On the order-up-to policy with intermittent integer demand and coherent forecasts

Bahman Rostami-Tabar [1] Stephen M. Disney [2]

[1]Logistics Systems Dynamics Group, Cardiff Business School, Cardiff University, UK

[2]Center for Simulation, Analytics, and Modelling, University of Exeter Business School, UK

22nd International Working Seminar on Production Economics, Innsbruck. 21st – 25th February 2022.

Background: Most Bullwhip research assumes *real valued* demand, inventory, and orders exist

- For some products sold by volume or weight (for example, powders, granules, or liquids) real valued demand may be appropriate.
- For other products, only integer valued demand makes sense.
  - For example, you can't buy half a bicycle or half a laptop.
  - In these situations, demand, orders, inventory must be integers.
- For some situations, with high volume demand, replenishment calculations can ignore integer effects as rounding becomes negligible.
- However, low volume demand settings may be more susceptible to integer effects.

Motivation: Many low volume products have intermittent and integer demand

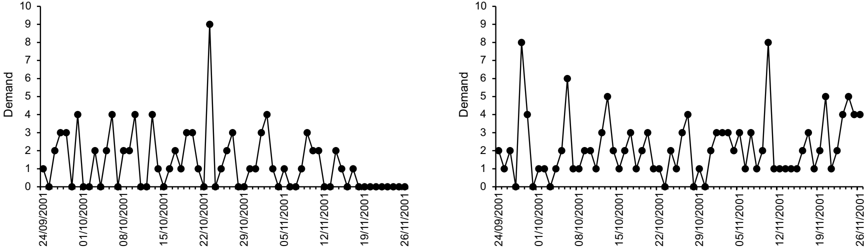

Figure 1. Low volume, occasionally zero, integer demand for two products from a single store of a UK grocery retailer

Research Question: How can we model and forecast low volume integer demand to understand how the order-up-to replenishment policy behaves in this situation?

Overview

- Demand is assumed to be an INteger Auto-Regressive, INAR(1), process.
- The order-up-to policy is used to make replenishment decisions.
- We consider two different forecasting mechanisms.
  - *Conditional mean* forecasts $\bar{d}$, which result in minimum mean squared error forecasts
  - *Conditional median* forecasts $\tilde{d}$, which minimize the expected absolute forecast error
- We measure the
  - *Bullwhip* effect, $\frac{\mathbb{V}[q_t]}{\mathbb{V}[d_t]}$,
  - Net Stock Amplification (*NSAmp*), $\frac{\mathbb{V}[i_t]}{\mathbb{V}[d_t]}$,

  and conduct an *economic analysis* of capacity and inventory costs.

First order integer auto-regressive (INAR(1)) demand

$$d_t = \phi \circ d_{t-1} + z_t, \tag{1}$$

- d_t is the demand in period t.
- $0 \leq \phi \leq 1$ is the auto-regressive parameter.
- z_t is a sequence of i.i.d. non-negative integer-valued Poisson distributed random variables, with mean λ and finite variance λ (Silva et al., 2009).
- The atomic expression $\phi \circ d_{t-1}$ is the binomial thinning operation,

$$\phi \circ d_{t-1} = \sum_{i=1}^{d_{t-1}} X_i. \tag{2}$$

Here, X_i is a sequence of i.i.d. Bernoulli indicators with parameter ϕ (i.e. with $\mathbb{P}(X_i = 1) = \phi$ for $i = \{1, 2, ..., d_{t-1}\}$).

The number of guests in an INAR hotel

$$d_t = \phi \circ d_{t-1} + z_t,$$

A natural interpretation of the INAR(1) demand is:

- d_t is the total number of guests in a hotel at time t,
- z_t is the number of new guests that arrived today,
- and $\phi \circ d_{t-1}$ is the number of guests that remained in the hotel from the day before (Ristić and Nastić, 2012).

Example INAR(1) demand processes

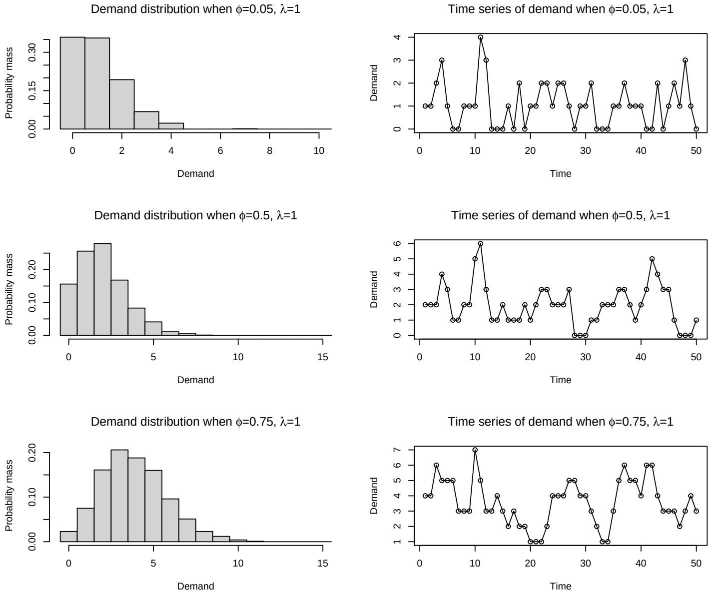

INAR algebra

Silva and Oliveira (2004) provide a number of relations and properties of the INAR(1) model. Notably, the relations

$$\mathbb{E}[\phi \circ d_t] = \phi \mathbb{E}[d_t], \tag{3}$$

$$\mathbb{E}[\phi \circ d_t]^2 = \phi^2 \mathbb{E}[d_t^2] + \phi(1-\phi)\mathbb{E}[d_t], \text{ and} \tag{4}$$

$$\mathbb{V}[\phi \circ d_t] = \mathbb{E}[\phi \circ d_t]^2 - (\mathbb{E}[\phi \circ d_t])^2 = \phi \mathbb{V}[d_t], \tag{5}$$

are useful. Here, $\mathbb{E}[\cdot]$ is the expectation operator and $\mathbb{V}[\cdot]$ is the variance operator.

Lemma 1: The INAR(1) demand process has mean, μ_d, and an auto-covariance with lag j, γ_j, of

$$\mu_d = \frac{\lambda}{1-\phi} \quad \text{and} \quad \gamma_j = \begin{cases} \frac{\lambda}{1-\phi}, & j = 0, \\ \phi^j \gamma_0, & j \geq 1. \end{cases} \tag{6}$$

The order-up-to replenishment policy

A retailer uses the OUT policy to order q_t items from the manufacturer;

$$q_t = s_t - s_{t-1} + d_t. \tag{7}$$

The order-up-to level, s_t, in time period t is determined by

$$s_t = \hat{d}_{t,L} + SS, \tag{8}$$

where

- $\hat{d}_{t,L} = \sum_{i=1}^{L} \hat{d}_{t+i}$ is the forecast of the demand over the lead-time L
- SS is a safety stock used to achieve a target inventory availability.
- As the inventory distribution is stationary, SS is a constant

The inventory balance equation is given by

$$i_t = i_{t-1} + q_{t-L} - d_t \tag{9}$$

where

- i_t is the inventory level at time t,
- q_{t-L} are the replenishment orders placed in period $t - L$,
- $L \geq 1$ is the integer lead-time that includes a sequence of event delay.

Determining the inventory variance

Vassian (1955) shows $\mathbb{V}[i_t]$ is given by the variance of forecast error over lead time:

$$\mathbb{V}[i_t] = \mathbb{V}\left[d_{t,L} - \hat{d}_{t,L}|d_t\right] = \mathbb{V}[d_{t,L}] + \mathbb{V}[\hat{d}_{t,L}|d_t] - 2\text{cov}[d_{t,L}, \hat{d}_{t,L}|d_t]. \tag{10}$$

We need three components:

- the variance of the demand over the lead time $\mathbb{V}[d_{t,L}]$,
- the variance of the forecast of the demand over the lead time $\mathbb{V}[\hat{d}_{t,L}]$,
- and the co-variance between the demand over the lead-time and the forecast of the demand over the lead time $\text{cov}[d_{t,L}, \hat{d}_{t,L}|d_t]$.

The conditional mean forecast of INAR(1) demand

Lemma 2: Under INAR(1) demand, over the lead time is

$$d_{t,L} = \frac{\phi(1-\phi^L)}{1-\phi} d_t + \frac{L\lambda(1+L)}{2}. \tag{11}$$

Proof. Start by deriving an exact expression of demand k periods ahead, d_{t+k}:

$$d_{t+k} = \phi \circ d_{t+k-1} + z_{t+k} \text{ and} \tag{12}$$

$$d_{t+k-1} = \phi \circ d_{t+k-2} + z_{t+k-1}. \tag{13}$$

Substituting (13) into (12) recursively (for $d_{t+k-2}, d_{t+k-3}, \dots$), and collecting ϕ terms (3) yields,

$$d_{t+k} = \phi^k \circ d_t + z_{t+1} + z_{t+2} + \dots + z_{t+k}. \tag{14}$$

Replacing future values of z_{t+i} in (14) with their expectation, $\mathbb{E}[z_{t+i}] = \lambda$, yields,

$$t + k = \mathbb{E}[\phi^k \circ d_t + z_{t+1} + z_{t+2} + \dots + z_{t+k} | d_t] = \phi^k d_t + k\lambda. \tag{15}$$

Summing $\sum_{k=1}^{L} \left(\phi^k d_t + k\lambda\right)$ provides (11). □

The three components of the inventory variance

Lemma 3: The variance of demand over lead time L

$$\mathbb{V}\left[d_{t,L}\right] = L\gamma_0 + 2\sum_{j=1}^{L-1}\sum_{i=1}^{j}\gamma_0\phi^i = \frac{\lambda\left(\left(\phi^2-1\right)L - 2\phi\left(\phi^L-1\right)\right)}{(\phi-1)^3} \tag{16}$$

Lemma 4: Variance of the forecast over the lead time

$$\mathbb{V}\left[\hat{d}_{t,L}|d_t\right] = \mathbb{V}\left[d_t\frac{\phi(1-\phi^L)}{1-\phi} + \frac{L\lambda(1+L)}{2}\right] = \frac{\lambda}{1-\phi}\left(\frac{\phi(1-\phi^L)}{1-\phi}\right)^2 \tag{17}$$

Lemma 5: Co-variance of lead time demand and its forecast over L

$$\text{cov}\left[d_{t,L},\hat{d}_{t,L}|d_t\right] = \text{cov}\left[d_{t+1}+d_{t+2}+\ldots+d_{t+L},\hat{d}_{t,L}|d_t\right] = \gamma_0\left(\frac{\phi\left(1-\phi^L\right)}{1-\phi}\right)^2 \tag{18}$$

Variance of the inventory levels: Conditional mean

Using Lemmas 3-5 in (10) yields:

Proposition 1: The inventory variance

$$\mathbb{V}[i_t] = \frac{\lambda}{1-\phi}\left(L + 2\phi\left(\frac{\phi^L + L(1-\phi) - 1}{(\phi-1)^2}\right) - \left(\frac{\phi(1-\phi^L)}{1-\phi}\right)^2\right) \quad (19)$$

Remark. Notice the *NSAmp* measure for the integer valued INAR(1) demand with conditional mean forecasting is the same as the *NSAmp* ratio the real valued AR(1) demand, Disney and Lambrecht (2008):

$$NSAmp = \frac{\mathbb{V}[i_t]}{\mathbb{V}[d_t]} = \left(L + 2\phi\left(\frac{\phi^L + L(1-\phi) - 1}{(\phi-1)^2}\right) - \left(\frac{\phi(1-\phi^L)}{1-\phi}\right)^2\right)$$

Variance of the replenishment orders: Condition mean

Using $q_t = \hat{d}_{t,L} - \hat{d}_{t-1,L} + d_t$ and the INAR algebra we obtain,

Proposition 2: Variance of the orders

$$\mathbb{V}\left[q_t\right] = \frac{\lambda}{1-\phi}\left(1 + 2\phi\left(1-\phi^L\right)\left(1 + \frac{\phi\left(1-\phi^L\right)}{1-\phi}\right)\right) \tag{20}$$

Remark. Notice the *Bullwhip* measure for the integer valued INAR(1) demand with conditional mean forecasting is the same as the *Bullwhip* ratio the real valued AR(1) demand, Disney and Lambrecht (2008):

$$Bullwhip = \frac{\mathbb{V}[q_t]}{\mathbb{V}[d_t]} = \left(1 + 2\phi\left(1-\phi^L\right)\left(1 + \frac{\phi\left(1-\phi^L\right)}{1-\phi}\right)\right)$$

Consequences forecasting INAR(1) demand with the conditional mean

- Under i.i.d. integer demand, the demand forecasts based on the conditional mean are constant over time,
  - $\bar{d}_{t,L} = \bar{d}_{t-1,L} = L\mu_d = L\lambda$,
  - orders equal demand, $q_t = d_t$
  - $\mathbb{V}[q_t] = \lambda$ and $\mathbb{V}[i_t] = L\lambda$
- The demand, orders, and inventory levels are all integer processes under i.i.d. demand.
- However, for correlated INAR(1) demand, the forecast is dynamic, (as $\bar{d}_{t,L} \neq \bar{d}_{t-1,L}$), and is also real valued ($\bar{d}_{t,L} \in \mathbb{R}$).
- This means real valued orders and inventory levels are present under correlated demand.
- These non-integer orders and inventory are incoherent with the integer demand assumption.

Forecasting INAR(1) demand over the lead-time with the conditional median

Let $\tilde{d}_{t+k|t}$ be an integer forecast of the demand k periods ahead, conditional upon d_t. The median k periods ahead forecast, $\tilde{d}_{t+k} = X$, where X is the smallest X such that

$$\sum_{x=0}^{X} \mathbb{P}[\tilde{d}_{t+k} = x | d_t] > 1/2. \tag{21}$$

Proposition 3: The pmf of $\tilde{d}_{t+k} = x$, given d_t is

$$\mathbb{P}[\tilde{d}_{t+k} = x | d_t] = \binom{d_t}{M_k} e^{\frac{2\lambda\left(\phi^k - 1\right)}{1-\phi}} \left(\phi^k\right)^{M_k} \Gamma(d_t + 1 - M_k) \left(1 - \phi^k\right)^{d_t - M_k} \left(\frac{\lambda\left(\phi^k - 1\right)}{\phi - 1}\right)^{x - M_k}$$

$$\times\ {}_2\tilde{F}_2\left[1, -M_k; d_t + 1 - M_k, x + 1 - M_k; \frac{\phi^{-k}\left(\phi^k - 1\right)^2 \lambda}{\phi - 1}\right], \tag{22}$$

where, $x = 0, 1, \ldots$ is a non-negative integer, $\Gamma(\cdot)$ is the Gamma function, $M_k = (\tilde{d}_{t+k} \wedge d_t)$ is the minimum of $\tilde{d}_{t+k}$ and d_t, and ${}_2\tilde{F}_2[a_1, a_2; b_1, b_2; z]$ is the regularized generalized hypergeometric function, Mathworld (2021).

Exploring the consequences of conditional median forecasts.

- Conditional median forecasts are integers and produce coherent integer orders and inventory levels.
- No further analytical work can be done (by us at least) with the conditional median forecasting technique.
- However (21) and (22) can be easily implemented and studied numerically in software such as Excel, Mathematica, and R.
- A simulation in R was ran on a HAWK High-Performance Computing Cluster at the super-computing facilities in Cardiff University.

NSAmp under INAR(1) demand with conditional mean and conditional median forecasting.

- The conditional mean (line) is exact. Each of conditional mean dots are numerical results from the simulation of 1M time periods.

- The conditional mean *NSAmp* is a lower bound of the conditional median *NSAmp*.

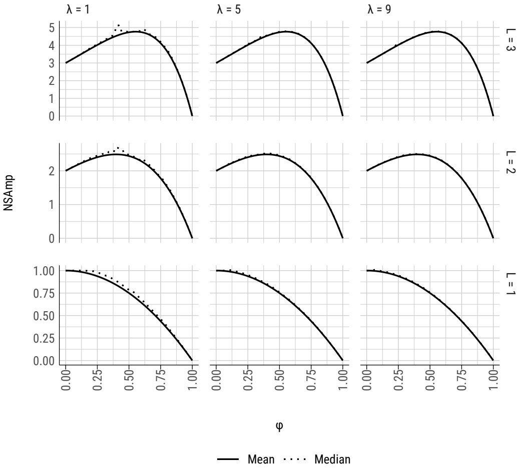

Bullwhip under INAR(1) demand with conditional mean and conditional median forecasting.

- When ϕ is near zero, the *Bullwhip* curves remain unchanged from the i.i.d. case.
- In the region $\phi \approx 0.1$ to $\phi \approx 0.8$ there is a period of seemingly erratic *Bullwhip*; above or below the conditional mean *Bullwhip*;
- Conditional median *Bullwhip* may also be above the AR(1) bullwhip upper bound (dashed line), Luong (2007).
- When ϕ is close to unity, above $\phi \approx 0.8$, the conditional median *Bullwhip* curves closely resembles the conditional mean *Bullwhip* curve.
- We do not know if conditional median forecasts ever produces *Bullwhip* < 1.

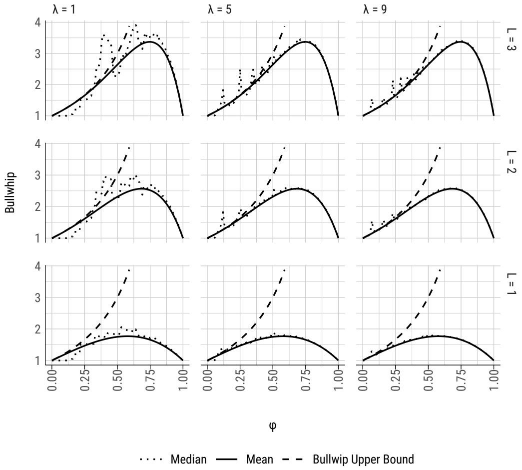

Inventory performance of the OUT policy under i.i.d. integer demands

When $\phi = 0$, we have i.i.d. Poisson distributed demands and constant demand forecasts. The inventory pmf is a simple reflection and translation of the pmf of the sum of demand over the lead-time,

$$\mathbb{P}[i_t = x] = \frac{(L\lambda)^{L\lambda+\bar{i}-x} e^{-L\lambda}}{(L\lambda + \bar{i} - x)!}. \tag{23}$$

The mean and variance of the inventory levels can be determined directly from the mean and variance of the sum of L Poisson distributed random variables:

$$\mathbb{E}[i_t] = \sum_{x=-\infty}^{L\lambda+\bar{i}} \frac{(L\lambda)^{L\lambda+\bar{i}-x} e^{-L\lambda}}{(L\lambda + \bar{i} - x)!} = \bar{i}, \tag{24}$$

and

$$\mathbb{V}[i_t] = \sum_{x=-\infty}^{L\lambda+\bar{i}} \frac{(L\lambda)^{L\lambda+\bar{i}-x} e^{-L\lambda}}{(L\lambda + \bar{i} - x)!} (x - \mathbb{E}[i_t])^2 = L\lambda. \tag{25}$$

Expected per period inventory holding and backlog costs

The expected per period inventory holding and backlog costs

$$C_{i,t} = h[i_t]^+ + b[-i_t]^+, \tag{26}$$

Here h is the per unit, per period inventory holding cost and b is the per unit, per period inventory backlog cost.

The expected per period inventory holding and backlog costs (see (26)) can be then obtained from (23) as follows:

$$\begin{aligned}\mathbb{E}[C_t^i] &= \sum_{x=-\infty}^{0} b(-x)\mathbb{P}[i_t = x] + \sum_{x=1}^{L\lambda+\bar{i}-1} hx\mathbb{P}[i_t = x] \\ &= \frac{(b+h)e^{-L\lambda}\left((L\lambda)^{\bar{i}+L\lambda+1} + \bar{i}e^{L\lambda}\Gamma[\bar{i}+L\lambda+1, L\lambda]\right)}{\Gamma[\bar{i}+L\lambda+1]} - b\bar{i}\end{aligned} \tag{27}$$

It is easy to conduct a line search on the integers for the optimal $i^\star$, the $\bar{i}$ that minimises (27).

OUT orders under i.i.d. integer demands

When $\phi = 0$, we have i.i.d. Poisson distributed demands, constant demand forecasts will be optimal and the order pmf equals the demand pmf,

$$\mathbb{P}[d_t = x] = \mathbb{P}[q_t = x] = \frac{\lambda^x e^{-\lambda}}{x!}, \tag{28}$$

with mean and variance of λ.

Note *Bullwhip* $= 1$ in this setting, under both the conditional mean and the conditional median forecasting.

Capacity costs under i.i.d. integer demands

The per period production cost C_t^q, with a nominal hours unit cost of u within a nominal capacity of K and flexible per unit overtime cost of um, where m is the overtime multiplier, (Boute et al. (2022)):

$$C_t^q = uK + um[q_t - K]^+. \tag{29}$$

The expected capacity costs can be obtained from the order pmf, (28):

$$\mathbb{E}[C_t^q] = u\left(\frac{e^{-\lambda}\lambda m\left(\lambda^k - e^{\lambda}\Gamma[k+1,\lambda]\right)}{\Gamma[k+1]} + \frac{m\Gamma[k+1,\lambda]}{\Gamma[k]} + K + m(\lambda - K)\right) \tag{30}$$

Eq. (30) can also be minimised efficiently via a line search on the integers for $K^\star$.

Optimal OUT policy settings in i.i.d. Poisson distributed demand

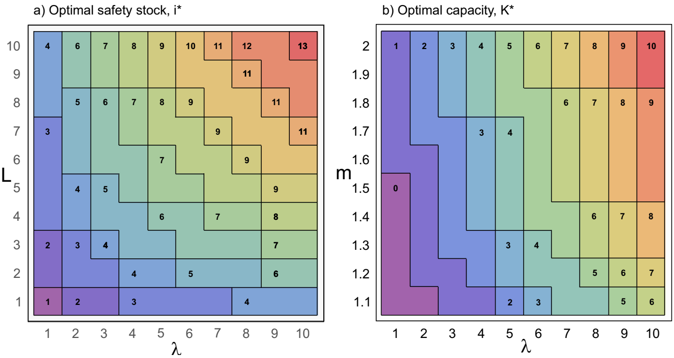

- **Panel a)** The safety stock $i^{\star}$, required to minimise the inventory costs when $h = 1, b = 9$ for different demand λ and lead times L.
- **Panel b)** The optimal capacity $K^{\star}$, required to minimise the capacity costs costs when $u = 4$ for different demand λ and over-time multiplier m.

Example time series of orders

Key:

- **Black:** Demand
- **Blue:** Orders with conditional mean forecasts
- **Red:** Orders with conditional median forecasts

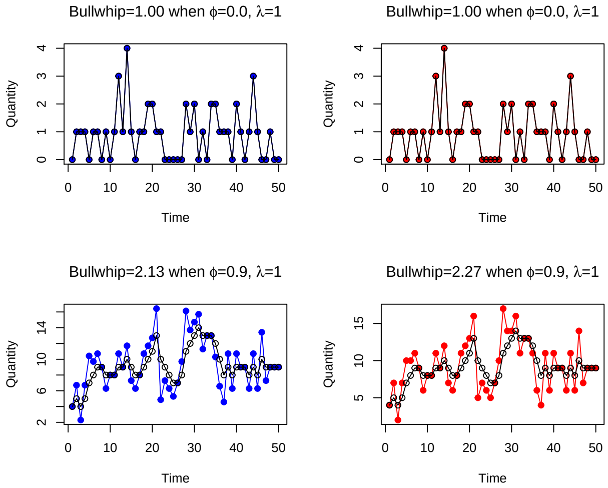

Concluding remarks:Conditional mean insights

- Variance ratios for conditional mean forecasts under integer demand and found to be the same as those for real-valued demand.
- This should be expected as the results for the real demand variance are *distribution free*.
- The Poisson distribution parameter λ had no impact on the variance ratios under INAR(1) demand.
- The real order and inventory levels produced are incoherent with the integer demand.
- *Bullwhip* always exists regardless of the auto-regressive parameter and the lead time.
- There exists a *lower bound*, *Bullwhip* > 1, and an *upper bound* which is a function of the auto-regressive parameter, ϕ. For a given value of ϕ, the upper bound represents the maximum value of the bullwhip effect regardless of the lead time L. The upper bound is tight when ϕ is small.

Concluding remarks: Conditional median insights

- *Bullwhip* is somewhat more erratic than *NSAmp* and can deviate significantly (positively and negatively) from both the conditional mean bullwhip and the upper bound.
- The *Bullwhip* and *NSAmp* expressions can be used with confidence in high volume settings.
- Low values of ϕ and λ lead to low volume intermittent demands where the integer effects become more significant.
- *Bullwhip seems* to always exists for INAR(1) demand.
- The inventory variance (and the *NSAmp* measure) is a lower bound for the inventory variance.
- When the demand is i.i.d., a constant forecast is produced by both forecasting methods that meant $q_t = d_t$.
- For this case we conducted an economic study that included a search for the target safety stock to minimise the expected inventory holding and backlog costs and the target capacity level to minimise the regular labour and overtime costs.

Thank you for listening

On the order-up-to policy with intermittent integer demand and coherent forecasts

Bahman Rostami-Tabar and Stephen Disney

Bibliography

Boute, R., S.M. Disney, J. Gijsbrechts, J. A. Van Mieghem. 2022. Dual sourcing and smoothing under non-stationary demand time series: Re-shoring with SpeedFactories. *Forthcoming in Management Science* **doi: 10.1287/mnsc.2020.3951**.

Disney, S.M., M.R. Lambrecht. 2008. On replenishment rules, forecasting, and the bullwhip effect in supply chains. *Foundations and Trends® in Technology, Information and Operations Management* **2**(1) 1–80.

Luong, H.T. 2007. Measure of bullwhip effect in supply chains with auto-regressive demand process. *European Journal of Operational Research* **180** 1086–1097.

Mathworld. 2021. Regularized hypergeometric function. `https://mathworld.wolfram.com/RegularizedHypergeometricFunction.html`.

Ristić, M.M., A.S. Nastić. 2012. A mixed INAR(p) model. *Journal of Time Series Analysis* **33**(6) 903–915.

Silva, M.E.D., V.L. Oliveira. 2004. Difference equations for the higher-order moments and cumulants of the INAR(1) model. *Journal of Time Series Analysis* **25**(3) 317–333.

Silva, N., I. Pereira, E. M. Silva. 2009. Forecasting in INAR(1) model. *Revstat – Statistical Journal* **7**(1) 119–134.