

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository:<https://orca.cardiff.ac.uk/id/eprint/98101/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Rekik, Yacine, Glock, Christoph H. and Syntetos, Aris A. 2017. Enriching demand forecasts with managerial information to improve inventory replenishment decisions: exploiting judgment and fostering learning. *European Journal of Operational Research* 261 (1) , pp. 182-194. 10.1016/j.ejor.2017.02.001

Publishers page: <http://dx.doi.org/10.1016/j.ejor.2017.02.001>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Enriching demand forecasts with managerial information to improve inventory replenishment decisions: exploiting judgment and fostering learning

Paper accepted for publication in the European Journal of Operational Research

Yacine Rezik^a, Christoph H. Glock^b, Aris A. Syntetos^{c,*}

^a*EMLYON Business School, DISP lab, France*

^b*Technische Universität Darmstadt, Germany*

^c*Panalpina Centre for Manufacturing and Logistics Research, Cardiff Business School, UK*

Abstract

This paper is concerned with analyzing and modelling the effects of judgmental adjustments to replenishment order quantities. Judgmentally adjusting replenishment quantities suggested by specialized (statistical) software packages is the norm in industry. Yet, to date, no studies have attempted to either analytically model this situation or practically characterize its implications in terms of ‘learning’. We consider a newsvendor setting where information available to managers is reflected in the form of a signal that may or may not be correct, and which may or may not be trusted. We show the analytical equivalence of adjusting an order quantity and deriving an entirely new one in light of a necessary update of the estimated demand distribution. Further, we assess the system’s behavior through a simulation experiment on theoretically generated data and we study how to foster learning to efficiently utilize managerial information. Judgmental adjustments are found to be beneficial even when the probability of a correct signal is not known. More generally, some interesting insights emerge into the practice of judgmentally adjusting order quantities.

Keywords: Inventory; Judgement; Judgmental adjustments; Newsvendor model; Learning

1. INTRODUCTION AND MOTIVATION

In most contemporary organizations, the size and complexity of the inventory management task at the individual Stock Keeping Unit (SKU) level necessitates

*Corresponding author

Email addresses: SyntetosA@cardiff.ac.uk (Aris A. Syntetos), +44 (0) 2920 876572 (Aris A. Syntetos)

the employment of statistical theory as far as both forecasting and stock control are concerned. The advantage of doing so is that the resulting methods can be fully automated. This, in conjunction with the recent IT processing developments, means that demand forecasts and replenishment decisions can be made for hundreds or thousands of SKUs in as little time as just a few seconds and as regularly as daily or even every half day, in the case of large stores or supermarkets. In principle, the automation of the inventory management task frees up managerial time that may be used at higher decision-making levels and/or for personally superimposing judgement on the quantities suggested by the automated system for the most important SKUs. Those would be items associated with a current promotion or any other special event (Goodwin & Fildes (1999)). However, research studies have demonstrated that managers personally intervene in the process far more than what one might expect. As an example, consider the following situation. In a study conducted by Syntetos et al. (2009) and Syntetos et al. (2010) for a branch of a major pharmaceutical company, sales forecasts were updated monthly for about 270 SKUs. The forecasts served both the inventory management task but also higher-level considerations and as such, extrapolation/estimation covered the subsequent 36 months (36 steps ahead). That is to say, at the end of every month, forecasts were produced for each single SKU for the subsequent 36 months. It is surprising that managers had adjusted about 65% of the forecasts examined in that research (Fildes et al. (2009)). Franses & Legerstee (2011b) looked at the linkage between judgmental adjustments of statistical forecasts and the forecast horizon. They also found, through an empirical study in the pharmaceutical industry, that all horizons (short term and long term) were associated with managerial interventions in forecasting.

The practice of judgmentally adjusting statistical sales/demand forecasts has received much attention in the academic literature in recent years. The increasing number of relevant studies reflects the importance of this area in terms of both necessary theory development and practical implications. Sales forecasts constitute an input into a stock control model that suggests when and how much to order. However, incorporating managerial judgment directly into such stock replenishment decisions has become the norm in industry (Kolassa et al. (2008)). Practitioners may directly adjust re-order and Order-Up-To (OUT) levels and/or order quantities (without this implying that forecasts are not also subject to such adjustments) for the purpose of achieving better trade-offs between achieved service levels and inventory costs.

Despite the fact that replenishment decisions may often be subject to judgmental intervention, the effect of judgmentally adjusting replenishment quantities has not attracted much attention in the academic literature, either in modelling or in empirical terms. Concerning the latter, we are aware of only one study that attempts to throw light on this area (Syntetos et al. (2016b)). The researchers considered the effects of superimposing judgement into statistically derived OUT levels and evaluated the implications of doing so through an empirical dataset from the electronics industry. Concerning the former, the only attempt we are aware of to model such a situation relied upon a System Dynamics (SD)

methodology (Syntetos et al. (2011)). The researchers considered, by means of SD simulation, the joint effects of adjusting forecasts and inventory replenishment decisions. They acknowledged the fact that analytical representation of the problem in hand, i.e. the joint consideration of judgmental adjustments at both the forecasting and inventory control process, is virtually impossible. Although this is true, analytical modelling of the effects of judgement replenishment decisions is feasible and very much needed also for the purpose of deriving insights into relevant situations. Conducting such modelling is our main objective.

This paper contributes to closing the research gap identified above and develops a single selling season model that provides insights into the practice of judgmentally adjusting order quantities. The scenario considered here can be described as follows. First, a forecast is produced based on past sales data and an initial order quantity is specified. Subsequently, a signal is observed that contains some important information not reflected in the historical sales data. Contrary to other modelling attempts presented in the literature (e.g. single season models with information updating), the signal may or may not be correct. Finally, the order quantity is adjusted based on the observed signal, it is released and the order is received prior to the beginning of the season. Due to long lead times, no further opportunities for ordering are available. The development of our model (and the specification of optimality conditions for the order quantity (Q) and the adjustment (A)) is followed by a numerical analysis that allows us to obtain some key insights into the process of adjusting the order quantity and an appreciation of how learning to efficiently utilize the signal can be fostered. An Excel file has been made available as an electronic companion to enable interested readers to experiment with the learning processes discussed in this paper. Instructions on how to use the Excel file are provided in an on-line supplement. Such material would enable other researchers to reproduce our results (Boylan (2016)) and ‘play’ with different control parameter combinations, but also extend our results should they wish to take this research further.

The remainder of the paper is structured as follows. The next section presents the research background of this work. We structure the section around two main issues: i) modelling research on (forecasting) judgmental adjustments and ii) research on single selling season models. The inventory model developed for our research purposes, along with the notation and assumptions used, is presented in Section 3. This is followed by a detailed numerical analysis in Section 4. Section 5 discusses some alternative techniques for learning to efficiently utilize the observed signal. Finally, the conclusions of our work along with its implications for Operational Research theory and practice and some natural avenues for further research are presented in Section 6.

2. RESEARCH BACKGROUND

This section discusses the thematic and methodological background of our work. We first focus on studies that relate to the practice of judgmentally adjusting sales forecasts and make the case for the extension of the current state of knowledge into the area of inventory control. We then move to the methodological

motivation behind our research by discussing the literature on single selling season modelling exercises that consider information updating and/or behavioral aspects.

2.1. Modelling of judgement in an inventory forecasting context

This section focuses on modelling attempts in the area of judgmentally adjusting statistical forecasts. As discussed in the previous section, there is a plethora of empirical studies in this area and we refer interested readers to Syntetos et al. (2016a) for a review of recent developments. Here we focus on the few studies that have attempted to model such practices statistically.

Franses & Legerstee (2009) examined whether we can predict expert adjustments using the forecasters' own past interventions and past model-based forecast errors at various lags. They did so by means of constructing an auxiliary regression model based on which they calculated the impact of one's own past adjustments (persistence) and their relevant size, taking into account the effects of past variance. They found that adjustments occurred a staggering 90% of the time, while such adjustments were more often than not, upward. Interestingly, they noted that the percentage size of expert adjustments is predictable for about 44% of its variation, and even the direction of those adjustments is predictable to some extent. The size of expert adjustments depends strongly on past adjustments, about three times as much as it depends on past model-based forecast errors.

Franses & Legerstee (2010) expressed model-based SKU-level forecasts as a linear function of past sales. The expert-adjusted forecasts were also assumed to be a linear function of past sales leading essentially to the forecasting scheme of the experts *nesting* the forecasting scheme of the model. They constructed a test regression to assess if the expert forecasts differ from the model forecast, and followed Clark & McCracken (2001)'s recommendation in conducting an ENC-NEW test to assess whether the Root Mean Squared Prediction Error (RMSPE) of the expert is significantly lower than that of the model. They concluded that more often than not experts' forecasts differ significantly from model forecasts. They also found that when the expert yields a significant positive contribution to forecast quality, the final forecast's improvement in terms of RMSPE is about equally large, as is the deterioration in case the expert does not significantly outperform the model. So, in general, expert forecasts are not necessarily better than the model forecasts.

In another paper, Franses & Legerstee (2011a) examined linear combinations of expert and model based forecasts, with the aim of improving the final forecast (judgmentally adjusted one). For this reason, they calculated the RMSPE of a cohort of different combinations of the forecasts. To gauge how the optimal weights can be explained by experts' characteristics and their behaviour (age, position, number of products, etc.) they created a weighted summation of such variables and estimated the weights through OLS. They concluded that the combination leads to improvements in 90% of the cases, with the unconditional weights being close to 50-50%, albeit with strong variation among the experts.

The same authors then moved to investigate whether the forecasting horizon affects the accuracy of the experts' adjustments (Franses & Legerstee (2011b)). They followed a similar methodology to that employed by Franses & Legerstee (2010), by calculating ENC-NEW values for the different horizons. They concluded that while all horizons experience forecasting interventions, the horizons that are most relevant to the expert involved in the process are those that show the greater overweighting of the expert adjustment. At the same time, experts tend to perform better for distant horizons rather than shorter ones.

Franses et al. (2011) proposed a method to estimate the key parameters of the linex and the lin-lin loss functions. They deduced that the experts use an asymmetric loss function, most likely the asymmetric absolute loss function, with forecasts that are too low having a weight (in the loss function) on average 40% higher than forecasts that are too high.

Note that the modelling attempts selectively discussed above are statistical rather than mathematical in nature; that is, there is a reliance upon regression type models to identify the effect of various parameters. No analytical modelling has been conducted in the area of adjustments of forecasts and this remains an important avenue for further research. Moreover, and as discussed in the previous section, there has been only one modelling attempt in the area of adjusting order replenishment decisions (Syntetos et al. (2011)). The emphasis though has been on the interactions between forecasting and stock control (rather than only the latter which constitutes the focus of our research) by means of using SD, rather than analytical modelling.

2.2. Single selling season modelling

A second stream of research that is relevant to this paper studies information updating and/or behavioral aspects in the newsvendor model. Concerning the first, most works assume that the newsvendor has two ordering opportunities for a single selling season, and that in-between the two ordering opportunities, the newsvendor can collect sales data and use it for updating the demand distribution. Two ordering opportunities for products with a short selling season are very common in the fashion industry, where so-called Quick Response Systems give fashion companies the possibility to order products again during the selling season in case the actual customer demand exceeds the expected demand.

An early work in this area is the one by Bradford & Sugrue (1990), who assumed that the selling season is divided into two equal time-periods, and that replenishments are allowed at the beginning of both periods. The proposed model utilizes a Bayesian forecasting procedure that uses demand information gathered in the first period to update the forecast for the second period. The forecast is then incorporated into the model to derive optimal stocking policies that maximize the expected profit over the season. Fisher & Raman (1997) studied a similar problem and took account of two constraints. First, they assumed that the capacity of the supplier at the second order opportunity is limited, which means that in many cases a significant share of the total order has to be placed in the first period. Secondly, they assumed that in case the newsvendor places an

order, a minimum lot size has to be ordered to ensure that the products can be produced economically.

Gurnani & Tang (1999) assumed that the unit cost at the second order opportunity is uncertain and could be higher (or lower) than the unit cost at the first order opportunity. To determine the profit-maximizing ordering strategies, the retailer has to evaluate the trade-off between a more accurate forecast and a potentially higher unit cost at the second order opportunity. Similar models were developed by Choi et al. (2003) and Serel (2009).

Milner & Rosenblatt (2002) studied the case where the newsvendor is able to adjust an order that was previously placed at the supplier. However, adjusting an order is associated with a per-unit adjustment penalty. The authors further assumed that the time of the adjustment of the order is flexible, and that the decision maker can decide how much demand to observe before adjusting the order.

Choi et al. (2004) assumed that the newsvendor has the option to use different delivery modes for placing a single order. The unit delivery cost (and hence the cost of the product) was formulated as a decreasing function of the lead-time. If the newsvendor decides to order later, s/he can collect information on the expected market demand and update the demand forecast using a Bayesian approach. Simultaneously, delaying the order entails that the newsvendor has to select a faster delivery mode, which will lead to higher purchase cost. The authors derived an optimal order policy that balances the cost reduction from reduced uncertainty and the cost increase from higher purchase cost. Li et al. (2009) studied the case where the newsvendor can order twice before the end of the selling season and where the timing of the second order is also a decision variable, and Kim (2003) and Yan et al. (2003) considered the situation where multiple orders are possible. A comparison of two different Bayesian updating models (one with the revision of an unknown mean, and the other with the revision of both an unknown mean and variance) may be found in the work of Choi et al. (2006).

Closely related to our work is also a research stream that studies behavioral aspects of newsvendor decisions, often employing laboratory experiments. Although some research has started to be published, this is still an area that requires further exploration according to Asgari et al. (2016).

Schweitzer & Cachon (2010), for example, investigated decision making in a newsvendor setting involving a group of MBA students. The authors identified several decision biases that led to smaller or larger order quantities than predicted by the newsvendor model and consequently less-than-maximum profits. Examples include a demand chasing and a pull-to-center behavior. Bolton et al. (2012) later showed that procurement professionals display the same ordering behavior, and that Schweitzer and Cachon's results are not limited to students. Several authors have since then tried to explain these decision biases and to find effective remedies, for example Bostian et al. (2008) and Bolton & Katok (2008), who explored the role of learning in the pull-to-center effect, or Ren & Croson (2013), who found evidence that overprecision displayed by the newsvendor may lead to wrong ordering decisions. In a comparison of decision

making of American and Chinese decision makers, Feng et al. (2011) found that cultural aspects may have an additional influence on the pull-to-center effect. Other works in this research stream are those of Kocabykoglu et al. (2016), Schiffels et al. (2014) and Moritz et al. (2013), among others.

3. THE MODEL

3.1. Experimental framework

The experimental framework employed in this paper is illustrated in Figure 1. First, based on a forecast of the mean and the standard deviation of the demand obtained from historical sales in conjunction with a hypothesized demand distribution (e.g. Normal, for instance), the inventory manager ends up with an estimated demand distribution of the random variable D_F . At the same time or earlier, the decision maker observes a signal modeled by the random variable D_S that conveys some information over and above what the historical data may indicate about the future demand behaviour. That is, the signal conveys some information that has not been included in the production of the statistical forecast. Such a signal could be sales data of related products (cf. Choi et al. (2004); Choi (2007)) or the feedback the decision maker receives from customers who perhaps indicate that they are planning to buy more or less than in past sales periods. Additionally, the signal may convey important information on the competition, e.g. on a lowering of the price of a competitive product potentially resulting in lower sales for the organization. Finally, the signal may also be not case-specific (i.e., not related to the product/organization under concern), but rather convey some more general information on the current state-of-the-world. Regardless, the signal will trigger a reaction on the part of the decision maker who will be expecting that demand in the selling period will be higher or lower than that originally predicted.




Figure 1: Experimental framework

In contrast with previous research in the area of information updating under a newsvendor model, we realistically assume that the signal observed is not necessarily correct. For example, high sales of a mobile phone of a certain brand could indicate that customers will also be interested in a new tablet PC of the same manufacturer. However, if customers were not satisfied with the mobile phone, they might refrain from buying the tablet PC, thus making the signal worthless. It is obvious that the decision maker has to consider the informativeness of the signal in making an adjustment, where such ‘informativeness’

could be conveniently expressed as the probability of the signal being correct (p). We also realistically assume that the signal informativeness modeled by the probability p is not accurately known by the decision maker when setting the order quantity. The quality of information could be improved over time if decision-making is associated with a learning process that permits the evaluation of the impact of the estimated informativeness and consequently improves knowledge about p .

By considering the estimated demand, the signal and its estimated probability, the decision maker uses a ‘judgment-based inventory policy’ that permits deriving an optimal ordering strategy taking into account the signal and its informativeness. The inventory policy may also facilitate learning to efficiently utilize the signal. Subsequently, demand is realized and the selling season starts. After the start of the selling season, no additional orders may be placed at the supplier due to the rather long lead-time associated with the product. Scenarios such as the one described here are often observed in practice.

3.2. Notation and assumptions

For the remainder of the paper the following notation is used.

- p : probability that the signal is correct
- D_k : random variable describing the demand k (k element of F, C, W, p) and the **S**ignal ($k = S$) :
 - $k = \mathbf{F}$, refers to an estimated demand distribution based on the results of a **F**orecasting procedure, for both the mean and standard deviation of demand, in conjunction with a hypothesized theoretical distribution (which we reasonably assume to be Normal - although the results can be applied to other distributions) - hereafter ”estimated demand distribution”, or simply ”estimated demand” when referring to the relevant variable.
 - $k = \mathbf{S}$, refers to an estimated **S**ignal distribution based on a judgmental estimate of the mean, which may be positive or negative and standard deviation of the signal, in conjunction with a hypothesized theoretical distribution (which we assume to be Normal - although the results can be applied to other distributions) - hereafter ”estimated signal distribution”, or simply ”signal” when referring to the relevant variable.
 - $k = \mathbf{C}$ if the signal is **C**orrect, i.e. $D_C = D_F + D_S$
 - $k = \mathbf{W}$ if the signal is **W**rong i.e. $D_W = D_F$
 - $k = \mathbf{p}$ if a mixture of D_C and D_W is considered: $D_k = pD_C + (1 - p)D_W$
- μ_k : the mean of D_k
- σ_k : the standard deviation of D_k

- $F_k(\cdot)$: distribution function of D_k
- $f_k(\cdot)$: probability density function of D_k
- Q_0 : order quantity if the signal is assumed to be 100% wrong, i.e., the order quantity is calculated without taking the signal into account
- Q_1 : order quantity if the signal is assumed to be 100% correct
- $Q_p = Q_0 + A$: order quantity that is calculated based on the forecast and the signal (that is assumed to be correct with probability p) where A is the adjustment of the order quantity
- C : cost per unit for the newsvendor, i.e. price charged by the supplier
- M : merchant price per unit charged by the newsvendor
- S : shortage penalty per unit short
- V : salvage value per unit
- $h=C-V$: the unit overage penalty
- $u = M - C + S$: the unit underage penalty.

Moreover, we make the following assumptions:

1. We consider a product with long lead times and a short selling season, which does not permit re-ordering. Products that cannot be sold during the season have to be sold at a salvage price V , whereas shortages lead to penalty costs of S per unit short.
2. The adjustment (A) made based on the signal can be either positive or negative; however, it may not be smaller than the negative value of Q_0 (i.e., the final order quantity $Q_0 + A$ has to be non-negative). Thus, $A \geq -Q_0$.
3. We assume that the signal is correct with probability p and incorrect with probability $(1 - p)$. The Bernoulli representation of the correctness of the signal may be restrictive when such information is only partially correct (incorrect). Such a representation of the problem would require a contextual qualification of the signal. In this paper, we wish to derive some initial insights on the practice of adjusting order quantities, and as such, we leave this issue for further research; please refer also Section 6.

3.3. Model development

To deal with the signal and its probability, the proposed judgment-based inventory policy illustrated in Figure 1 can take two distinct functional forms (see Figure 2): 1) A process according to which the signal is used to adjust the initially calculated order quantity; or 2) A process that relies on an update of the estimated demand distribution (that takes the signal into account) followed by a

classical newsvendor-type optimization. The former corresponds to a situation where in light of new information, the originally calculated classical newsvendor order quantity is adjusted. The latter reflects claims originated in the forecasting literature (Goodwin (2002)) that in light of new information perhaps, it may be best to produce an entirely new order quantity based on an updated estimated demand distribution without anchoring on the originally calculated one¹.




Figure 2: Modelling of the signal implications

For the **first functional form**, we propose two approaches to model and optimize the modified newsvendor problem.

- The first approach models and optimizes the adjusted order quantity Q_p taking into account the signal and its informativeness. We rewrite the classical newsvendor as a weighted double newsvendor problem where the expected profit function is expressed as the sum of two newsvendor-type expected profits, weighted with probability p . This approach allows us to derive a new adjusted order quantity (different from the classical newsvendor one).
- The second approach models and optimizes the adjustment, A , to be superimposed on the classical order quantity, Q_0 , by integrating the signal distribution and the probability p associated with that signal. Hence, the classical newsvendor problem is rewritten to introduce a new decision variable, A , representing the adjustment of the order quantity as a reaction to the signal.

The **second functional form** relies upon a classical newsvendor formulation where the estimated demand distribution is updated and written as a distribution mixture of the estimated demand and the signal distributions. In contrast with the first functional form, the probability p is included in the estimated demand distribution.

Although the two functional forms are different from a modelling and mathematical development point of view, we show that they are equivalent in terms

¹The actual argument offered in the forecasting literature is that in the light of some new information (e.g. under a disruption to the demand pattern) it may be best to replace the statistical forecast with a purely judgmental one rather than adjusting it judgmentally.

of providing the same optimal ordering (and adjustment) decisions. As we will show in Section 4, such an equivalence permits the design of two complementary learning processes about the key parameter p , which is not known exactly by the inventory manager when deciding on the ordering and adjustment policy. The first and second functional forms are now discussed in detail in sections 3.3.1 and 3.3.2, respectively.

3.3.1. A judgment-based newsvendor model

Practical parametric approaches to inventory management rely upon an explicit demand distribution assumption and the employment of a forecasting method for estimating the moments of such a distribution (typically mean and variance). The decision parameters in the inventory model are then based on the estimated demand distribution, obtained from the results of the forecasting procedure. Here we assume that the demand distributional function (say Normal) remains unchanged over the planning horizon and is known (or reasonably decided upon at the beginning of the planning horizon either based on past data or the type of product we deal with), but both its moments (mean and standard deviation, for a 2-parameter distribution such as the Normal) are estimated (based on past data). Thus, the decision maker receives an estimated demand distribution of D_F , which consists of the type of distribution (e.g. Normal) and a forecast of the first two moments. Based on this information, the decision maker calculates an initial order quantity Q_0 using a classical newsvendor approach. The profit of the newsvendor associated with this decision is as follows:

$$\pi_F = \begin{cases} (M-C)Q_0 - S(x-Q_0) & \text{if } x \geq Q_0 \\ Mx + V(Q_0 - x) - CQ_0 & \text{if } x < Q_0 \end{cases} \quad (1)$$

where x is the realized demand.

After observing a signal, which is modelled with the random variable D_S (of the same type, e.g. Normal) with a mean μ_S and a standard deviation σ_S , the decision maker may adjust the order quantity. As discussed above, two approaches could be applied to rewrite the classical newsvendor problem in order to take into account the signal and its informativeness measured by the probability p :

1. The first modelling approach extends the profit function provided in Eq. (1) to optimize a new order quantity Q_p which takes into account the fact that the signal could be correct with a probability p and wrong with a probability $1 - p$.
2. The second modelling approach extends Eq. (1) by adding a profit component enabled by an adjustment quantity A . Hence, this approach uses the optimal order quantity solving Eq. (1), i.e. the classical newsvendor solution Q_0 , and derives an adjustment quantity A which is optimized.

Although these extensions to Eq. (1) are clearly different, they are equivalent in terms of the resulting order quantity since Q_p could be written as $Q_0 + A$. Here, we discuss the first modelling approach in detail and we present the second one in *Part A of the on-line supplement*.

In the presence of a signal with a level of informativeness p , the expected profit realized by the newsvendor is equal to the sum of the classical newsvendor profit subject to demand D_C if the signal is correct, and another newsvendor profit subject to demand D_W if the signal is wrong. The former (or latter) expected profit is conditional on the fact that the signal is correct (or wrong) and, consequently, is obtained with a probability p (or $1-p$). Thus, the expected profit function under judgment is written as:

$$\pi_p(Q_p) = p\pi_C(Q_p) + (1-p)\pi_W(Q_p) \quad (2)$$

where the newsvendor-type expected profit functions for $i = C, W$ are as follows:

$$\begin{aligned} \pi_i(Q_p) &= (M - C + S)\mu_i - (M - C + S) \int_0^{Q_p} (x - Q_p) f_i(x) dx \\ &- (C - V) \int_{Q_p}^{+\infty} (Q_p - x) f_i(x) dx \end{aligned} \quad (3)$$

It is straightforward to observe that the second derivative of the expected profit function is the sum of two newsvendor-type second derivative functions and by using the fact that $pf_C(x) + (1-p)f_W(x)$ is positive, it is a straightforward exercise to show that the expected profit function is concave and consequently the optimal order quantity can be derived from the first-order derivative condition. The optimal order quantity (Q_p^*) for the case where the signal and its probability are taken into consideration is given by:

$$pF_C(Q_p^*) + (1-p)F_W(Q_p^*) = \frac{M - C + S}{M - V + S} \quad (4)$$

We note that the result provided in Eq. (4) is valid for any assumption regarding the distributions of the estimated demand and the signal since $F_W(\cdot) = F_F(\cdot)$ and $F_C(\cdot)$ could be derived using the convolution of $f_F(\cdot)$ and $f_S(\cdot)$.

3.3.2. Estimated demand distribution updates

A second functional form to model the process of judgmental intervention is by means of including the signal probability p not in the expected profit function as written in Eq. (2), but rather in the estimated demand distribution the decision maker faces. By doing so, the decision maker adjusts the estimated demand distribution to include the signal and then applies the classical newsvendor problem in conjunction with the adjusted demand estimate.

We assume that the adjusted demand random variable is D_C with a probability p and D_w with probability $(1-p)$. Thus, we write the adjusted demand random variable D_p as a mixture of the estimated demand and the signal:

$$\begin{aligned} D_p &= pD_C + (1-p)D_W \\ &= p(D_F + D_S) + (1-p)D_F \end{aligned} \quad (5)$$

Independent of the assumption regarding the estimated demand and the signal distributions, if the inventory manager is able to derive the density function $f_p(\cdot)$ associated with the mixture variable D_p , s/he could apply the classical newsvendor solution subject to D_p . That is, by considering the variable D_p with its distribution function $F_p(\cdot)$ and its density function $f_p(\cdot)$, the decision maker could solve the problem in hand by optimizing the expected profit of the classical newsvendor subject to the demand D_p :

$$\begin{aligned} \pi_p(Q_p) = (M - C + S) \mu_p - (M - C + S) \int_0^{Q_p} (x - Q_p) f_p(x) dx \\ - (C - V) \int_{Q_p}^{+\infty} (Q_p - x) f_p(x) dx \end{aligned} \quad (6)$$

Thus:

$$Q_p^* = F_p^{-1} \left[\frac{M - C + S}{M - V + S} \right] \quad (7)$$

$$\pi_p(Q_p^*) = (M - V + S) \int_0^{Q_p^*} x f_p(x) dx \quad (8)$$

If D_C and D_W are, for instance, normally distributed, D_p is known as the Gaussian Mixture Model (GMM), which is defined as a parametric probability density function represented by a weighted sum of Gaussian component densities. Such a GMM is not very often used in inventory control, but is widely used in signal treatment (speaker recognition systems), as well as in the fields of economics and finance (Titterton et al. (1985); McLachlan & Peel (2009)). Under the normality assumption, D_p is described by its distribution function $F_p(\cdot) = pF_C(\cdot) + (1 - p)F_W(\cdot)$ and its density function $f_p(\cdot) = pf_C(\cdot) + (1 - p)f_W(\cdot)$.

It is important to note that the GMM modelling approach is mathematically, and consequently numerically, equivalent to the modelling approach provided in the previous section where the probability p was included in the expected profit function. In fact, Eq. (2) becomes Eq. (6) by applying in the former the expression of the GMM density $f_p(\cdot) = pf_C(\cdot) + (1 - p)f_W(\cdot)$.

As we will show in the simulation study, – despite being mathematically equivalent – the two approaches permit the development of two different and complementary learning processes about the actual value of the parameter p , which might not be known by the decision maker. In fact, p measures the probability that the signal information is correct. In contrast to the estimated demand which is stochastically known (by its estimated mean and standard deviation), the signal comes from a non-official source: its probability p is unknown and its realization will be embedded, if it occurs, in the demand recorded by the newsvendor. Over subsequent periods, the decision maker will improve his/her knowledge about the estimation, denoted hereafter \hat{p} , of the probability p .

By applying a judgment-based inventory policy based on the estimate \hat{p} , the newsvendor orders:

$$Q_{\hat{p}}^* = F_{\hat{p}}^{-1} \left[\frac{M - C + S}{M - V + S} \right] \quad (9)$$

The actual expected profit is, however, a function of the actual value of p , and it is given by

$$\pi_{\hat{p}} = \pi_p(Q_{\hat{p}}^*) \quad (10)$$

In the extreme cases the decision maker could either totally ignore the signal by setting $\hat{p} = 0$, or s/he could totally trust it by setting $\hat{p} = 1$. By 100% ignoring (100% trusting) the signal, the decision maker orders Q_0^* (Q_1^*), which corresponds to the classical newsvendor solution associated with the demand random variable D_W (D_C), and the actual expected profit is given by $\pi_0 = \pi_p(Q_0^*)$ ($\pi_1 = \pi_p(Q_1^*)$). All the following results will relate to the adjustment case with an estimate \hat{p} of the probability p and the extreme cases: $\hat{p} = 0$ when the signal is ignored and $\hat{p} = 1$ when it is 100% trusted.

As previously discussed, the analytical results provided in this section are independent of the assumption regarding the estimated demand and signal distributions. In the following numerical analysis, we assume without loss of generality that the estimated demand and signal are normally distributed.

4. NUMERICAL ANALYSIS

The aim of this section is threefold:

1. To illustrate the implications of the judgement-based policy on inventory performance. In particular, we compare the adjustment case with the extreme cases where the signal is 100% ignored or 100% trusted.
2. To provide insights on the impact of the different input parameters on the judgement-based inventory policy and the associated benefits.
3. To derive insights on the inaccurate estimation of the signal probability p .

More importantly, the study of these three issues will help us appreciate some aspects of learning which can be used by the decision maker when estimating the unknown probability p .

For this purpose, we assume that the decision maker faces a normally distributed estimated demand with mean $\mu_F = 100$ and standard deviation $\sigma_F = 20$. The signal is also normally distributed and we consider three possibilities for the signal mean, $\mu_S = \{-30, 0, 60\}$, in conjunction with three possible values of $\sigma_S = \{10, 20, 50\}$. Regarding the cost, we consider two families of products: i) high margin products where the unit costs are set such that $u = 10$ and $h = 5$, and ii) low margin products such that $u = 0.5$ and $h = 1$. Notice that other settings of the parameters inputs were tested but lead to similar results to those discussed below. The choice of the above parameters is motivated by the fact that they allow for the collective illustration of various types of behavior.

It is important to note that the newsvendor derives the order quantity under the judgement-based policy, Q_p^* , by using the GMM demand distribution whose density function depends on the probability p .

Depending on p as well as the given moments of the estimated demand and signal distributions, the resulting demand mixture shape changes. For some particular cases (such as the one shown on the left-hand side pane of Figure

3), the two demand densities ($f_C(x)$ and $f_W(x)$) overlap and consequently the newsvendor could not easily distinguish between a demand realization resulting from the correct signal or the wrong one. In such cases, the learning process about the probability p is not straightforward, as we will show in the simulation study. In contrast, in cases such as the one shown on the right hand side pane of Figure 3, the density functions do not overlap, making it easier to distinguish between alternative demand realizations. Notice that the case of demand multi-modality (illustrated on the right hand side pane of Figure 3) has been previously observed and commented upon in the academic literature, albeit in a different experiment setting than ours. Providing more detail, the investigation of Hanasusanto et al. (2015) motivates the possible multimodality assumption of a newsvendor demand distribution from four practical cases: *i*) the launch of new products; *ii*) the existence of large customers among the newsvendor customers; *iii*) the emergence of a new market entrant and *iv*) the case of different popularity states for some apparel products.




Figure 3: Mixture demand shapes

When relying upon the observed demand realization to learn about the probability p , the newsvendor has to qualify the degree of similarity between the correct (estimated demand distribution + signal distribution) and wrong (i.e. estimated demand distribution only) distributions. To measure this similarity and the degree of overlap between the two distributions, we could use the squared Hellinger parameter, which, if both distributions are normally distributed is expressed as follows:

$$H^2 = 1 - \sqrt{\frac{2\sigma_C\sigma_W}{\sigma_C^2 + \sigma_W^2}} e^{-\frac{1}{4} \frac{(\mu_C - \mu_W)^2}{\sigma_C^2 + \sigma_W^2}} \quad (11)$$

The optimal order quantity under the judgment-based policy is a function of the probability p , and takes values in-between those corresponding to the extreme cases of the order quantities calculated when the signal is either 100% ignored or 100% trusted (cf. Figures 4a and 4b, respectively). Adjusting with the actual value of p (if it is accurately known) permits the decision maker to improve the expected profit compared to the *ignore* or *trust* policies (cf. Figures 4c and

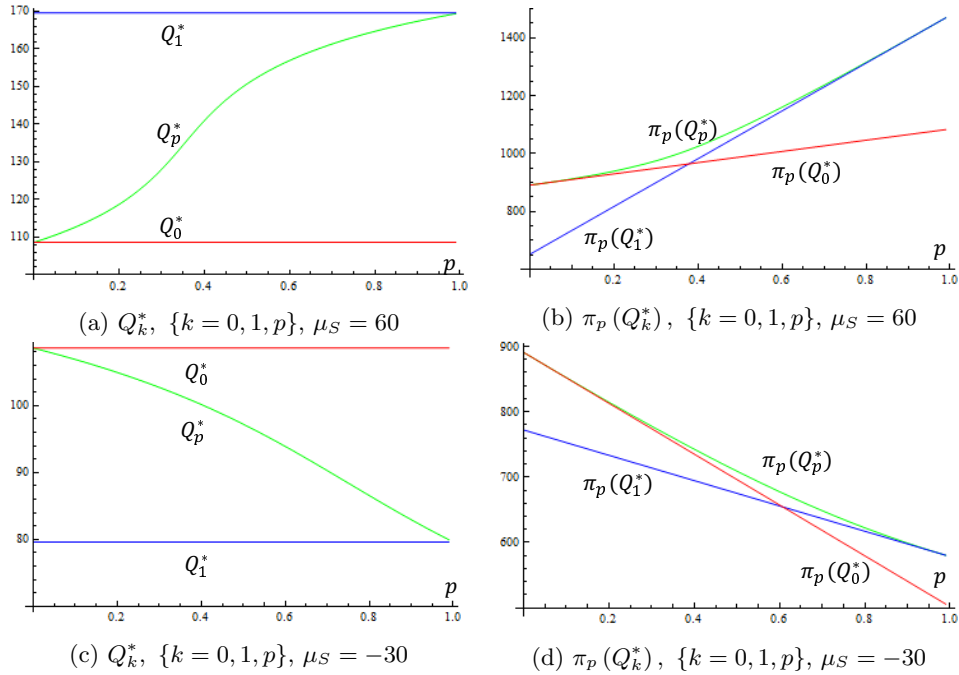


Figure 4: Optimal order quantity (and associated profit) as a function of p

4d). The latter policies are suboptimal if the value of p is known. Note that depending on the input parameters, particularly the moments of the estimated demand distribution, the order quantity under the *ignore* policy could be higher or lower than that related to the *trust* one, and the order quantity under adjustment could consequently be decreasing or increasing in p . It is also worth noting that there exists a threshold value of p below which it is better for the decision maker to ignore the signal, while above it trusting the signal is the best alternative if the adjustment is not applied. Such a threshold, denoted hereafter as $p_{trust-ignore}$, could be calculated by the decision maker and would be useful in the learning process about the unknown probability p . $p_{trust-ignore}$ solves:

$$\pi_{p_{trust-ignore}}(Q_0^*) = \pi_{p_{trust-ignore}}(Q_1^*) \quad (12)$$

Both Eqs. (2) and (6) permit calculating $p_{trust-ignore}$: the latter calculates it numerically, whereas the former allows the derivation of an analytical expression:

$$p_{trust-ignore} = \frac{\pi_W(Q_1^*) - \pi_W(Q_0^*)}{(\pi_W(Q_1^*) - \pi_W(Q_0^*)) - (\pi_C(Q_1^*) - \pi_C(Q_0^*))} \quad (13)$$

The existence and the uniqueness of the threshold $p_{trust-ignore}$ can be proven by observing that the functions $\pi_p(Q_0^*)$ and $\pi_p(Q_1^*)$ are linear in p as stated in Eq. (2) and illustrated in Figs. 4(b) and 4(d). By contrasting the starting and the ending points of these two linear functions, we can prove the existence

and the uniqueness of the threshold $p_{trust-ignore}$. As shown in Figure 5, the threshold $p_{trust-ignore}$ depends on the signal moments and on the product cost structure. For a high margin product, the threshold decreases with the signal average: the space where trusting is better than ignoring the signal is more important. In fact, for a high-margin product, the shortage penalty is more important and the newsvendor tends to the less risky policy of ordering a higher quantity when trusting the signal. For the case of a low-margin product, the threshold encouraging trusting the signal increases with the signal average, since the overstock risk is more penalizing.




Figure 5: $p_{trust-ignore}$ as a function of signal parameters and unit costs

Over successive periods, the decision maker could measure and update the frequency of periods when *trusting* the signal provides a higher profit than *ignoring* it: $freq(trust \geq ignore)$, as well as the opposing tendency: $freq(ignore \geq trust)$. S/he may do so by contrasting the observed demand at the end of the period with Q_0^* and Q_1^* and deducing if trusting is better or not than ignoring the signal.

If $freq(trust \geq ignore) \geq freq(ignore \geq trust)$ (or the difference between the two is higher than a determined value), then the actual p is higher than $p_{trust-ignore}$; otherwise, it is lower than this threshold. We note that the rule under concern helps position the actual p on the interval $[0,1]$ and thus avoiding ending up with a situation where for instance the signal is ignored when the probability p could be shown, thanks to this rule, to belong to the interval $[p_{trust-ignore}, 1]$.

To illustrate the simplicity and the benefit of such a basic rule, let us consider a numerical example ($\mu_F = 100$, $\sigma_F = 20$, $\mu_S = -30$, $\sigma_S = 20$) for a high margin product ($u = 10$ and $h = 5$). Based on these control parameter values the order quantities under the *ignore* and *trust* policies are $Q_0^* = 108$ and $Q_1^* = 82$, respectively. By applying Eq. (13) with this numerical setting, $p_{trust-ignore} = 0.64$. Let us assume an actual value of $p = 0.8$ and generate 10 random demand realizations based on this signal informativeness: demand={91, 66, 94, 75, 77, 68, 96, 32, 26,100}. By contrasting these demand realizations with the values of Q_0^* and Q_1^* , the inventory manager could verify that trusting the signal (using Q_1^* as an order quantity) brings a higher profit than ignoring it (using Q_0^* as an order quantity) in 6 realizations over 10. Consequently, the actual value of p is

most likely to belong to the interval $[0.64,1]$.

In the case where the decision maker estimates the probability p inaccurately, we show the existence of a second threshold denoted by \hat{p}_{adj} above/below which the adjustment policy – even with an inaccurate estimate – allows increasing the expected profit compared with both the 100% *trust* and *ignore* policies.

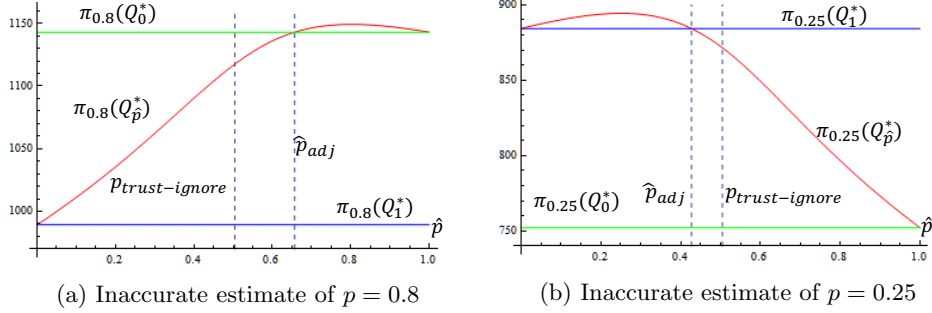


Figure 6: The impact of inaccurate estimation of the probability p

For $\mu_F = 100$, $\sigma_F = 20$, $\mu_S = 60$, $\sigma_S = 20$, $u = 10$ and $h = 5$, let us consider the case where the actual p is set equal to 0.8 (0.25) as illustrated in Figure 6a (6b). The decision maker is able to calculate $p_{trust-ignore}$, and based on the comparison $freq(trust \geq ignore)$ vs. $freq(ignore \geq trust)$, s/he could establish that p should be in the interval $[p_{trust-ignore}, 1]$ ($[0, p_{trust-ignore}]$). Within such an interval, if the newsvendor uses a value of \hat{p} higher (lower) than \hat{p}_{adj} , the adjustment policy offers a positive benefit compared with the *trust* (*ignore*) policy even if the estimate is not accurate. The reader is referred to *Part C of the on-line supplement* for a further illustrative example on the impact of adjustment on economic performance.

\hat{p}_{adj} , which solves $\pi_p(Q_{\hat{p}}^*) = \pi_p(Q_1^*)$ if $freq(trust \geq ignore) \geq freq(ignore \geq trust)$ and $\pi_p(Q_{\hat{p}}^*) = \pi_p(Q_0^*)$ in the opposite case, cannot be calculated by the decision maker since p must be known to derive the solution. In the next section, in order to assess the quality of the learning process from which \hat{p} is obtained we will contrast the estimate \hat{p} with the value \hat{p}_{adj} .

5. IMPLEMENTATION AND FOSTERING LEARNING

In this section, we consider a multi-period setting of the problem composed of N newsvendor problems where the decision maker tries to dynamically learn and apply the judgement-based inventory policy to take into consideration the signal and its probability. This scenario could occur in the consumer electronics industry or the publishing industry, for example, where companies face a new newsvendor problem whenever a new version of an existing product (say newspaper, new issue of a magazine, etc.) is brought into the market, which takes place during a short selling period (e.g. daily, weekly) basis. So essentially we consider a newsvendor framework with short selling periods (to allow

for an update of the learning process at the end of each selling period), with a long lead-time (that justifies the unique replenishment opportunity) where the unsold items cannot be transferred between successive periods.

For the purpose of accurately estimating p , we propose and compare three techniques:

1. An estimate based on the mixture demand observation: in each period t , the decision maker observes the demand realized at the end of the period and dynamically derives the \hat{p}_t that provides the best “fit” between the set of demands observed from period 1 to t and the mixture demand $D_p = pD_C + (1 - p)D_W$. Such a learning process results in an estimate \hat{p}_t^{demand} , which is a pure statistical estimate obtained independently of the inventory costs. It is important to note that the estimate obtained in period t is independent of the ones calculated from period 1 to $t - 1$ since it only relies on the observed demand realizations which are independent of the learning process.
2. An estimate based on the profit observation: as shown in the first functional form to model judgment (Section 3.3.1), the expected profit under the judgment-based policy is a mixture of the profits obtained under 100% trusting and 100% ignoring the signal: $\pi_p(Q_p) = p\pi_C(Q_p) + (1-p)\pi_W(Q_p)$. For each period t , the decision maker calculates the realized profit and contrasts it with what that profit would be if s/he 100% ignored or 100% trusted the signal. Such a learning process results in an estimate of \hat{p}_t^{profit} , which is (in contrast to \hat{p}_t^{demand}) a function of the inventory costs as well as the estimates of p realized before period t (\hat{p}_k^{profit} for $k = 1 \dots t - 1$).
3. An estimate based on any combination of \hat{p}_t^{demand} and \hat{p}_t^{profit} . As mentioned earlier, the outcome of the first two learning techniques relies on different types of observations. The rationale behind the proposal of this technique is to combine the results of the two other techniques discussed above. The weights assigned to each technique may of course vary but we will experiment with the simpler option of a straight average combination, by using the estimate $\hat{p}_t^{average} = (\hat{p}_t^{demand} + \hat{p}_t^{profit})/2$.

\hat{p}_t^{demand} and \hat{p}_t^{profit} are motivated by the two functional forms presented in Section 3 (Figure 2): \hat{p}_t^{demand} stems from the modelling approach assuming that the demand distribution is adjusted before applying the classical newsvendor problem, whereas \hat{p}_t^{profit} stems from the modelling approach where the expected profit is written as the sum of two weighted newsvendor profits.

Although the two functional forms are equivalent, from an optimality point of view, the associated \hat{p}_t^{demand} and \hat{p}_t^{profit} provide different and complementary learning techniques which will permit better and faster convergence to the actual value of the probability p . The rationale behind combining both as proposed in $\hat{p}_t^{average}$ is to mix the two functional forms in the learning process.

The calculation of \hat{p}_t^{demand} is derived by using Maximum Likelihood Estimation, which is adapted in our case to permit such a calculation by assuming that the moments of the estimated demand and the signal moments are known (cf. *Part*

B of the on-line supplement). For a given data set $D=\{d_n\} n= 1...t$ associated with the observed demand realizations up to period t , \hat{p}_t^{demand} verifies:

$$\sum_{n=1}^t \frac{f_C(d_n) - f_W(d_n)}{\hat{p}_t^{demand} f_C(d_n) + (1 - \hat{p}_t^{demand}) f_W(d_n)} = 0 \quad (14)$$

We present the results of two simulation instances by using the numerical values employed in Figure 6a (Figure 6b), i.e., $\mu_F = 100$, $\sigma_F = 20$, $\mu_S = 60$, $\sigma_S = 20$, $p = 0.8$ ($p = 0.25$, respectively) with the cost parameter values being $u = 10$ and $h = 5$. According to the estimated demand and signal moments, the Hellinger parameter is $H = 0.37$. The learning process about p depends on the normally generated estimated demand and signal realizations; for each of the simulated p values we present a good (Figures 7a and 8a) and, a not particularly successful, learning process (Figures 7b and 8b).

As illustrated in Figures 7a and 7b (Figures 8a and 8b), the decision maker could calculate $p_{trust-ignore} = 0.5$ and, starting from the first periods, be aware that the actual value of p is in the interval $[0.5,1]$ ($[0,0.5]$), respectively. At the end of each period, after observing the demand realization and the associated realized profit, the three estimates \hat{p}_t^{demand} , \hat{p}_t^{profit} and $\hat{p}_t^{average}$ could be updated. In order to choose among the three estimates, the decision maker calculates the posterior profit associated with each estimate and chooses the one providing the highest profit. Such a posterior calculation permits the newsvendor to switch between the three estimates (as illustrated, for instance, in period 4 in Figure 7a and in period 20 in Figure 7b).

It is important to note that for both studied values of p :

- The two learning techniques complement each other, which is accelerating convergence to the actual p .
- By applying the judgement-based inventory policy, the newsvendor starts improving the profit compared with the 100% trust (for $p = 0.8$) or 100% ignore (for $p = 0.25$) policies from the first simulation periods, even if the estimate is inaccurate. As illustrated, the value of p applied is higher (lower) than \hat{p}_{adj} for $p = 0.8$ ($p = 0.25$ respectively) from the third period onwards in the selling horizon.

For the remainder of this section, and in order to obtain better insights about the learning process, we average the results of the previously described simulations over 30 runs.

We propose measuring the learning process quality by two indicators:

- The relative estimate error measuring the difference between the estimate and the actual value of p calculated as $err^k = |\hat{p}_t^k - p|/p$, for $k=\{demand, profit, average, used\}$;
- The value of the estimate \hat{p}_t^k when contrasted with \hat{p}_{adj} . We denote this indicator by pos_t^k and set it equal to 1 if \hat{p}_t^k generates an expected profit higher than these obtained under 100% ignore and trust policies.

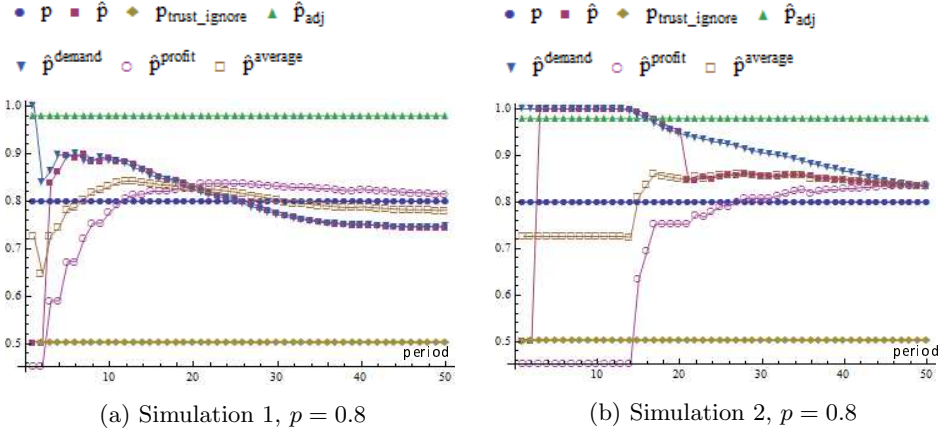


Figure 7: Learning processes for $p = 0.8$

The superscript $k = used$ describes the learning technique adopted by the newsvendor since, as we mentioned previously, s/he could switch between the three learning processes $k = \{demand, profit, average\}$ by performing the posteriori profit calculation.

Over a horizon composed of 50 periods, Figures 9a and 9b illustrate the quality of the learning process as well as the value of the estimate when contrasted with \hat{p}_{adj} . The value of \hat{p} deviates by a maximum of 10% from the actual p . Such an estimation permits the newsvendor to increase his/her profit compared with the extreme cases (signal ignored or trusted) in about 80% of the periods. To measure the convergence speed of the estimation, we illustrate in Figures 10a and 10b the error estimate value and the value of the estimate when contrasted with \hat{p}_{adj} for 10-period blocks of the whole horizon.

As can be seen, starting from the 10th period, the learning technique enables a good estimate of the parameter p with a relative error less than 5% and more than 80% of the used estimates offer a certain benefit to the newsvendor compared with the extreme cases where the signal is ignored or trusted.

5.1. The impact of the Hellinger parameter on the learning process

The speed of convergence of the learning process for the Gaussian-mixture-model fitting is known to depend on the amount of overlap among the mixture components, which is measured by the Hellinger parameter. Our learning proposal stems from a standard likelihood maximization (\hat{p}_t^{demand}) complemented by the profit observation learning measure (\hat{p}_t^{profit}). It is worth noticing that both estimates are technically easy to deploy in practice using an Excel worksheet that we provide as an on-line supplement for interested readers. Figure 11 illustrates the quality of the learning process for different values of the signal average $\mu_S = \{-20, 0, 20, 40, 60, 80, 100\}$ leading to the Hellinger parameter changing from 0 to 0.75.

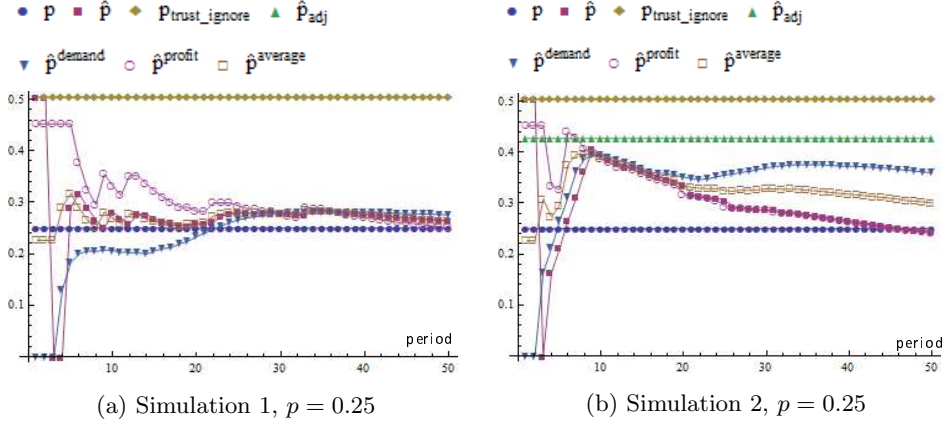


Figure 8: Learning processes for $p = 0.25$

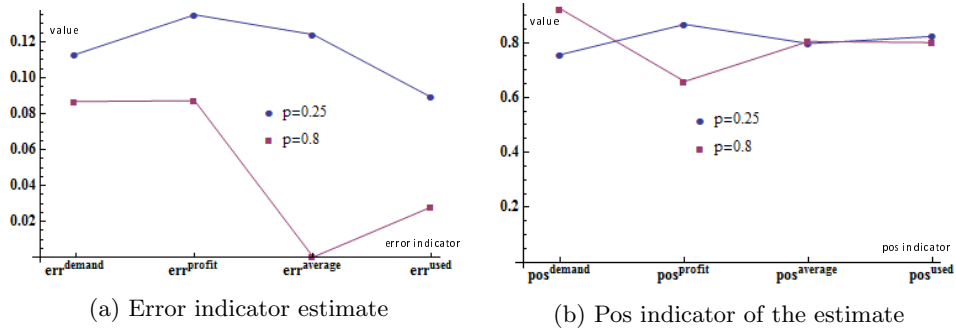


Figure 9: The learning process quality: comparison of the techniques

For a signal average outside the interval $[-20, 20]$, the estimation error, err^{used} , is less than 15% and the learning process permits the newsvendor to derive a benefit (compared with the extreme cases) with a value of the indicator pos^{used} higher than 65%. For high values of H , the error tends towards 0 and pos^{used} to 1. We also notice that the posterior rule used to choose between the three estimators (\hat{p}^k , for $k = \{demand, profit, average\}$) is relatively efficient since the one used by the decision maker \hat{p}^{used} tends to be closer to the estimator minimizing the error value.

5.2. The learning process for a dynamically changing probability p

The signal source may change over periods and consequently the moments of the signal distribution as well as the probability p should be expressed in dynamic terms.

Figure 12 illustrates the quality of the learning process for three values of the signal average $\mu_S = \{-20, 0, 60\}$ by assuming that the probability p in each period is randomly chosen in the interval $\{[0.1, 0.3], [0.3, 0.7] \text{ or } [0.7, 0.9]\}$.

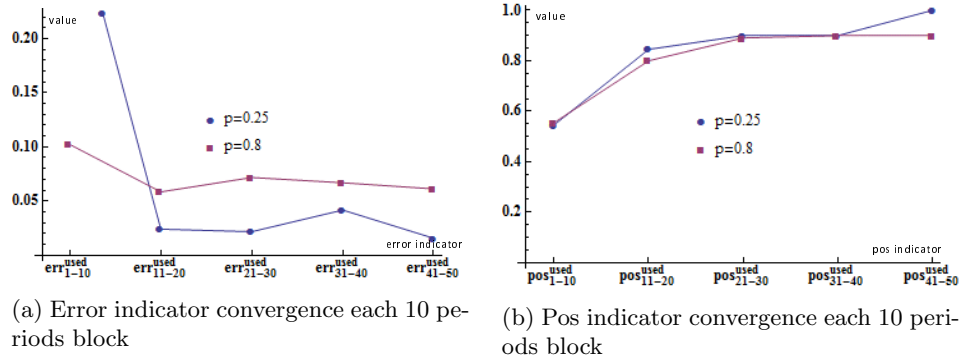


Figure 10: The learning process quality: evolution over periods

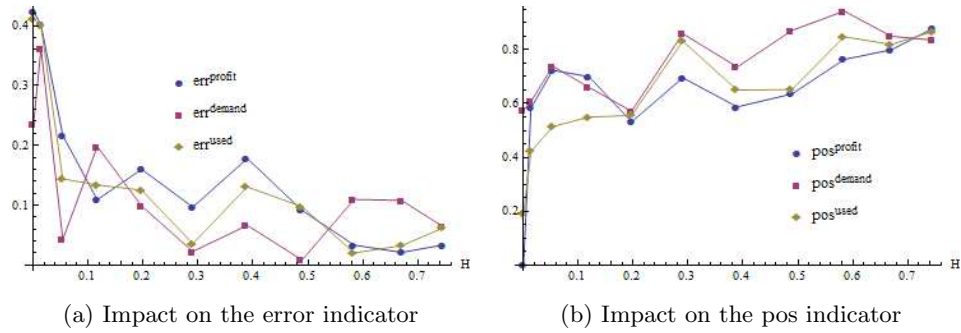


Figure 11: The learning process quality: impact of the Hellinger parameter

All illustrated by the simulation, realized over 50 periods, except for the case of a signal average equal to 0, the learning process is relatively good, with an error indicator less than 20% and a certain adjustment benefit higher than 60% even with inaccurate estimates. The scenario where $\mu_S = 0$, which sets the Hellinger parameter close to zero, is associated with a rather poor estimation of p : such a scenario is associated with a signal resulting in a more variable demand for the newsvendor.

In contrast to the last figure, if each period we fix the probability p and allow the estimated demand and signal moments to be chosen randomly in an interval, we observe that the proposed learning process is no longer efficient. Nevertheless, the newsvendor is better off than 100% trusting or ignoring in 30% of the cases on average, even when applying the inaccurate estimation of p . Additionally, the proposed learning process assumes that the moments of the signal distribution are known and the aim is to learn about its probability p . The signal distribution parameters may be unknown in some cases, or the newsvendor may not be able to quantify them. In such a case, a new judgement-based policy should be designed, along with relevant leaning elements.

As previously discussed, an Excel file that contains details of the learning pro-

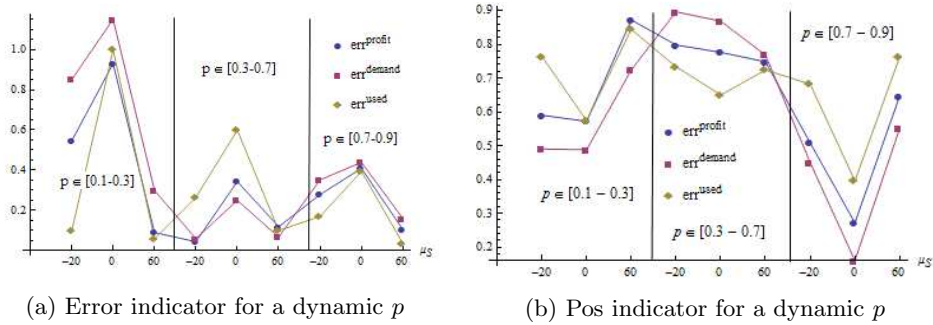


Figure 12: The learning process quality: the case of a dynamic p

cesses, along with instructions on how to use it (please refer to the on-line supplement, Part D) are introduced as supplementary material to this paper. Any inquiries with regards to this material may be directed to the authors of the paper.

6. CONCLUSION

We make a first attempt to model and analytically characterise the process and implications of judgmentally intervening into the specification of replenishment orders. We point out the scope for learning to efficiently utilise potentially important information available to the decision maker and offer some simple, spreadsheet implementable, rules to do so. We conduct our analysis under a simple newsvendor setting and show the analytical equivalence of adjusting an order quantity and deriving an entirely new one in light of the necessary update of the estimated demand distribution. Such an equivalence demonstrates the lack of sensitivity of the system to the way information is handled, as long as it is captured in a structured manner. It also signifies that the decision maker may select the procedure that s/he better understands or finds more intuitive. This is distinctly different from what we know about the process of judgmentally adjusting statistical sales forecasts, in which case (see also Section 2.1) the results depend very much on the background of the decision makers and the approach taken to 'explain' the added value of adjustments. It does, however, explain previous results obtained based on a System Dynamics approach that show the inventory-forecasting system to be more sensitive to the forecasting rather than the inventory optimisation part of the process (Syntetos et al., 2011). Consequently, it points out the need for more analytical research in this area to appreciate the scope for improving the performance of real world inventory systems.

In addition to the ease of implementing the methods proposed here, it is also important to note the minimal information that needs to be stored to enable such an implementation: i) information about price and cost of the product;

ii) prior demand data; iii) information on the signal we estimated in the past². Although we are not able to comment on profit increases resulting from the implementation of those methods (as this would depend on many assumptions on price and cost) we can point out the improvements in the calculation of the order quantity (in terms of lower error) and the very promising results we have obtained with regards to the percentage of instances where adjusting and using our recommended procedure considerably improves matters.

Our analysis is admittedly constrained, in terms of generalizability, by the newsvendor setting, especially so when it comes to the multi-period formulation to assess learning processes and their effects in practice. It does open the way though to further research that complements analytical modelling with empirical investigations or empirical based analytical modelling. Further research could also be concerned with the consideration of more general two-parameter inventory systems that in addition to an order quantity look at safety stock calculations. Our model does not apply to a situation where a product is consumed continuously over many periods with multiple ordering opportunities. If backorders were allowed, then misinterpreting the signal and ordering an incorrect quantity in one period would ultimately influence the inventory levels in the next period, and this would have to be reflected in the model. Extending our model to a newsvendor setting with multiple order opportunities (cf. also the works discussed in Section 2.2) would, however, be possible. An interesting scenario could be a situation where a demand signal is observed before the first ordering opportunity, and where the reliability of the demand signal (true, wrong) materializes before the second order opportunity. The second order opportunity would then enable the newsvendor to amend a possible misinterpretation of the

²Despite the relatively little information needed to apply our model, we do acknowledge that implementation may often be challenged, by, e.g., the required data (such as information on previously estimated signals) being readily available. To overcome such challenge, decision makers would have to raise awareness on the considerable advantages that improved forecasting and ordering policies may offer to the company and thus ensure that the required data is routinely collected. In addition, the IT infrastructure that is needed to gather such data would need to be established. A second challenge may arise from the use of ERP systems or software packages the company may have acquired to support its forecasting and ordering decisions. Using the proposed model would render necessary an integration of its implementation within such systems (or, less preferably, updating forecasts and/or order quantities outside these systems). These challenges could again be overcome by highlighting the benefits of improved forecasting and ordering policies, and by ensuring that the proposed model is effectively integrated within existing software systems to avoid extra handling effort and unnecessary disruptions. An important possibility to overcome implementation challenges also arises from using our model for training and professional development purposes. That is, the company could use the model to evaluate the adjustments of forecasters, and demonstrate to them in which cases (and under what conditions, e.g., in terms of increasing/decreasing order quantities) adjusting forecasts or order quantities improve the performance of the inventory system. As such, the proposed model would then be used to train people to understand which types of adjustments are beneficial to the company and thus should be pursued, and which ones are not. It could similarly be used for introducing internal competition (and relevant rewarding schemes) amongst colleagues to encourage improved forecasts (e.g., less bias) by means of acknowledging best performance.

signal before placing the first order.

The newsvendor setting can also be extended by considering more realistic representations of the correctness of a signal or the decision making behavior of the newsvendor and, in the spirit of inter-disciplinary research, be complemented by primary qualitative (interview obtained) data on the way managers treat and utilise external information. Regarding the representation of the correctness of the signal, this work assumed that the signal is correct with probability p and incorrect with probability $(1 - p)$. The Bernoulli representation of the correctness of the signal may be restrictive when such information is only partially correct (incorrect). Such a representation of the problem would require a contextual qualification of the signal. In addition, it is true that since the signal stems from an informal source of information or is indeed related to personal interpretation, it is very difficult to quantify or represent it by means of a stochastic distribution. Modelling the signal by a possibility distribution, or as a fuzzy variable, and mixing it in the stochastic representation of the estimated demand distribution derived from a forecasting process could be an interesting avenue for further research. With respect to the decision making behavior of the newsvendor, we note that taking account of risk aversion in our model would be interesting. In contrast to existing research on risk averse newsvendors that investigates how the newsvendor should replenish its inventory in case demand variations reduce the newsvendor's utility, investigating risk aversion in our problem setting would also pose the question of whether the newsvendor should trust the (possibly incorrect) signal if an additional risk (of the signal being incorrect) further reduces the newsvendor's utility. Finally, and with regards to some possible empirical follow-up research, it would be interesting to investigate how decision makers weight demand signals as the one discussed in this paper in making inventory replenishment decisions. While prior research has shown that practitioners often display a 'pull-to-center' effect in newsvendor type scenarios (see Section 2.2), it would be interesting to analyze under which conditions decision makers tend towards a 'pull-to-signal' behavior. The above discussed lines of research should be valuable towards extending the insights obtained from this study.

ACKNOWLEDGEMENTS

We acknowledge the financial assistance provided by TU Darmstadt and the Carlo and Karin Giersch Stiftung that permitted a research visit of Aris Syntetos to Germany and a follow-up research visit of Christoph Glock to Cardiff University. The research described in this paper was completed whilst Yacine Rekik was visiting Cardiff University during a sabbatical leave. We would also like to thank the 4 anonymous reviewers for their constructive comments that greatly helped to improve the content of this paper and its presentation.

REFERENCES

- Asgari, N., Nikbakhsh, E., Hill, A., & Farahani, R. Z. (2016). Supply chain management 1982-2015: a review. *IMA Journal of Management Mathematics*, *27*, 353–379.
- Bolton, G., & Katok, E. (2008). Learning by doing in the newsvendor problem: A laboratory investigation of the role of experience and feedback. *Manufacturing & Service Operations Management*, *10*, 519–538.
- Bolton, G., Ockenfels, A., & Thonemann, U. (2012). Managers and students as newsvendors. *Management Science*, *58*, 2225–2233.
- Bostian, A., Holt, C., & Smith, A. (2008). Newsvendor "pull-to-center" effect: Adaptive learning in a laboratory experiment. *Manufacturing & Service Operations Management*, *10*, 590–608.
- Boylan, J. E. (2016). Editorial: Reproducibility. *IMA Journal of Management Mathematics*, *27*, 107–108.
- Bradford, J., & Sugrue, P. (1990). A bayesian approach to the two-period style-goods inventory problem with single replenishment and heterogeneous poisson demands. *Journal of the Operational Research Society*, *41*, 211–218.
- Choi, T. (2007). Pre-season stocking and pricing decisions for fashion retailers with multiple information updating. *International Journal of Production Economics*, *106*, 146–170.
- Choi, T., Li, D., & Yan, H. (2003). Optimal two-stage ordering policy with bayesian information updating. *Journal of the Operational Research Society*, *54*, 846–859.
- Choi, T., Li, D., & Yan, H. (2004). Optimal single ordering policy with multiple delivery modes and bayesian information updates. *Computers & Operations Research*, *31*, 1965–1984.
- Choi, T., Li, D., & Yan, H. (2006). Quick response policy with bayesian information updates. *European Journal of Operational Research*, *170*, 788–808.
- Clark, T., & McCracken, M. (2001). Tests of equal forecast accuracy and encompassing for nested models. *Journal of Econometrics*, *105*, 85–110.
- Feng, T., Keller, R., & Zheng, X. (2011). Decision making in the newsvendor problem: Across-national laboratory study. *OMEGA: The International Journal of Management Science*, *30*, 41–50.
- Fildes, R., Goodwin, P., Lawrence, M., & Nikolopoulos, K. (2009). Effective forecasting and judgmental adjustments: an empirical evaluation and strategies for improvement in supply-chain planning. *International Journal of Forecasting*, *25*, 3–23.

- Fisher, M., & Raman, A. (1997). Reducing the cost of demand uncertainty through accurate response to early sales. *Operations Research*, *44*, 87–99.
- Franses, P., & Legerstee, R. (2009). Properties of expert adjustments on model-based sku-level forecasts. *International Journal of Forecasting*, *25*, 35–47.
- Franses, P., & Legerstee, R. (2010). Do experts' adjustments on model-based sku-level forecasts improve forecast quality? *Journal of Forecasting*, *29*, 331–340.
- Franses, P., & Legerstee, R. (2011a). Combining sku-level sales forecasts from models and experts. *Expert Systems with Applications*, *38*, 2365–2370.
- Franses, P., & Legerstee, R. (2011b). Experts adjustments to model-based sku-level forecasts: does the forecast horizon matter? *Journal of the Operational Research Society*, *62*, 537–543.
- Franses, P., Legerstee, R., & Paap, R. (2011). *Estimating loss functions of experts*. Technical Report Erasmus University Rotterdam, Tinbergen Institute.
- Goodwin, P. (2002). Integrating management judgement and statistical methods to improve short-term forecasts. *OMEGA: The International Journal of Management Science*, *30*, 127–135.
- Goodwin, P., & Fildes, R. (1999). Judgmental forecasts of time series affected by special events: does providing a statistical forecast improve accuracy? *Journal of Behavioral Decision Making*, *12*, 37–53.
- Gurnani, H., & Tang, C. (1999). Note: Optimal ordering decisions with uncertain cost and demand forecast updating. *Management Science*, *45*, 1456–1462.
- Hanasusanto, G., Kuhn, D., Wallace, S., & Zymler, S. (2015). Distributionally robust multi-item newsvendor problems with multi-modal demand distributions. *Mathematical Programming*, *152*, 1–32.
- Kim, H. (2003). A bayesian analysis on the effect of multiple supply options in a quick response environment. *Naval Research Logistics*, *50*, 937–952.
- Kocabykoglu, A., Gogus, C., & Gonul, M. (2016). Decision making and the price setting newsvendor: Experimental evidence. *Decision Sciences*, *47*, 157–186.
- Kolassa, S., Shtz, W., Boylan, J., & Syntetos, A. (2008). Judgmental changes to forecasts: higher inventories, unchanged out-of-stocks. In *The 28th International Symposium on Forecasting*. Nice, France.
- Li, J., Chand, S., Dada, M., & Mehta, S. (2009). Managing inventory over a short season: models with two procurement opportunities. *Manufacturing & Service Operations Management*, *11*, 174–184.
- McLachlan, G., & Peel, D. (2009). *Finite Mixture Models*. Wiley.

- Milner, J., & Rosenblatt, M. (2002). Flexible supply contracts for short life-cycle goods: the buyers perspective. *Naval Research Logistics*, *49*, 25–45.
- Moritz, B., Hill, A., & Donohue, K. (2013). Individual differences in the newsvendor problem: Behavior and cognitive reflection. *Journal of Operations Management*, *31*, 72–85.
- Ren, Y., & Croson, R. (2013). Overconfidence in newsvendor orders: An experimental study. *Management Science*, *59*, 2502–2517.
- Schiffels, S., Fugener, A., Kolisch, R., & Brunner, O. (2014). On the assessment of costs in a newsvendor environment: Insights from an experimental study. *OMEGA: The International Journal of Management Science*, *43*, 1–8.
- Schweitzer, M., & Cachon, G. (2010). Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science*, *46*, 404–420.
- Serel, D. (2009). Optimal ordering and pricing in a quick response system. *International Journal of Production Economics*, *121*, 700–714.
- Syntetos, A., Babai, M., Boylan, J., Kolassa, S., & Nikolopoulos, K. (2016a). Supply chain forecasting: theory, practice, their gap and the future. *European Journal of Operational Research*, *252*, 1–26.
- Syntetos, A., Georgantzias, N., Boylan, J., & Dangerfield, B. (2011). Judgement and supply chain dynamics. *Journal of the Operational Research Society*, *62*, 1138–1158.
- Syntetos, A., Kholidasari, I., & Naim, M. (2016b). The effects of integrating management judgement into out levels: in or out of context? *European Journal of Operational Research*, *249*, 853–863.
- Syntetos, A., Nikolopoulos, K., & Boylan, J. (2010). Judging the judges through accuracy-implication metrics: the case of inventory forecasting. *International Journal of Forecasting*, *26*, 134–143.
- Syntetos, A., Nikolopoulos, K., Boylan, J., Fildes, R., & Goodwin, P. (2009). The effects of integrating management judgement into intermittent demand forecasts. *International Journal of Production Economics*, *118*, 72–81.
- Titterton, D., Smith, A., & Makov, U. (1985). *Statistical analysis of finite mixture distributions*. Wiley.
- Yan, H., Liu, K., & Hsu, A. (2003). Optimal ordering in a dual-supplier system with demand forecast updates. *Production and Operations Management*, *12*, 30–45.

Supplement to “Enriching demand forecasts with managerial information to improve inventory replenishment decisions: exploiting judgment and fostering learning”

Yacine Rekik • Christoph H. Glock • Aris A. Syntetos

EMLYON Business School, DISP Lab, France

Technische Universität Darmstadt, Germany

Panalpina Centre for Manufacturing and Logistics Research, Cardiff Business School, UK

rekik@em-lyon.com • glock@pscm.tu-darmstadt.de • SyntetosA@cardiff.ac.uk

Part A: Joint ordering and adjustment optimization

If A is the adjustment of the order quantity, then the order that is placed at the supplier after the adjustment has been made is $Q_0 + A$. If the signal is assumed to be correct, then the newsvendor takes his/her decision based on the random variable D_C , which is the sum of the random variable representing the estimated demand and the random variable representing the signal ($D_C = D_F + D_S$). In this case, by adjusting the order quantity, the decision maker realizes an additional profit equal to:

$$\pi_C = \begin{cases} (M-C)A - S(x - Q_0 - A) & \text{if } x - Q_0 \geq A \\ M(x - Q_0) + V(Q_0 + A - x) - CA & \text{if } 0 \leq x - Q_0 < A \\ (V - C)A & \text{if } x - Q_0 < 0 \end{cases} \quad (\text{A1})$$

In contrast, if the signal is wrong, then the newsvendor realizes the additional profit π_W , which is π_C subject to the random variable D_W which is equal to the random variable characterizing the estimate demand distribution ($D_W = D_F$). The total profit of the newsvendor can now be formulated as:

$$\pi_{adj} = \pi_F + p\pi_C + (1 - p)\pi_W \quad (\text{A2})$$

The expected value of (A2) is calculated as

$$\begin{aligned}
E(\pi) &= (M + S - C) \left(\int_{Q_0}^{\infty} Q_0 f(x) dx + p \int_{Q_0+A}^{\infty} A f_C(x) dx + (1-p) \int_{Q_0+A}^{\infty} A f_W(x) dx \right) \\
&- S \left(\int_{Q_0}^{\infty} x f_W(x) dx + p \int_{Q_0+A}^{\infty} (x - Q_0) f_C(x) dx + (1-p) \int_{Q_0+A}^{\infty} (x - Q_0) f_W(x) dx \right) \\
&+ (M - V) \left(\int_0^{Q_0} x f_W(x) dx + p \int_{Q_0}^{Q_0+A} (x - Q_0) f_C(x) dx + (1-p) \int_{Q_0}^{Q_0+A} (x - Q_0) f_W(x) dx \right) \\
&+ (V - C) \left(\int_0^{Q_0} Q_0 f_W(x) dx + p \int_0^{Q_0+A} A f_C(x) dx + (1-p) \int_0^{Q_0+A} A f_W(x) dx \right) \quad (A3)
\end{aligned}$$

We can easily show that the profit function is concave and the optimal adjustment quantity could be derived from the first derivative condition:

$$\begin{aligned}
\frac{\partial E(p)}{\partial A} &= (C - V) p F_C(0) + M + S - C - (V - C) (1 - p) F_W(0) \\
&- (M + S - V) ((1 - p) F_W(Q + A) + p F_C(Q + A)) = 0 \quad (A4)
\end{aligned}$$

Part B: Likelihood function optimization

For a given data set $D = \{d_n\}$ $n = 1 \dots t$ associated with the observed demand realizations up to period t , the likelihood function of the GMM is given by:

$$P[D \setminus p] = \prod_{n=1}^t [p f_C(d_n) + (1-p) f_W(d_n)] \quad (B1)$$

The log likelihood function is derived as:

$$\ln[P[D \setminus p]] = \sum_{n=1}^t \ln[p f_C(d_n) + (1-p) f_W(d_n)] \quad (B2)$$

By applying the first derivative condition, the \hat{p}_t^{demand} estimator derived from the maximisation of the log likelihood function should consequently verify:

$$\sum_{n=1}^t \frac{f_C(d_n) - f_W(d_n)}{\hat{p}_t^{demand} f_C(d_n) + (1 - \hat{p}_t^{demand}) f_W(d_n)} = 0 \quad (B3)$$

Without the need of an advanced calculation software, \hat{p}_t^{demand} could be found by using the Solver add-in of Microsoft Excel.

Part C: Impact of the adjustment on economic performance

As illustrated in Figures (6a) and (6b) of the paper, the profit difference between the 100% *trust* and the 100% *ignore* policy might be important. Such a difference depends on the signal distribution moments, its probability of being correct, its position compared to the estimated demand as well as the inventory unit costs. For a given actual signal probability p estimated by \hat{p} by the decision maker, we could measure the impact of the judgement-enabled inventory policy by its relative benefit compared with the *ignore* and *trust* policies:

$$benefit_0(p, \hat{p}) = \frac{\pi_p(Q_{\hat{p}}^*) - \pi_p(Q_0^*)}{\pi_p(Q_0^*)} \quad (C1)$$

$$benefit_1(p, \hat{p}) = \frac{\pi_p(Q_{\hat{p}}^*) - \pi_p(Q_1^*)}{\pi_p(Q_1^*)} \quad (C2)$$

For perfect knowledge about p , i.e. $\hat{p} = p$, Figure C1 illustrates the relative adjustment benefit compared with the *trust* and *ignore* policies.

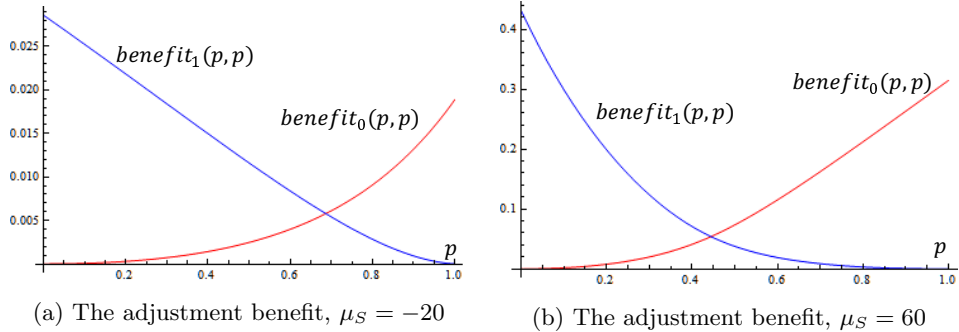


Figure C1: The impact of the adjustment policy

When we estimate p , the adjustment benefit is lower since the decisions made by the newsvendor are suboptimal. For instance, if $\mu_S = 60$ and $p = 0.8$, Figure C2 contrasts the benefit resulting from perfect information and an estimated value of p in the interval $[p_{trust-ignore}, 1]$:

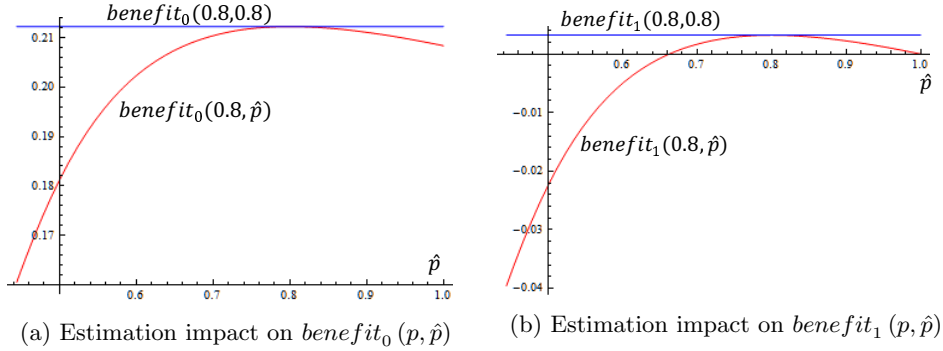


Figure C2: Perfect information versus estimation

It is worthwhile noting that adjusting is always beneficial compared with 100% ignoring the signal. However, adjustment could be harmful compared to a 100% trust if the estimated p is lower than \hat{p}_{adj} .

Part D: A note on the simulation spreadsheet

The Excel file simulation accompanying this paper is composed of six worksheets which are described in detail in this simulation user manual. The “Description” worksheet introduces the simulation file and defines all columns used during the simulation run. To run the simulation, the user first needs to complete the “SimulationInputs” worksheet with two types of data:

1. Data related to the simulation itself:
 - (a) The total number of periods to run;
 - (b) The average of the forecast demand (minimum and maximum values);
 - (c) The standard deviation of the forecast demand (minimum and maximum values);
 - (d) The average of the signal (minimum and maximum values);
 - (e) The standard deviation of the signal (minimum and maximum values);
 - (f) The actual value of the probability p (minimum and maximum values);
 - (g) The number of periods enabling the initialization of the learning process: during these periods, a random value of p is used to derive the order quantity, which allows the decision maker to collect initial observations;
 - (h) The unit cost values: purchase cost, selling price, salvage cost, shortage cost and the associated underage and overage unit costs.

The main results provided in the paper (with the exception of those in Section 5.2) are associated with a deterministic assumption with regard to the forecast demand and signal moments as well as the probability p . By allowing the reader to enter the minimum and maximum values of these

parameters in the simulation file, we permit him/her to test the quality of the learning process under a uniform distribution of these parameters (presented in the paper in Section 5.2).

2. Data related to the link between the simulation file and a Mathematica calculation server:
 - (a) If Mathematica is locally installed on the users machine, the Mathematica mode should be set to “Local” and the user needs to enter the path of the “math.exe” file on his/her machine.
 - (b) If the user does not have Mathematica installed on his/her machine, the Mathematica mode should be set to “Remote”. In such a case, the user needs to enter the parameters of a remote server where Mathematica is installed.
 - i. The IP address of the Mathematica server;
 - ii. The user name of the Mathematica server;
 - iii. The user password or the key file permitting to connect to the Mathematica server;
 - iv. The path on the remote server where the calculation files are transferred.

In the case where the user does not have either a local or a remote Mathematica server, he/she could contact the first author of the paper, who could provide the parameters of his Mathematica server for the purpose of testing the simulation experiment. We note that the calculation time for each period under the “Local” Mathematica mode is almost instantaneous (around 1 second), whereas it takes much more time to execute the simulation under the “Remote” mode. In the latter case, each calculation query is firstly written in a file, then transferred to the server, then executed remotely on the server with its results written in a file, which is then downloaded to the local machine and finally incorporated in the Excel file. The upload and the download tasks increase the time of calculation considerably. Besides, during one period calculations, some calculation queries cannot be launched before having the results of previous queries. That is, the calculation queries are not launched in a batch, but they are performed one by one which increases the number of upload and download tasks. For this reason, and as noted above, the calculation time under the “Remote” Mathematica mode is considerably higher than under the “Local” mode. Once the Mathematica mode and its associated parameters have been entered, the user needs to validate the link between the Excel file and the Mathematica calculation server. When validated, the user can launch the simulation: one run or multiple runs. The main simulation worksheet, named “Simulation_Running”, illustrates the evolution of the learning process in one simulation run. The user can graphically follow the learning process in the worksheet “Simulation_Chart”. Figures 7 and 8 of the paper are derived from this latter worksheet. The results of each run are summarized in the worksheet “Simulation_Results”, and when running multiple runs, the average of these runs are presented in the worksheet “Simulation_Runs_Summary”. The “Simulation_Running” worksheet illustrates the learning process presented in the paper. All columns of this worksheets are

defined in the “Description” worksheet. Particularly, the outcome of the learning process is presented in columns AI (the p resulting from the first learning technique), AL (the p resulting from the second learning technique), and AQ (the p resulting from the mix between the first and the second learning techniques).

Technical Note: Notice that the simulation file was developed and tested on a 64 bit machine running Microsoft Windows 10 and Microsoft Office 2016 both of them under the 64 bit architecture. In the case of a remote Mathematica link, the connection with the calculation server is established by using the ssh protocol and its associated port (generally the port 22). The user needs to be sure that his/her internet provider (particularly at his/her organisation) allows the use of this port. Finally, in order to coordinate the decimal representation between Mathematica and Microsoft Excel, it is important to configure the decimal parameters under the Microsoft Windows configuration panel with a point “.”. For some international versions of Windows, French for instance, this parameter could be configured with a comma “,” which creates a communication problem with Mathematica. The latter is configured to represent decimal values by a point.