	-	-	-	-

The expected total cost estimated by Tarim and Kingsman's model is 9989.07. We simulated this policy and estimated its expected total cost with a margin of error of $\pm 0.001\%$ at 95% confidence; the resulting cost is 9993.74 ± 0.1.

Static-dynamic uncertainty

MILP model under penalty cost

S. Armagan Tarim and Brian G. Kingsman. Modelling and computing (Rn,Sn) policies for inventory systems with non-stationary stochastic demand.
European Journal of Operational Research, 174(1):581–599, October 2006

We introduce two new sets of variables $\tilde{B}_t^{lb}$ and $\tilde{B}_t^{ub}$ for $t = 1, \ldots, N$, which represent a lower and upper bound, respectively, for the true value of $\mathsf{E}[-\min(\tilde{I}_t, 0)]$ and thus allow us to compute lower and upper bounds for the expected backorders in each period.

$$\tilde{B}_t^{lb} \geq -\tilde{I}_t + \tilde{I}_t \sum_{k=1}^{i} p_k - \sum_{j=1}^{t} \left(\sum_{k=1}^{i} p_k \mathsf{E}\{Z|\Omega_i\} \right) P_{jt}\sigma_{d_{j\ldots t}} \qquad \begin{array}{l} t = 1, \ldots, N, \\ i = 1, \ldots, W; \end{array}$$

where $\tilde{B}_t^{ub} \geq -\tilde{I}_t$ and

$$\tilde{B}_t^{ub} \geq -\tilde{I}_t + \tilde{I}_t \sum_{k=1}^{i} p_k - \sum_{j=1}^{t} \left(\sum_{k=1}^{i} p_k \mathsf{E}\{Z|\Omega_i\} \right) P_{jt}\sigma_{d_{j\ldots t}} + \sum_{j=1}^{t} e^W P_{jt}\sigma_{d_{j\ldots t}} \qquad \begin{array}{l} t = 1, \ldots, N, \\ i = 1, \ldots, W; \end{array}$$

where $\tilde{B}_t^{ub} \geq -\tilde{I}_t + e^W$.

The objective function then becomes

$$\mathsf{E}[\mathsf{TC}] = -vI_0 + v \sum_{t=1}^{N} \tilde{d}_t + \min \sum_{t=1}^{N} (a\delta_t + h\tilde{I}_t^{lb} + b\tilde{B}_t^{lb}) + v\tilde{I}_N \tag{7}$$

if our aim is to compute a lower bound for the cost of an optimal plan, or

$$\mathsf{E}[\mathsf{TC}] = -vI_0 + v \sum_{t=1}^{N} \tilde{d}_t + \min \sum_{t=1}^{N} (a\delta_t + h\tilde{I}_t^{ub} + b\tilde{B}_t^{ub}) + v\tilde{I}_N \tag{8}$$

if our aim is to compute an upper bound for the cost of an optimal plan.

Static-dynamic uncertainty

Numerical example under penalty cost

We demonstrate our approach on an instance originally discussed in Charles R. Sox. Dynamic lot sizing with random demand and non-stationary costs.
Operations Research Letters, 20(4):155–164, May 1997.

The instance comprises $N = 8$ periods in the planning horizon. Demand d_t in period t is normally distributed with mean μ_t and standard deviation σ_t.

Piecewise linear approximation - E[TC]$\in [1024.70, 1034.24]$

t	1	2	3	4	5	6	7	8	9	10
δ_t	1	1	0	1	0	1	1	1		
S_t^{ub}	130.2	57.072	-	85.597	-	102.363	156.103	185.484		
S_t^{lb}	130.2	57.072	-	85.597	-	102.363	156.103	185.484		

Tarim and Kingsman - E[TC]=1031 (simulated: 1036.30)

t	1	2	3	4	5	6	7	8	9	10
δ_t	1	1	0	1	0	1	1	1		
S_t	128.5	56.9	-	84.6	-	101.9	155.4	165.6		

The expected total cost estimated by Tarim and Kingsman's model is 1031. We simulated this policy and estimated its expected total cost with a margin of error of $\pm 0.01\%$ at 95% confidence; the resulting cost is 1036.30 ± 0.1.

Policy parameters obtained via our MILP approximation converge for eleven segments; the linearisation gap is however 0.92%, reflecting the fact that the actual cost of this policy lies somewhere between 1024.70 and 1034.24.

We simulated this policy and estimated its expected total cost with a margin of error of $\pm 0.01\%$ at 95% confidence; the resulting cost is 1034.14 ± 0.1.

Static-dynamic uncertainty

MILP model under β^{cyc} service level as defined in Tempelmeier (2007)

We introduce constraints

$$\tilde{B}_t^{lb} \leq (1 - \beta^{\text{cyc}}) \sum_{j=1}^{t} P_{jt} \mu_{d_{j...t}} \quad t = 1, \ldots, N, \tag{9}$$

if our aim is to compute a lower bound for the cost of an optimal plan; or with

$$\tilde{B}_t^{ub} \leq (1 - \beta^{\text{cyc}}) \sum_{j=1}^{t} P_{jt} \mu_{d_{j...t}} \quad t = 1, \ldots, N, \tag{10}$$

if our aim is to compute an upper bound for the cost of an optimal plan. Finally, the objective function becomes

$$\text{E[TC]} = -vI_0 + v \sum_{t=1}^{N} \tilde{d}_t + \min \sum_{t=1}^{N} (a\delta_t + h\tilde{I}_t^{lb}) + v\tilde{I}_N \tag{11}$$

if our aim is to compute a lower bound for the cost of an optimal plan, or

$$\text{E[TC]} = -vI_0 + v \sum_{t=1}^{N} \tilde{d}_t + \min \sum_{t=1}^{N} (a\delta_t + h\tilde{I}_t^{ub}) + v\tilde{I}_N \tag{12}$$

if our aim is to compute an upper bound for the cost of an optimal plan.

Static-dynamic uncertainty

Numerical example under β^{cyc} service level as defined in Tempelmeier (2007)

We solved the same instance discussed for the case of an α service level, however we now enforced a β^{cyc} service level of 0.95. By using eleven segments in the linearisation, the optimality gap is very narrow, i.e. 0.23%.

Piecewise linear approximation - E[TC]$\in [8347.40, 8367.03]$

t	1	2	3	4	5	6	7	8	9	10
δ_t	1	0	0	1	0	0	0	0	0	0
S_t^{ub}	373.95	-	-	1150.85	-	-	-	-	-	-
S_t^{lb}	372.84	-	-	1149.17	-	-	-	-	-	-

Tempelmeier - E[TC]=8348 (simulated: 8347.10)

t	1	2	3	4	5	6	7	8	9	10
δ_t	1	0	0	1	0	0	0	0	0	0
S_t	373	-	-	1149	-	-	-	-	-	-

We simulated both the policies obtained and estimated their expected total cost with a margin of error of ± 0.001% at 95% confidence; the resulting costs are 8347.71 ± 0.08 and 8361.31 ± 0.08, respectively.

Static-dynamic uncertainty

MILP model under β service level

We introduce two new set of nonnegative variables $\tilde{C}_t^{lb}$ and $\tilde{C}_t^{ub}$ for $t = 0, \ldots, N$. These variables express the expected total backorders within the replenishment cycle that ends at period t, if there is one.

We set $\tilde{B}_0^{lb} = \tilde{B}_0^{ub} = \tilde{C}_0^{lb} = \tilde{C}_0^{ub} = I_0$, then we enforce

$$\tilde{C}_t^{lb} \geq \tilde{B}_t^{lb} - \delta_{t+1} \sum_{k=1}^{t} \tilde{d}_t \quad t = 0, \ldots, N-1, \tag{13}$$

$$\tilde{C}_t^{ub} \geq \tilde{B}_t^{ub} - \delta_{t+1} \sum_{k=1}^{t} \tilde{d}_t \quad t = 0, \ldots, N-1. \tag{14}$$

Finally, we must ensure that $\tilde{C}_N^{lb} = \tilde{B}_N^{lb}$ and $\tilde{C}_N^{ub} = \tilde{B}_N^{ub}$. We then use these new variables to build constraint

$$\sum_{t=1}^{N} \tilde{C}_t^{lb} \leq (1-\beta) \sum_{t=1}^{N} \tilde{d}_t \tag{15}$$

which will replace (9), if our aim is to compute a lower bound for the cost of an optimal plan; and constraint

$$\sum_{t=1}^{N} \tilde{C}_t^{ub} \leq (1-\beta) \sum_{t=1}^{N} \tilde{d}_t \tag{16}$$

which will replace (10), if our aim is to compute an upper bound for the cost of an optimal plan.

Static-dynamic uncertainty

Numerical example under β service level

We solved the same instance discussed for the case of an α service level, however we now enforced a β service level $\beta = 0.6$ and the setup costs are reduced to $a = 1000$.

Piecewise linear approximation - E[TC]$\in [2602.58, 2612.69]$

t	1	2	3	4	5	6	7	8	9	10
δ_t	0	0	0	1	0	0	0	0	0	0
S_t^{ub}	-	-	-	903.49	-	-	-	-	-	-
S_t^{lb}	-	-	-	902.43	-	-	-	-	-	-

We simulated both the policies and estimated their expected total cost with a margin of error of $\pm 0.01\%$ at 95% confidence; the resulting costs are 2603.63 ± 0.26 and 2609.11 ± 0.26, respectively.

The cost reduction with respect to the policies obtained under a β^{cyc} service level is substantial and amounts to 6.4%.

Computational experience

Instances

We generated a total of 810 instances.

10 demand patterns	
ordering cost	[500,1000,2000]
unit cost	[2,5,10]
coefficient of variation	[0.10,0.20,0.30]
penalty cost	[2,5,10]

Computational experience

Demand patterns

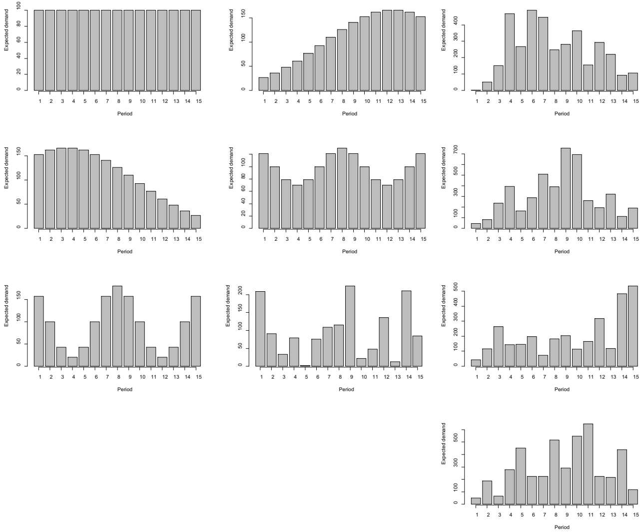

Computational experience

Optimality gap

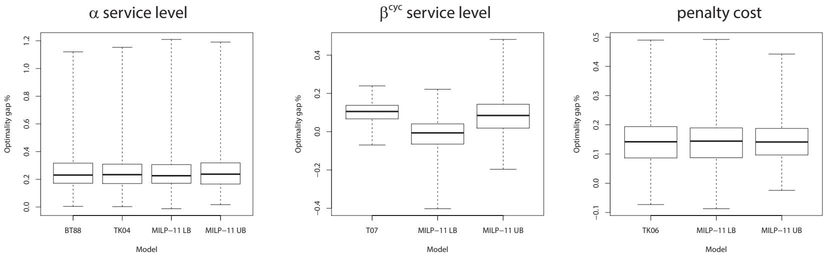

Figure: Boxplots illustrating the optimality gap for different models considered in our study.

Computational experience

Optimality gap

		α service level				β^{cyc} service level			penalty cost		
		BT88	TK04	MILP-11		T07	MILP-11		TK06	MILP-11	
				LB	UB		LB	UB		LB	UB
c	0.1	0.202	0.198	0.198	0.202	0.098	-0.002	0.033	0.118	0.107	0.119
	0.2	0.256	0.249	0.257	0.256	0.106	-0.023	0.067	0.160	0.162	0.153
	0.3	0.313	0.314	0.300	0.311	0.098	-0.024	0.165	0.160	0.169	0.166
a	500	0.351	0.346	0.341	0.345	0.099	-0.064	0.088	0.147	0.147	0.156
	1000	0.240	0.231	0.233	0.244	0.103	-0.007	0.080	0.155	0.155	0.148
	2000	0.180	0.184	0.180	0.180	0.100	0.022	0.097	0.137	0.136	0.134
b	2								0.133	0.131	0.121
	5								0.149	0.148	0.155
	10								0.157	0.159	0.162
service	0.8	0.222	0.230	0.226	0.233	0.110	-0.008	0.066			
	0.9	0.258	0.242	0.245	0.240	0.097	-0.005	0.092			
	0.95	0.291	0.288	0.284	0.297	0.095	-0.036	0.107			

Table: Average optimality gaps (%) of methods for different pivoting parameters

Computational experience

Linearisation gap - α service level

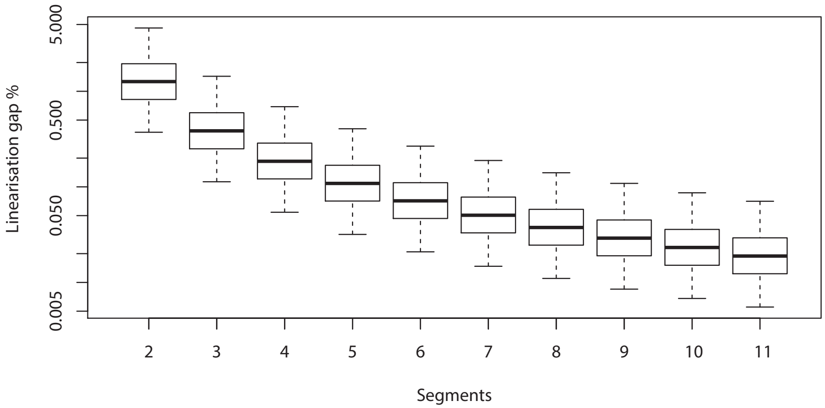

Computational experience

Linearisation gap - Penalty cost

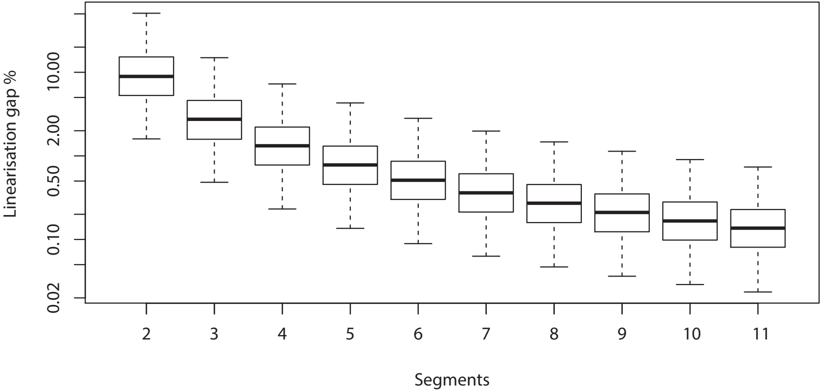

Computational experience

Linearisation gap - Cycle β service level

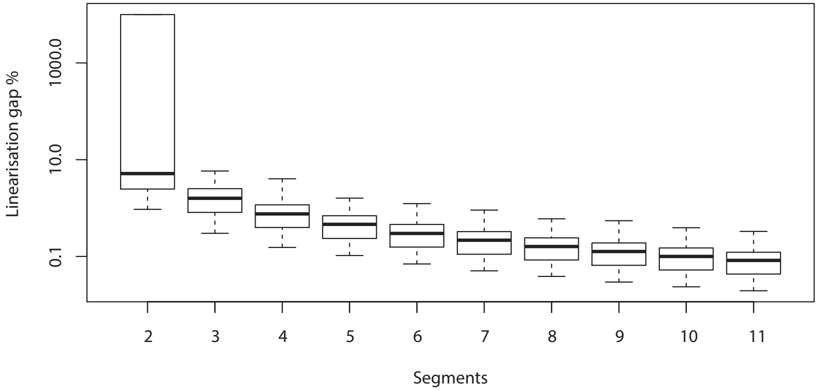

Computational experience

Linearisation gap - β service level

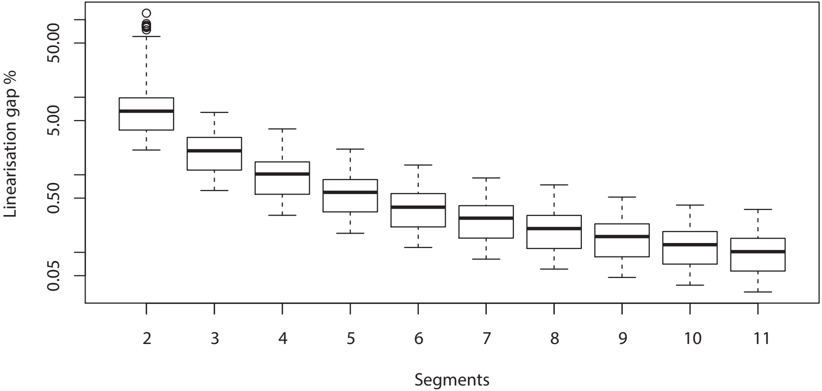

Conclusions

Related literature (sketch)

J. H. Bookbinder and J. Y. Tan. Strategies for the probabilistic lot-sizing problem with service-level constraints.
Management Science, 34:1096–1108, 1988

S. A. Tarim and Brian G. Kingsman. The stochastic dynamic production/inventory lot-sizing problem with service-level constraints.
International Journal of Production Economics, 88(1):105–119, March 2004

S. Armagan Tarim and Brian G. Kingsman. Modelling and computing (Rn,Sn) policies for inventory systems with non-stationary stochastic demand.
European Journal of Operational Research, 174(1):581–599, October 2006

H. Tempelmeier. On the stochastic uncapacitated dynamic single-item lotsizing problem with service level constraints.
European Journal of Operational Research, 181(1):184–194, August 2007

Conclusions

Questions