Newsvendor decisions under supply uncertainty

Anssi Käki[*], Juuso Liesiö, Ahti Salo, and Srinivas Talluri

September 25, 2013

Abstract

We analyze the impact of supply uncertainty on newsvendor decisions. First, we derive a solution for a newsvendor facing stochastic supply yield, in addition to stochastic demand. While earlier research has considered independent uncertainties, we derive the optimal order quantity for interdependent demand and supply and provide a closed-form solution for a specific copula based dependence structure. This allows us to give insights into how dependence impacts the newsvendor's decision, profit, and risk level. In addition to theory, we present experimental results that show how difficult newsvendor decisions under supply uncertainty are for human subjects. In our experiment, the control group replicated a well-known newsvendor experiment whereas the test group faced additional supply yield uncertainty. Comparison of these results shows that under low-profit condition, subjects are able to incorporate supply uncertainty quite well in their decisions. Under high-profit condition, the deviation from the optimum is much more significant. We discuss this asymmetry and also propose some ways to improve newsvendor decision making.

[*]Corresponding author (anssi.kaki@aalto.fi)

1 Introduction

The newsvendor problem characterizes situations where a costly commitment (e.g., procurement order, production order, capacity plan) must be made before a realization of unknown demand occurs. Both excess commitments and shortages are costly. A well-known result (Porteus 1990) is that the newsvendor should select her commitment level by matching a so called critical fractile to the uncertain demand, which means ordering more (less) than mean demand when the profit margin is high (low). Besides the traditional newsvendor setting, various extensions of the problem have been studied: Petruzzi and Dada (1999) analyze pricing decisions in the newsvendor setting, Cachon and Lariviere (2005) study supply chain coordination with revenue sharing contracts using the newsvendor model, and Dada et al. (2007) study a newsvendor who procures from multiple suppliers, of whom some are unreliable. These, and many other studies serve as a foundation for various inventory policies, pricing rules, and supplier selection methods.

Even if the newsvendor problem has a simple structure, experimental studies (e.g., Schweitzer and Cachon 2000; Bolton and Katok 2008; Bostian et al. 2008) have shown that human decision makers systematically make non-optimal decisions. Most typically, subjects anchor their order quantities to the expected demand in cases where maximal profits could be achieved by ordering much more than the expected demand (if product has a high profit margin, i.e., shortage is relatively more costly than excess) or less than the expected demand (if product has a low profit margin). This behavior is known as the *pull-to-center* effect. There is experimental evidence it is (at least partially) context-specific, and that with lottery type framing the decision biases are less significant (Kremer et al. 2010). Various explanations for this have been proposed, but conclusive explanations have not been provided (Katok 2011).

In this paper, we study how uncertain supply affects newsvendor decision making. We first examine the theoretical implications of stochastic supply yield on the newsvendor model under both independent and *interdependent* demand and supply uncertainties. We find that uncertain supply, even if it has a simple structure, has a non-trivial impact on the optimal order quantity.

We then present an experiment which we carried out to find out whether earlier behavioral results of the newsvendor problem hold under uncertain supply. The results show that the pull-to-center effect becomes more pronounced (orders are further away from the optimal) when the profit margin is high, but the deviation from the optimum is only moderate when the profit margin is low. We propose several possible explanations for this outcome, and conclude with some managerial guidelines for newsvendor decision making under uncertain supply.

2 Literature review

We structure our review of earlier research into theoretical papers that consider newsvendor facing uncertain supply and experimental papers in the newsvendor setting. For comprehensive reviews of the newsvendor problem, we refer to Porteus (1990); Khouja (1999); and Qin et al. (2011). We present the basic newsvendor formulation in the beginning of Section 3.

2.1 Newsvendor and uncertain supply

Qin et al. (2011) call for more studies of the newsvendor problem under both stochastic demand and stochastic supply. In general, uncertain supply –either related to procurement (order) or production (production batch order)– has been discussed at length in the inventory management literature. Supply uncertainty can be divided into two types: disruptional and operational. Supply disruptions (supplier is either in operation or not) have recently received more attention; see Tomlin (2006) and references therein. We focus on operational supply uncertainty in the form of stochastic yield. Yano and Lee (1995) review random yield models, and present five basic approaches: i) a Bernoulli process; ii) stochastically proportional yield (our model); iii) stochastic yield proportional to order quantity; iv) random capacity; and v) general model that specifies the probability of each output for each order quantity. Stochastic yield is widely in use as indicated by reviews Henig and Gerchak (1990) and Gerchak and Grosfeld-Nir (2004).

Stochastic yield has been considered in the newsvendor setting for both single and multi supplier case. For example, Keren (2009) presents how stochastic yield impacts supply chain coordination. Yang et al. (2007) solve the newsvendor problem under multiple suppliers (with stochastic yield) and multiple products. They find that the newsvendor optimum is a function of both supplier cost and reliability. Burke et al. (2009) provide analytical results and conclude that cost is the key order qualifier when the newsvendor can select from multiple suppliers. They also discuss how demand uncertainty impacts the newsvendor decision in this setting; the key finding is that under low mean demand, single supplier is preferred whereas under high demand, multiple suppliers strategy is preferred. The paper closest to our approach is Inderfurth (2004), who presents a closed-form solution of a single product single supplier newsvendor problem under uniform demand and uniform supply yield. We extend this to interdependent demand and supply uncertainty, which to our knowledge has not been studied before.

We also consider risk-averse newsvendor decisions. Keren and Pliskin (2006) present how risk aversion impacts a newsvendor who faces uniform demand. Among other things, they find that increasing demand variability can both increase or decrease the optimal order quantity. Chen et al. (2009) derive a generic solution to newsvendor problem with random demand when both pricing and ordering are decisions are taken by minimizing the Conditional-Value-at-Risk (CVaR) measure. The results show that the risk-averse newsvendor behaves differentially from her risk-neutral counterpart. CVaR is widely used because it accounts for both reward and risk and is thus attractive from decision maker's viewpoint. It is also coherent risk measure and has attractive technical characteristics (Chen et al. 2009, Rockafellar and Uryasev 2000).

2.2 Newsvendor behavior in experiments

Our newsvendor experiment is a single factor design where supply is either perfect (baseline) or random (treatment). The perfect supply case replicates the seminal laboratory experiment by Schweitzer and Cachon (2000) (S&C). To our knowledge, no other experiment has explicitly accounted for random supply; Gavirneni and Isen 2010 mention imperfect yield in their setup, but it is deterministic and not in the focus of their experiment.

S&C consider a classic newsvendor setup with random (uniform) demand and simple pricing where only order cost and selling price are non-zero. In their experiment, S&C compare two products: one with high profit margin and one with low profit margin. They find that the average orders made by human subjects are located between the average demand and the optimal order quantity. S&C have a mean demand of 150 and the optima for the two profit conditions are 75 and 225; the subjects ordered on average 134 and 177. This pull-to-center bias has been observed in many later experiments (e.g., Benzion et al. 2008; Bostian et al. 2008; Lurie and Swaminathan 2009; Bolton et al. 2012).

S&C tested a number of explanations for the pull-to-center behavior and suggested that subjects try to minimize *ex post* inventory error, i.e., the deviation between order and realized demand. This leads the subjects to pull their orders towards the mean demand, or anchor to the mean. S&C also found some evidence of demand-chasing in which the order decision is adjusted to the direction of the previous round's demand.

Risk aversion, in general, is a plausible explanation whenever the subjects' behavior does not appear risk neutral. While risk aversion does not explain newsvendor's pull-to-center bias in the low-profit condition, De Véricourt et al. (2013) argue that it could explain the results in the high-profit newsvendor setting. They find that males make higher (more risky) orders than females, which can be explained due to differences in risk attitude. S&C also considered prospect theory (Kahneman and Tversky 1979) as one candidate explanation for their results, but found that it was inadequate by testing a case in which strict losses were not possible.

Bostian et al. (2008) repeated the S&C experiment with more decision rounds, and found that subjects can learn during the experiment and become less responsive to demand (as is rational) as they gain experience. Bolton and Katok (2008) also found that experience improves performance, in particular under the high profit condition. Also Benzion et al. (2008) reported similar results. But more information does not necessarily lead to better results. The experiment by Lurie and Swaminathan (2009) shows that frequent feedback (seeing detailed results of every round) can actually detoriate performance. They also found that subjects, when given the choice to access any information they want, spent most of their time analyzing the demand of the previous round when making order decisions. This leads to demand-chasing, which is not optimal in the newsvendor case. Lurie and Swaminathan (2009) found little, if any evidence of learning during the rounds.

In addition to De Véricourt et al. (2013) (who compared males and females), also Bolton et al. (2012) use heterogeneous subject pools: they test whether experienced procurement managers perform better than students. They conclude that there are no differences between these groups. In other words, managerial experience does not improve decision making in the newsvendor setting. Lau et al. (2013) analyze two newsvendor experiments and find that less than half of subjects suffer from pull-to-center bias. They conclude that there can be significant differences between the subjects within one experiment and urge researchers not to make too strong conclusions from averages across experiments; similar conclusions can be drawn from Gavirneni and Isen (2010).

In conclusion, it is not clear exactly individuals behave in the newsvendor setting or why biases occur. Su (2008) proposes a bounded rationality model to explain non-optimal decisions and achieves a rather good fit with the experimental data. However, Kremer et al. (2010) conduct a newsvendor experiment in a neutral (lottery type) setting and in operations (ordering under uncertain demand) setting. The results between these two experiments differ significantly, indicating that bounded rationality alone does not explain the non-optimal behavior in the original newsvendor context. As noted by Katok (2011), "a full explanation for the newsvendor behavior is elusive and it seems likely that there is no single explanation for the observed behavior."

3 Newsvendor model with uncertain demand and supply

We assume a profit-maximizing newsvendor who faces random demand D with mean μ_D and support $[0,\overline{D}], \overline{D} \in \mathbb{R}^+$. Supply is also uncertain: if q products are ordered, $Z \cdot q$ orders are received, where Z is a random variable with expectation μ_Z and support $Z \subseteq [0,1]$. Essentially, this corresponds to a stochastic (multiplicative) yield model, which can have several interpretations: for example, the product in question can be technologically challenging and Z can describe the share of products that are of adequate quality; or the supplier can have resource constraints, and Z represents the share of order that the customer receives. The joint distribution function of demand and supply yield is denoted with $f_{ZD}(z,d) = f(z,d)$

The product's selling price is p, unit order cost per received product is $c < p$, and the salvage value is $s < c$. For order q, $\min\{D, Z \cdot q\}$ units are sold and $\max\{Z \cdot q - D, 0\}$ units are salvaged. Thus, the expected profit is

$$\begin{aligned}
\mathbb{E}[\Pi(q)] &= \mathbb{E}[p \cdot \min(D, Z \cdot q) + s \cdot \max(Z \cdot q - D, 0) - c \cdot Z \cdot q] \\
&= p \cdot \mathbb{E}[D - \max(D - Z \cdot q, 0)] + s \cdot \mathbb{E}[\max(Z \cdot q - D, 0)] - c \cdot \mathbb{E}[Z \cdot q] \\
&= p \cdot \mathbb{E}[D] - \frac{p}{2} \cdot \mathbb{E}[|D - Z \cdot q| + D - Z \cdot q] + \frac{s}{2} \cdot \mathbb{E}[|Z \cdot q - D| + Z \cdot q - D] - c \cdot \mathbb{E}[Z \cdot q] \\
&= \frac{1}{2}\Big(\mu_D(p-s) + \mu_Z q(p+s-2c) - (p-s)\mathbb{E}[|Z \cdot q - D|]\Big). \qquad (3.1)
\end{aligned}$$

This is a concave function and the optimum $q^* = \arg\max_{q \geq 0} \mathbb{E}[\Pi(q)]$ can be found by examining the first order derivative

$$\pi(q) := \frac{\partial \mathbb{E}[\Pi(q)]}{\partial q} = (s-c)\mu_Z + (p-s)\int_{\underline{Z}}^{\min\{1,\overline{D}/q\}} z \int_{z \cdot q}^{\overline{D}} f(z,d)\mathrm{d}d\mathrm{d}z. \qquad (3.2)$$

For details of he calculation, see the Appendix. With deterministic supply, i.e. $\underline{Z} = 1$, (3.2) becomes

$$\pi(q) = (s-c) + (p-s)\int_{1}^{\min\{1,\overline{D}/q\}} z(1 - F_D(q))\mathrm{d}z$$

and, since $\pi(\overline{D}) = s - c < 0$ and $\pi(q)$ is a monotonically decreasing function, the optimal order $q^* < \overline{D}$ and $\min\{1, \overline{D}/q\} = 1$. Thus, solving $\pi(q^*) = 0$ yields

$$(s-c)+(p-s)(1-F_D(q^*)) = 0 \Rightarrow F_D(q^*) = \frac{p-c}{p-s}, \tag{3.3}$$

which is the solution of the basic newsvendor problem.

3.1 Independent uniformly distributed demand and supply

We consider first the case of independent demand and supply with uniform marginal distributions. The joint distribution function for this case is:

$$f(z,d) = \begin{cases} \left(\overline{D}(1-\underline{Z})\right)^{-1} & \text{if } d \in [0,\overline{D}], z \in [\underline{Z},1] \\ 0 & \text{otherwise} \end{cases} \tag{3.4}$$

and the expected yield $\mu_Z = (\underline{Z}+1)/2$. Inserting these into (3.2) gives

$$\pi(q) = \frac{(s-c)(\underline{Z}+1)}{2} + (p-s)\int_{\underline{Z}}^{\min\{1,\overline{D}/q\}} \frac{z}{1-\underline{Z}} \int_{z\cdot q}^{\overline{D}} \frac{1}{\overline{D}}\, \mathrm{d}d\mathrm{d}z$$

and evaluating this at $q = \overline{D}$ yields

$$\pi(\overline{D}) = \frac{3(s-c)(\underline{Z}+1)+(p-s)(1+\underline{Z}-2\underline{Z}^2)}{6} \tag{3.5}$$

Hence, if cost parameters p, c and s and the minimal supply yield $\underline{Z}$ are such that (3.5) is non-positive. Then, $q^* \leq \overline{D}$ and the optimal order is obtained from:

$$\begin{aligned} &\frac{(s-c)(\underline{Z}+1)}{2} + (p-s)\int_{\underline{Z}}^{1} \frac{z}{1-\underline{Z}} \int_{z\cdot q}^{\overline{D}} \frac{1}{\overline{D}}\, \mathrm{d}d\mathrm{d}z = 0 \\ \Rightarrow \quad & q^* = \frac{3}{2}\frac{(p-c)\overline{D}(1+\underline{Z})}{(p-s)(1+\underline{Z}+\underline{Z}^2)}. \end{aligned} \tag{3.6}$$

If, however, (3.5) is positive, then $q^* > \overline{D}$ and $\min(1,\overline{D}/q^*) = \overline{D}/q^*$. In this case, the optimality condition is

$$\begin{aligned} &\frac{(s-c)(\underline{Z}+1)}{2} + (p-s)\int_{\underline{Z}}^{\overline{D}/q^*} \frac{z}{1-\underline{Z}} \int_{z\cdot q}^{\overline{D}} \frac{1}{\overline{D}}\, \mathrm{d}d\mathrm{d}z = 0 \\ \Rightarrow \quad & \frac{(\overline{D}-q^*\underline{Z})^2(\overline{D}+2q^*\underline{Z})}{6\overline{D}q^{*2}(1-\underline{Z})} - \frac{(\underline{Z}+1)(c-s)}{2(p-s)} = 0. \end{aligned} \tag{3.7}$$

For a numerical example, we use the same parameters as Schweitzer and Cachon (2000): the selling price $p = 12$, and for a *high-profit product*, the order cost $c = 3$, and for a *low-profit product*, $c = 9$. Salvage value $s = 0$ in both cases. Demand follows the uniform distribution $[0, 300]$ and, if supply is deterministic, $q^* < \overline{D}$ and (3.6) gives $q^*_H = 225$ and $q^*_L = 75$.

Consider now uncertain supply and a high-profit product: with the given cost parameters, the derivative (3.5) is positive for $\underline{Z} < (1+\sqrt{33})/16 \approx 0.42$. Thus, (3.7) is the optimum if lower bound for supply is less than 0.42, otherwise the optimum is (3.6). The results for varying $\underline{Z}$ are in Figure 1(a). The figure shows that the optimal order in the fully random case ($\underline{Z} = 0$) is around 346, and decreases almost linearly towards 225, which is the optimal order under deterministic supply. The cost of supply uncertainty can be calculated with profit difference between the fully random supply ($\underline{Z}$=0) and deterministic ($Z = 1$); the latter is 33% higher. Under the low-profit condition, (3.5) is negative for all $\underline{Z} \in [0, 1]$ and (3.6) is the optimum. The results in Figure 1(b) are similar to the high-profit condition.

Note that the optimal order cannot be derived from the optimum under deterministic supply by inflating with the expected supply yield, i.e., using the following heuristic:

$$q^h = q^*/\mu_Z, \tag{3.8}$$

where q^* is the optimum from (3.3). For example, if one assumes $Z \sim U[0, 1]$, the newsvendor expects to receive μ_Z=0.5 of her order and the heuristic above suggests doubling the optimal orders, i.e., $q^h_H = 450$ and $q^h_L = 150$. However, as shown in Figure 1, where the heuristic is marked with dashed lines, the corresponding true optima are different ($q^*_H = 346$ and $q^*_L = 113$). Especially under high supply uncertainty (low $\underline{Z}$), the difference between the heuristic and the true optimum is large.

The poor performance of the heuristic cannot be attributed to the non-symmetric cost structure. If we again assume $Z \sim U[0, 1]$, $D \sim U[0, 300]$, and, in addition, $p = 2c$ ("mid-profit" product) and $s = 0$, the deterministic optimum is 150 and the heuristic suggests $q^h = 150/0.5 = 300$ although this is not optimal, but the optimal order is 225. Indeed, in case $p = 2c$ and $s = 0$, the only term in the expected profit (3.1) that is dependent on q is $-(p-s)\mathbb{E}[|Z \cdot q - D|]$. Thus,

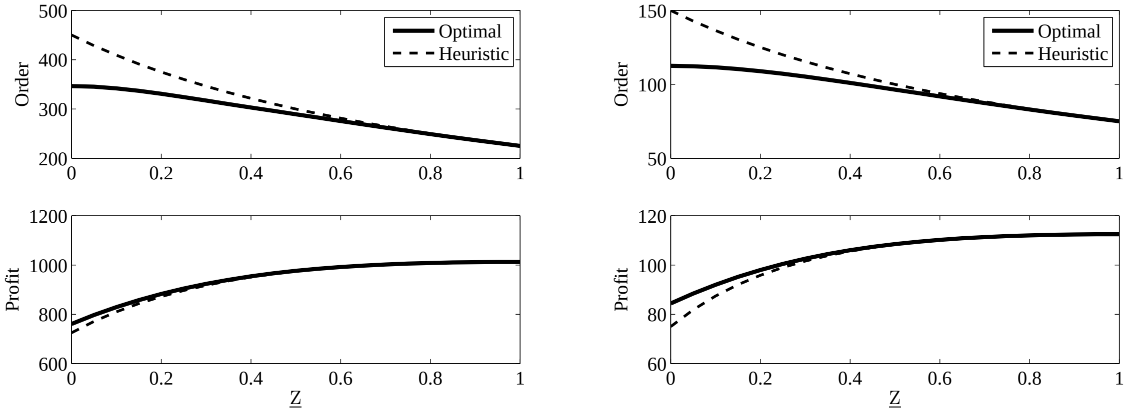

(a) High-profit product

(b) Low-profit product

Figure 1: The optimal order quantity and expected profit vs. heuristic where the deterministic supply optimum is divided by the expected yield. Supply uncertainty grows from right to left ($\underline{Z} : 1 \to 0$).

the order quantity q that minimizes $\mathbb{E}[|Z \cdot q - D|]$ maximizes profit as well. But in general, this absolute deviation is not minimized by setting q such that $\mathbb{E}[Zq] = \mathbb{E}[D]$. For example here, the expected deviation is minimized by $q^* = 225 \neq \mathbb{E}[D]/\mathbb{E}[Z]$.

In case $Z \sim [0.4, 1]$ (we use this in our experimental setup), the newsvendor's optimal orders are $q_H^* = 303$ and $q_L^* = 101$. With these order, as Figure 2 shows, the profit distributions are quite different for different profit conditions. For example, the coefficient of variation of profit in the high-profit case is $cv_H = \sigma_H/\mu_H = 835/955 \approx 0.87$ whereas in the low-profit case, $cv_L = 231/105 \approx 2.17$. Also skewness varies significantly: $Skew_H = -0.73 > Skew_L = -2.07$. Both cases have have roughly equal (17%) probability for making a loss. In qualitative terms, in the high-profit case, there is a lot of upside potential, but, In the low-profit case, the distribution is heavily skewed towards left, which also makes the variation relatively high.

3.2 Interdependent uniformly distributed demand and supply

The assumption of independent demand and supply is not always realistic. For example, consider a single supplier for N newsvendors. The supplier plans capacity based on orders $q_1, ..., q_N$.

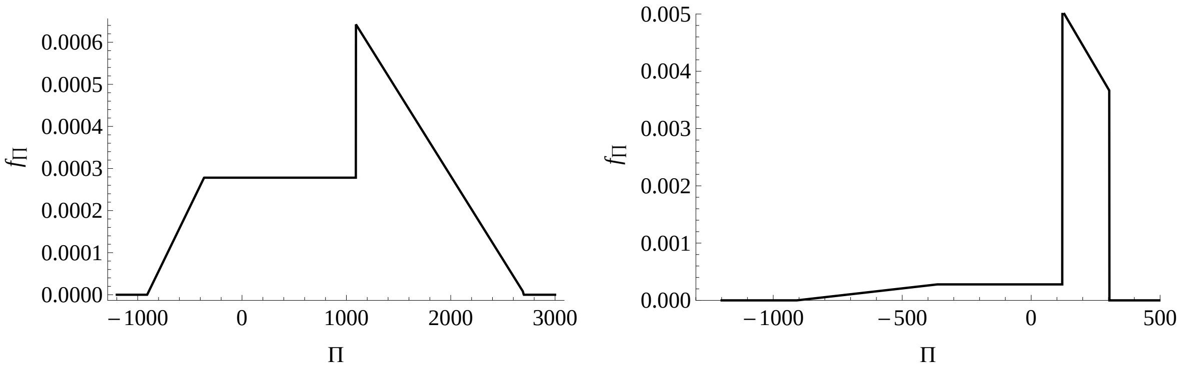

(a) High-profit product

(b) Low-profit product

Figure 2: Profit distributions for optimal orders $q_H^* = 303$ and $q_L^* = 101$ ($Z \sim [0.4, 1]$). See the Appendix for details of calculating the distributions.

Assume that our newsvendor is a non-prioritized customer, and that there also are prioritized newsvendors in the market. If the actual demand now exceeds expectations by a large margin, the supplier's capacity could become constrained because, e.g., the prioritized newsvendors can put in extra orders. Our newsvendor is likely to see large demand, too, but the supplier delivers only a portion of the order due to constrained capacity, implying *a low supply yield under high demand*. Then again, if the actual demand is disappointing, there will be ample capacity and a high yield is expected. Altogether, this implies *a negative demand-supply dependency.*

On the other hand, assume now that the supplier has flexible capacity. If the actual demand is way below the level indicated by initial orders, the supplier might cut its capacity significantly by, e.g., dedicating a production line or plant to a different product. This would benefit both the supplier and its customers, who are facing the disappointing demand and do not mind receiving less than what they asked for. In a similar manner, in case of high demand the supplier might ramp up new capacity to ensure that extra orders from prioritized newsvendors can be met. This kind of dynamic would imply *a positive demand-supply dependency.*

In general, solving the newsvendor problem analytically becomes intractable if the joint density function $f(z,d)$ is cumbersome, which is the case with most joint densities with dependency. One exception is the case of uniform margins and Farlie–Gumbel–Morgenstern (FGM) copula (Gumbel 1960; for summarized details, see Balakrishnan and Lai 2009). The copula (distribution) function for FGM is:

$$C(u,v) = uv\big(1+\theta(1-u)(1-v)\big), \tag{3.9}$$

where $\theta \in [-1,1]$ is the parameter for dependency strength and u and v follow the standard uniform distribution. If measured with linear correlation coefficient, the dependency strength of FGM varies between $[-\theta/3, \theta/3]$ (Balakrishnan and Lai 2009). Thus, the FGM copula can only be used in case of weak dependency. Differentiating (3.9) with respect to u and v gives the copula density function $c(u,v) = 1+\theta(1-2u)(1-2v)$ and the corresponding probability density function for uniform marginals (with $\underline{D}=0$ and $\overline{Z}=1$) becomes

$$\begin{aligned} f(z,d) &= c\big(d/\overline{D}, z/(1-\underline{Z})\big) \\ &= \frac{1+\theta}{\overline{D}(1-\underline{Z})} - \frac{2d\theta}{\overline{D}^2(1-\underline{Z})} - \frac{2\theta(z-\underline{Z})}{\overline{D}(1-\underline{Z})^2} + \frac{4d\theta(z-\underline{Z})}{\overline{D}^2(1-\underline{Z})^2}. \end{aligned} \tag{3.10}$$

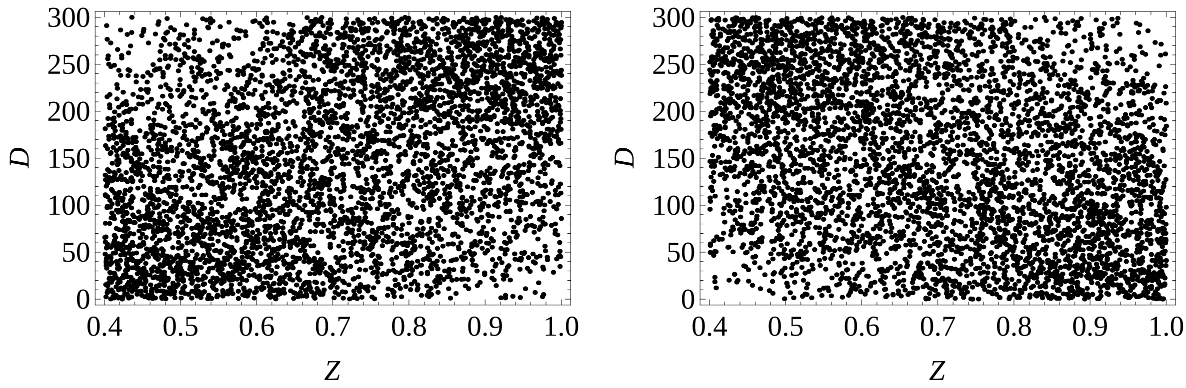

Figure 3: Samples from FGM copula with uniform margins and $\theta = 1$ (left) and $\theta = -1$ (right).

Using the same parameters as before and the FGM copula, the joint density becomes $f(z,d) = \big[450+\theta(d-150)(10z-7)\big]/81000$. Samples of this density are illustrated in Figure 3 for $\theta \pm 1$. The optimal orders can be calculated using the first order derivative (3.2) and the expected profit with (3.1). The results are in Figure 4 and the optimal orders in Figure 5. Note that with the given parameters, θ increases the expected profit for all $0 < q \leq 750$, i.e., the more positive the dependency, the more profit the newsvendor is expected to make (see the Appendix for details).

The results show that interdependency has a clear impact: under the high-profit condition, the optimal expected profit is 7% more in the case of positive dependency, compared to negative dependency. In the low-profit case, the expected profit is 50% higher in the positive dependency case. Also the optimal newsvendor decision is affected: in the positive dependency case, the optimal order is 6% lower compared to the negative case; with low-profit it is 40% larger. We also note that the results are based on the FGM copula, which attains maximum 1/3 dependency strength (in terms of linear correlation). With stronger dependency, these differences would be emphasized.

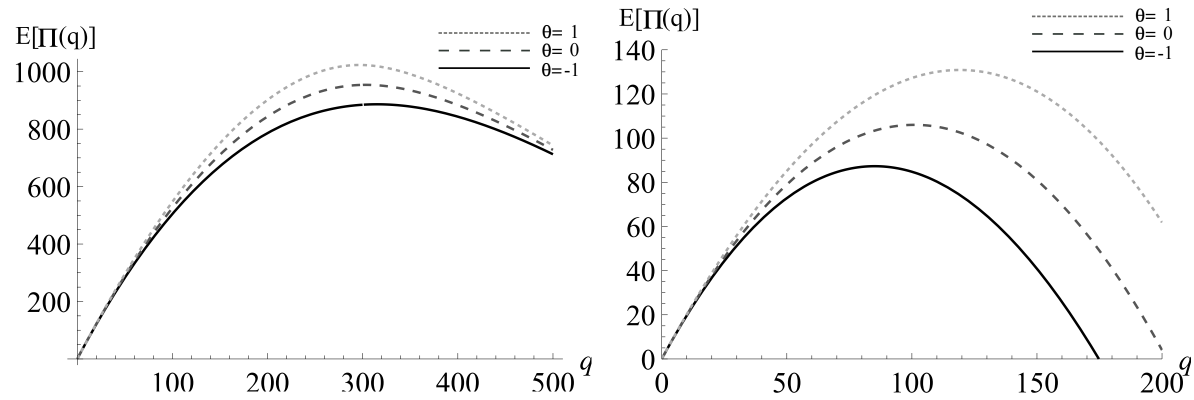

(a) High-profit product

(b) Low-profit product

Figure 4: Expected profit for varying q with negative dependency ($\theta = -1$), independency ($\theta = 0$) and positive dependency ($\theta = 1$).

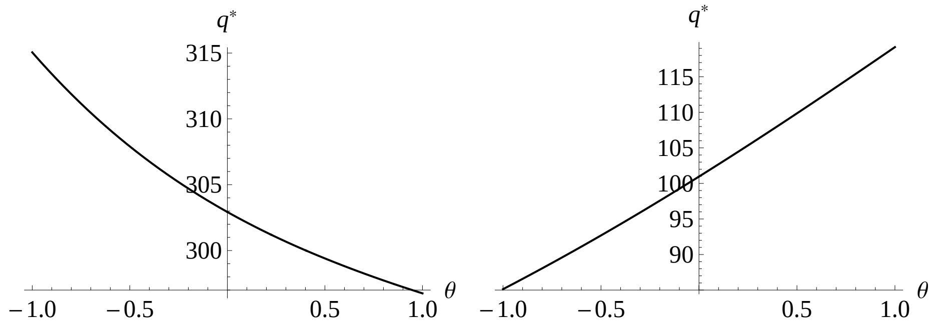

(a) High-profit product

(b) Low-profit product

Figure 5: Optimal order quantities under varying dependency.

3.3 Risk-averse newsvendor

In newsvendor applications, products typically have short selling period and high demand uncertainty. In this kind of setting, it is natural to consider also criteria other than profit maximization. We analyze Conditional-Value-at-Risk, which is the mean profit among the worst (1-β)% profits, denoted with $CVaR_\beta$. Calculation of $CVaR_\beta$ can be formulated as an optimization problem, where α denotes Value-at-Risk, i.e., the treshold value at worst (1-β) of profits (Rockafellar and Uryasev 2000):

$$\begin{aligned} CVaR_\beta(q) &= \max_{\alpha \in \mathrm{R}} \Big\{ \alpha + \frac{1}{1-\beta} \mathrm{E}[\min\{\Pi(q)-\alpha, 0\}] \Big\} \\ &= \max_{\alpha \in \mathrm{R}} \Big\{ \alpha + \frac{1}{1-\beta} \int_{\underline{Z}}^{1} \int_0^{\overline{D}} \min\{\Pi(q)-\alpha, 0\} f(z,d) \mathrm{d}d\mathrm{d}z \Big\}. \end{aligned} \tag{3.11}$$

In our case, the profit function $\Pi(q) = \frac{1}{2}\Big(d(p-s) + zq(p+s-2c) - (p-s)|zq-d|\Big)$.

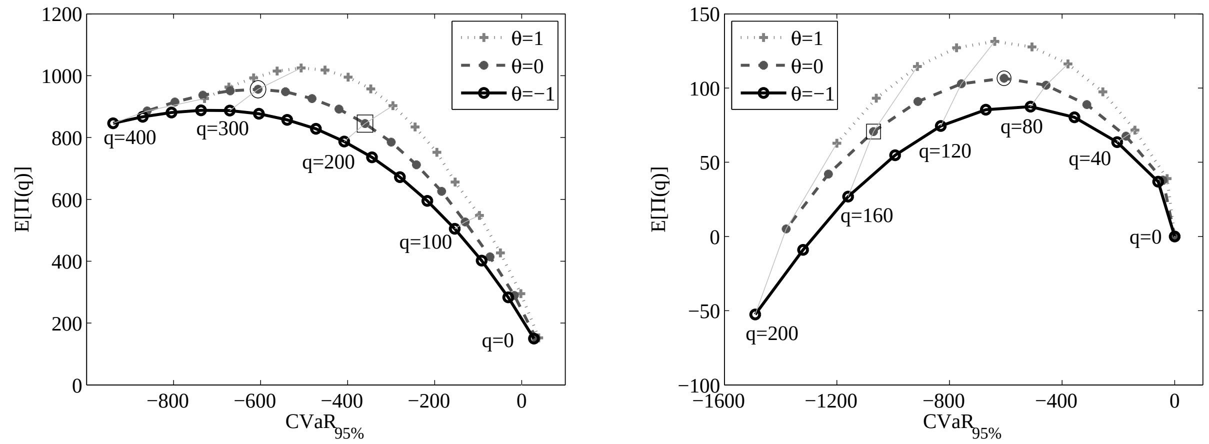

(a) High-profit

(b) Low-profit

Figure 6: Risk (CVaR) and expected profit tradeoffs for varying q under negative dependency ($\theta = -1$), independency ($\theta = 0$) and positive dependency ($\theta = 1$). The circled and squared points are discussed in experimental results section.

CVaR minimization can result in trivial solutions: for example, a risk averse newsvendor solution in many cases is $q = 0$ which guarantees zero-losses. Instead of minimizing CVaR, the newsvendor can calculate the expected profit (3.1) and CVaR (3.11) for varying q to assess the tradeoff between profit and risk.

Figure 6 illustrates the risk-profit tradeoffs for uniform marginals and the FGM copula with negative dependency, positive dependency, and the independent case. The numerical parameters were the same as in the earlier examples, and CVaR was calculated for 95% level. Two observations can be made: First, the expected profit is flat around the optimum, and, under both profit conditions, the newsvendor can achieve close-to-optimal profit with a significantly lower risk level by *ordering less*. Second, positive dependency structure implies less risks and more profit in both cases.

4 Newsvendor experiment

Previous experiments with human subjects have shown that newsvendor decisions are difficult for decision makers. In particular, Schweitzer and Cachon (2000) observed a pull-to-center bias: under a low-profit condition, when excess orders are costly, subjects place their orders below mean demand but the adjustment is too small – and vice versa under high-profit. S&C identified anchoring to the mean demand as the best explanation for the non-optimal behavior.

In our experiment, we sought to explore *whether pull-to-center behavior applies also under uncertain supply*. After testing the experimental setup with eight test subjects, it became apparent that independent supply uncertainty alone is difficult to account for. Thus, we did not consider dependencies in the experiment. As shown in the previous section, adding supply uncertainty makes it more difficult to calculate the optimal order. Because humans have limited information processing capacity, one would expect even more severe departure from the optimal order compared to the simpler setup of S&C (Katok 2011). However, it is also possible that the introduction of supply uncertainty does not make things worse: in case anchoring on the mean demand holds, the low-profit case should not be problematic because the optimal order is still below mean demand. Under high-profit condition, though, anchoring on the mean would cause even more severe deviation from the optimum.

4.1 Experimental setup

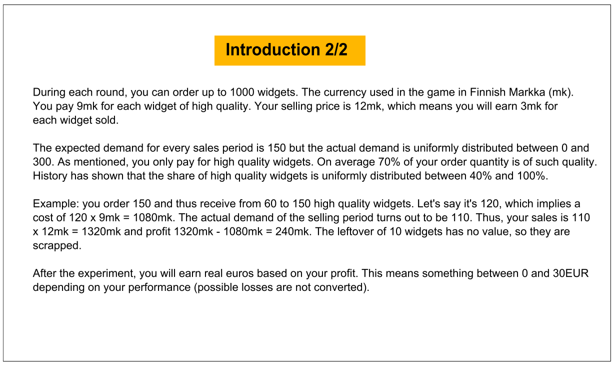

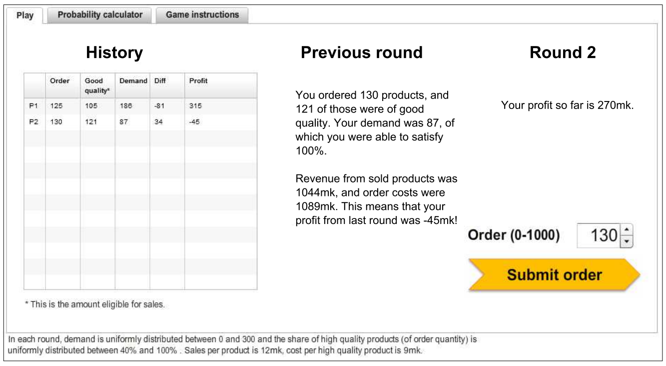

Figure 7: Initial instructions and the main screen.

The experiment was conducted in two computer classes supervised by a facilitator. Communications was restricted, but the subjects were allowed (but not instructed) to use pen and paper for calculations. The experiment was labeled as "Widget selling game" and implemented with on-line gaming tool Forio Simulate (screenshot in Figure 7). In the beginning, the subjects were introduced to the game and a simple numerical newsvendor example was presented. These instructions were available to the subjects during the game, and, in addition, they had access to a probability calculator to test how any order quantity translates into: i) the probability of demand exceeding the order; ii) the probability of shortage; and iii) the probability of making profit. The same random sample of demand and supply (Figure 8) was used for all subjects.

After the game was over, we asked the subjects to explain their ordering strategy. We also invited feedback related to the experiment, which revealed that most subjects enjoyed playing the game and that there were no difficulties with the instructions or user interface. In two out of four experiments, the subjects were participants in a course on risk analysis and they received one extra point for the course from participation. Also, cash based on performance was offered as incentive: the students knew that they can get 0-30 EUR based on how they play. Eventually, on average 12.50 EUR was paid to each subject. The subjects are described in more detail in the following sections.

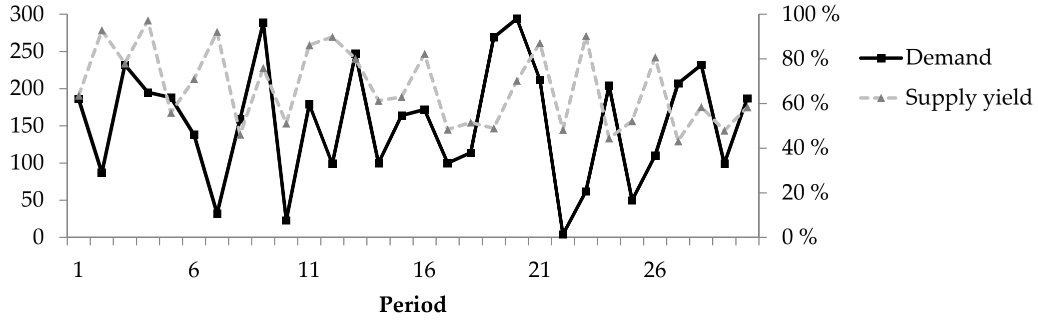

Figure 8: Random samples of demand and supply yield. Mean demand is 154, mean supply is 68%, and the series are uncorrelated.

4.2 Group *control*

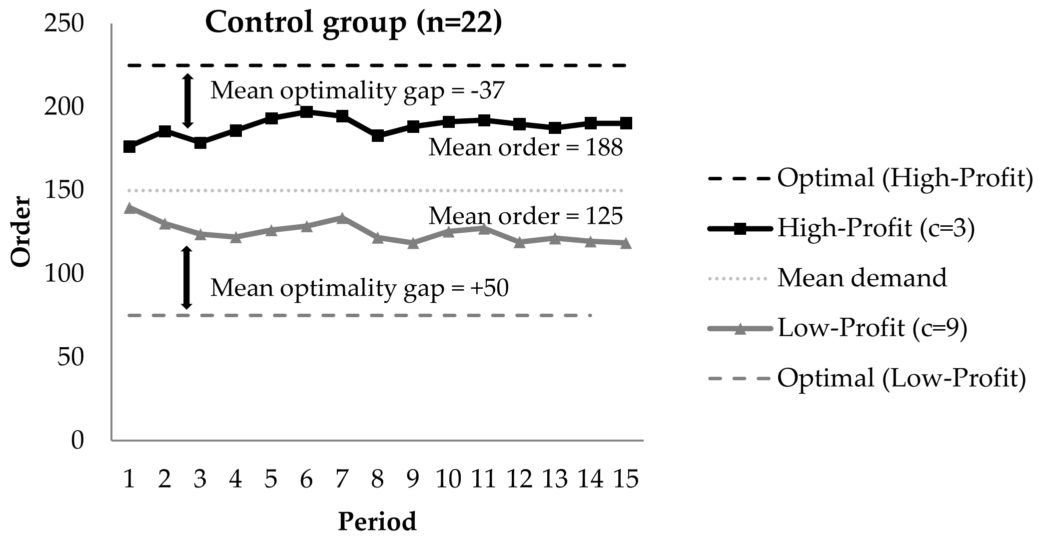

Figure 9: Results of group *control*.

The first group, labeled as *control*, was recruited from course on risk analysis offered by the Department of Mathematics and Systems Analysis at Aalto University. They were B.Sc. and M.Sc. students from varying fields, including engineering mathematics, industrial management, electrical engineering, and civil engineering. Altogether, 22 students were randomized to the *control* group.

The group replicated the S&C experiment with same parameters, i.e, $D \sim U[0,300]$, $p = 12$, $s = 0$, and $c = 3$ (high-profit) and $c = 9$ (low-profit). These imply optimal order quantities of

$q_H^* = 225$ and $q_L^* = 75$. During the 30-round experiment, randomly selected subjects (p=0.5) encountered the high-profit condition for the first 15 rounds and low-profit for the last 15 rounds (and vice versa with p=0.5), and the instructions and information available during the game were replicated from S&C as accurately as possible.

The results of this, shown in Figure 9, are very similar to those of S&C. The average order under high-profit condition is 188 (in S&C experiment it was 177), and under low-profit condition it is 125 (134 in S&C). We also observe similar pull-to-center asymmetry as S&C (and, e.g., Bostian et al. 2008): the optimality gap is smaller under the high-profit condition.

4.3 Group *test*

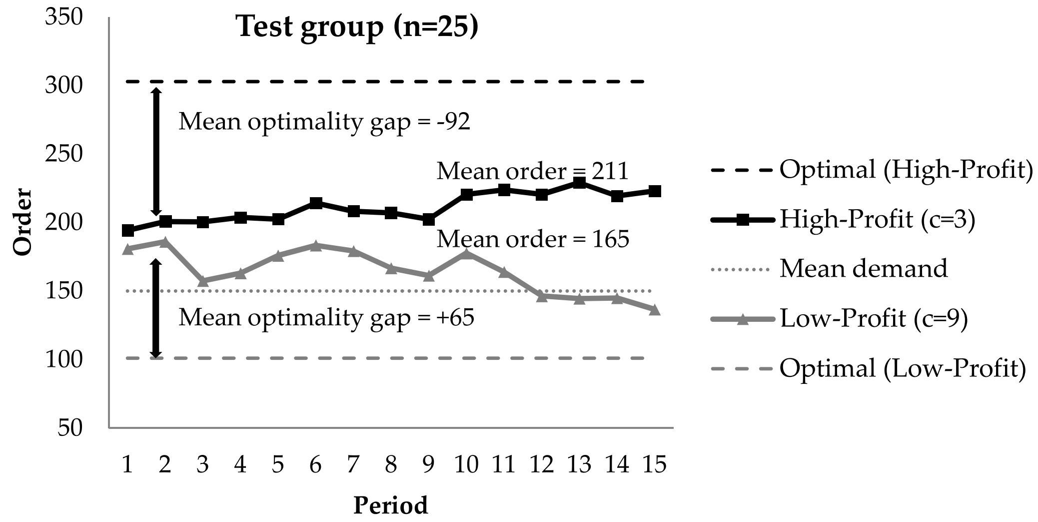

Figure 10: Results of group *test*.

The *test* group ($n = 25$) was recruited from the same course as the *control* group. The subjects played an otherwise identical game, but they were also told to expect "quality problems" in 0% to 60% their orders (i.e., share of usable products was 40% to 100%). In this setting, the optima are $q_H^* = 303$ and $q_L^* = 101$. The incentive structure was identical to the *control* group, and because supply uncertainty decreases the profit, the experiment fees were normalized to be at the same level among these two groups.

The results of the *test* group (Figure 10) show that the ordering behavior is more erratic than in the *control* group. Accounting for uncertain supply seems difficult as can be observed from the first round decisions: in the high-profit case, the average order is 194 vs. 180 of the low-profit case – this very small difference of 14 units is very small compared to difference of 202 units between the real optima. However, a learning impact can be observed: fitting a linear regression model to the data shows that in the high-profit case, orders increase towards the optimum by 2.1 per period (p=0.00) and in the low-profit case, orders decrease towards the optimum by 2.7 per period (p=0.00). In the *control* group, similar learning impact is either insignificant (p=0.09 under high-profit), or rather small (-0.9, p=0.01, under low-profit).

In the *test* group, subjects adjusted their order quantity between the rounds more often than not: only 42% of *test* orders were based on repeat choice whereas in the *control* group, 61% were repeated orders (which is in line with S&C; 64%). The amount of adjustments per subject varied from 2 to 25. If we classify strategies to static (6 adjustments or less out of 28 not-first-rounds) and dynamic (more than 6 adjustments), in the *control* group only 12% (3 out of 25) of subjects had a static strategy, as opposed to 36% (8 out of 22) in the *control* group. When adjusting, the subjects exhibited chasing demand heuristic. If the received order was higher (lower) than demand and the consequent order is adjusted downwards (upwards), it is treated here as demand chasing. In the *test* group, 68% (224 out of 330) of adjustments (does not include 16th round decisions) can be classified as demand chasing; in the *control* group 73% (154 out of 211) – again a comparable result to S&C (69%).

Three representative sequences of orders and corresponding commentaries are in Figure 11. Some conclusions can be drawn from all commentaries: 20% (5 out of 25) of the *test* group and 32% (7 out of 22) of the *control* group subjects explicitly mention mean demand in their strategies, if only some of them describe it as the initial point for their strategy that was later adapted. Thus, anchoring to the mean demand seems one plausible cause for deviation from the optima. At least 16% (4 out of 25) in the *test* group directly stated that their strategy was based on demand forecasting and few others mention this as a part of their strategy. One subject from the *control* group used reverse demand chasing, i.e., placed a low order after high

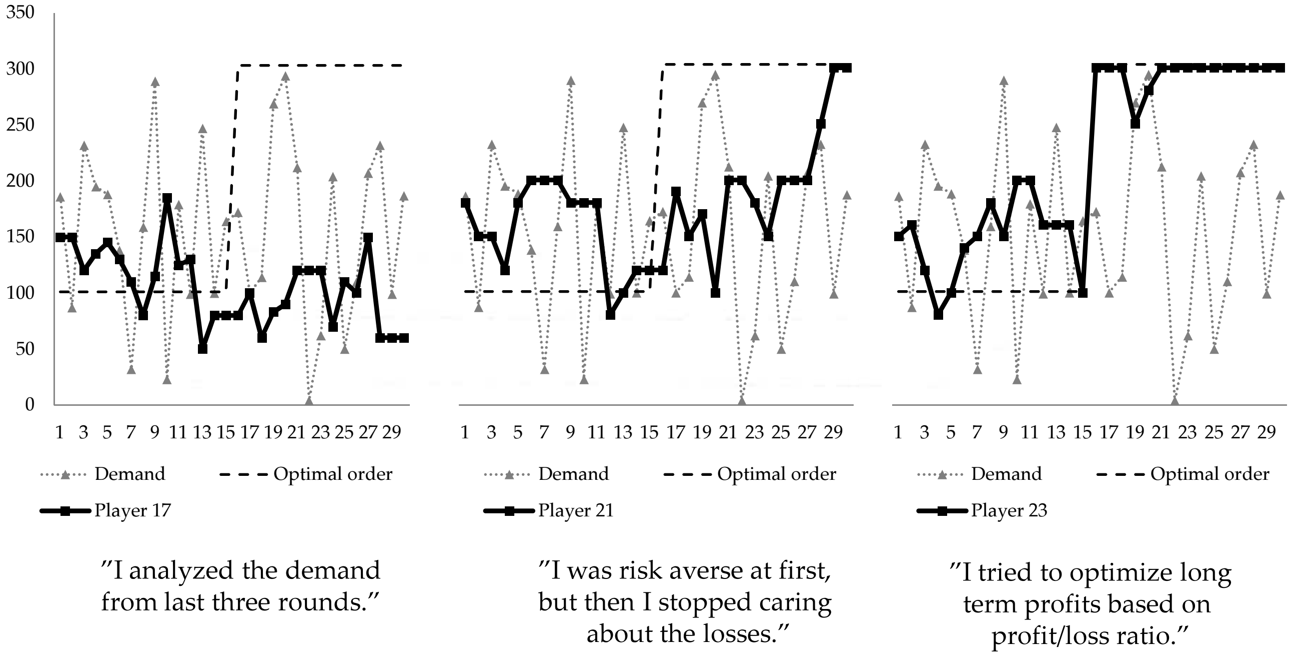

Figure 11: Example results and feedbacks from three *test* group subjects.

demand and vice versa. Also, 24% (6 out of 25) of the *test* group and 18% (4 out of 22) of the *control* group subjects report that their strategy was based on some initial guess and learning during the rounds.

4.4 Group *PhD*

We also collected data from 13 PhD students of a cross-disciplinary postgraduate course at Aalto University. These subjects comprise an international group with a diverse academic and cultural background and little experience with operations management. The group conducted the same experiment as the *test* group above. The results are not directly comparable for two reasons: i) the subjects conducted the experiment in an uncontrolled on-line environment; and ii) instead of a monetary reward, peer comparison (a high-score table) was used as incentive. The latter should not compromise the usability of results: there is evidence that social comparison provides strong incentive for subjects in a newsvendor setting (see Avci et al. 2012 and references therein). The lack of control, however, is a concern in that direct comparison with the laboratory results may not be warranted. For example, it was also impossible to ensure that

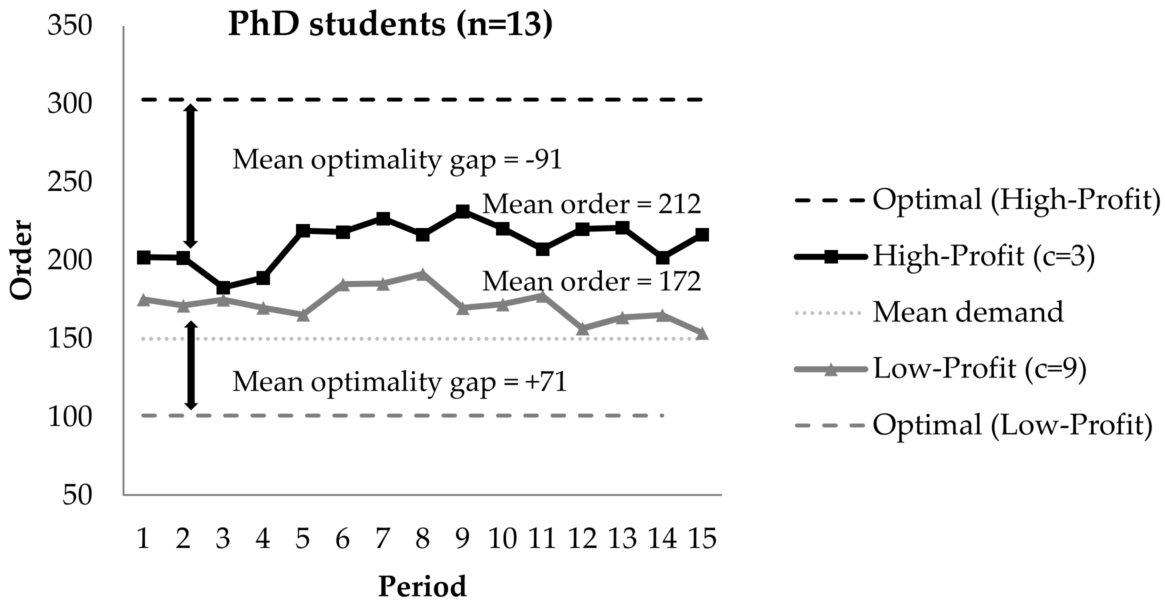

Figure 12: Results of group *PhD*.

there was no communication between subjects, but we doubt that it would happen because of the competitiveness aspect. In any case, these results are more suggestive than conclusive.

The results are plotted in Figure 12 and they are very similar to results of the *test* group. The pull-to-center effect is clearly observable, and based on the qualitative feedback, the rationale for ordering decisions are similar to the (undergraduate) *test* group results, too.

4.5 Group *consultants*

The fourth group consisted of 11 consultants employed by a mid-sized supply chain consulting company. The subjects were skilled experts in the areas of supply chain planning, decision support systems, and project management. We note that the caveats related to non-monetary incentives and use of on-line environment apply also here.

The results (Figure 13) show that even though the orders are still non-optimal, the group performed significantly better under the high profit condition: the mean order was 257 (compared to 211 of the *test* group and 212 of the *PhD* group). For this reason, this group was able to generate over 20% more profit than the *test* group or the *PhD* group.

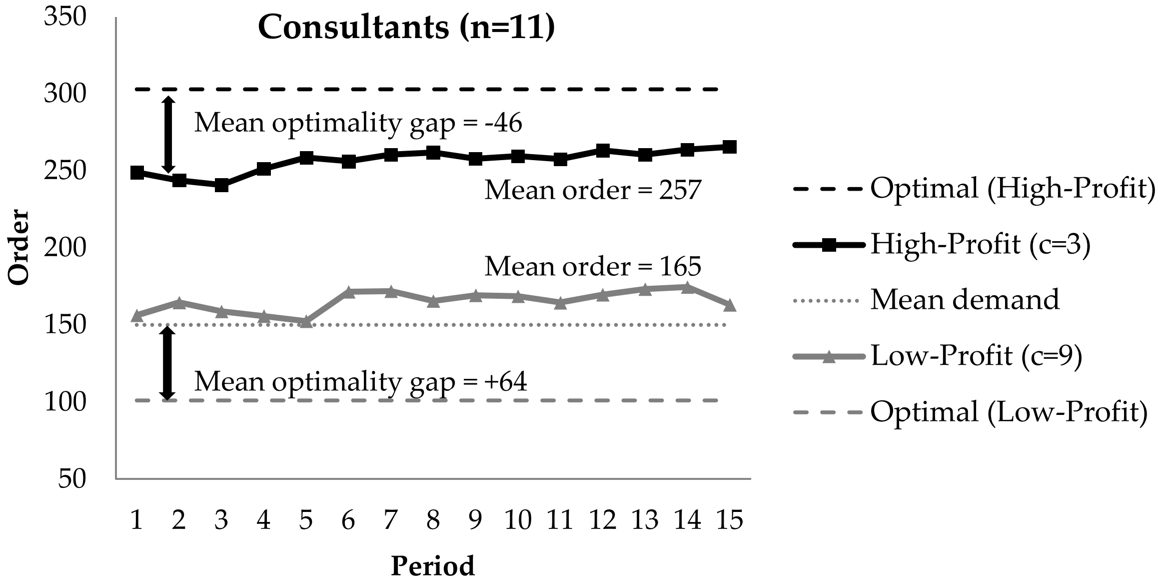

Figure 13: Results of group *consultants*.

In this group, 27% (3 out of 11) of the subjects were familiar with the newsvendor problem, compared to 1 or 2 subjects in the other (larger) groups. Even though in-depth understanding of the newsvendor problem can significantly improve the results (Bolton et al. 2012) the presence of uncertain supply in our experiment makes the decision making more complicated, compared to the standard newsvendor problem. In any case, altogether four subjects (the three familiar with the newsvendor problem and one that was not) had near optimal strategies. Two of these four subjects reported a solution that was based on critical ratio and divided the resulting order quantity by 0.7, which in this case yields orders close to real optima. The two other subjects reported using some sort of simulation based spreadsheet solution.

However, even without these four optimal performances in the *consultants* group, the average order in the high-profit case is 234 which is significantly higher than in the other groups. In their comments about ordering strategy, 73% (8 out of 11) of the subjects in the group explained that they tried to balance risk based on high and low profit margin (i.e., define the critical ratio), either using decision support, experience, or some simple heuristic. While it is not possible to make exact inference from these subjects' free verbal comments, based on the replies from other groups it seems apparent that less than 50% of subjects in those groups were able to define or estimate the critical ratio (Gavirneni and Isen (2010) report similar observations).

4.6 Discussion of results

In all three groups that faced uncertain supply (*test, PhD, consultants*), the average orders are either below the optimum (high-profit) or above the optimum (low-profit). This too low/too high pattern is consistent with the results of S&C. Yet it is unclear how this relates to earlier results under deterministic supply. Even though the actual orders themselves are between the optima and the mean demand and thus indicate a pull-to-center effect, the subjects should have expected to receive only a share of their order due to supply uncertainty. It is possible to convert the actual orders to expected received orders by multiplying with $\mathbb{E}[Z] = 0.7$. Consider now the results of the *test* group: under the high-profit condition, they ordered 211 items, and under low-profit condition 165 items. These translate into expected received order of $211 \cdot 0.7 = 148$ and $165 \cdot 0.7 = 116$. So actually, the expected recieved orders in the *test* group were below the mean demand not only under low-profit but also *under high-profit condition.* These results can be compared with the *control* group results of 188 (high-profit) and 125 (low-profit). Interestingly enough, the high-profit condition is clearly more problematic under uncertain supply, but in the low-profit condition, the order amounts under supply uncertainty are somewhat similar.

Many subjects used the term *risk* in their comments, albeit in varying contexts. This indicates that risk aversion could have had an impact on the subjects' behavior. In particular under the high-profit condition, the profit curve is rather flat around the optimum and risk can be reduced by ordering less. In Figure 6(a), the optimal order is denoted with a circle, and the test group result with a square. The expected profit for the (severely non-optimal) *test* group order is 864, which is only 10% smaller than the optimal expected profit 960. At the same time, the 95%-level CVaR was reduced by over 40%. Perhaps the subjects were satisfied with the (relatively) high profits that resulted from undersized orders and did not want to take the risk of a loss by ordering more.

Yet risk aversion is only a partial explanation because it should lead to less-than-optimum orders also in the low-profit case – which did not happen. Instead, as noted in Figure 6(b), the *test* group order (the square between $q = 140$ and $q = 160$) was significantly more than the optimum (the circle between $q = 80$ and $q = 100$) so that the expected profit was 67, which is 38% smaller than the optimal expected profit 101. At the same time, however, also risk increased with a significant 77%. So, even though the subjects seem to perform better under the low-profit condition when measured with order quantities, it is actually a lot worse if measured with profit or risk.

The commentaries by the *test* group subjects indicate that these subjects paid very little attention to uncertain supply, or at least this was not mentioned in the feedback. Only one subject reported first estimating the demand to be fulfilled, and then inflating the order by diving by the expected supply yield. However, many subjects mentioned mean demand in one way or another in their commentaries which suggests anchoring behavior. Thus, while risk aversion could partially explain the results, anchoring to the mean demand together with ignorance of supply uncertainty also play a big role.

In addition to these findings, we found some evidence that general formal education does not improve newsvendor decisions (*PhD* group results) but instead, context specific experience may give an advantage (*consultants* group results). There is already evidence from Kremer et al. (2010) that newsvendor behavior is context-sensitive. While Kremer et al. found that subjects perform better in a neutral (lottery) context than in an operational one (demand-supply matching), they used students who probably were not experienced in operations management. It might be that context also provides an advantage to those who solve operations management problems in their profession.

5 Managerial insights

We have shown how random supply yield increases the newsvendor order quantity in a non-trivial manner. In particular, the optimal order under supply uncertainty cannot be derived from the optimal order under deterministic supply by dividing by the expected supply yield – this can be severely non-optimal in case yield uncertainty is high. In case of interdependent demand and supply, it was found that the newsvendor can benefit from positive interdependency between demand and supply yield, and that it also has a big impact on the optimal newsvendor decision especially in the low-profit case. If the newsvendor is risk averse, we note that Conditional-Value-at-Risk can be reduced significantly by reducing the order quantity and this has only small impact on the expected profit.

These results have management implications: especially in arms-length relationships in which quality problems and other reasons for imperfect yield cannot be influenced by the buyer, order size is the only way to control the amount of products available for the selling season. In this case, a buyer who anticipates yield problems can make significantly higher profits than one that assumes a perfect yield. Also, increasing positive dependency between demand and supply yield is recommended but perhaps an abstract advise. However, some procurement policies could promote such dependency. For example, if the supplier is flexible, the buyer can introduce contractual incentives that motivate high yield under high demand and allows lower yields when the demand is low.

Alongside these theoretical observations, managers should be aware of the behavioral challenges in the newsvendor problem. We found that the mean demand is a strong anchor, which causes severely non-optimal behavior especially in case of a high-profit product when the optimal order can be very far from the mean demand – and even above the maximum demand. However, this deviation does reduce risks. Thus, risk aversion is a possible explanation for the undersized orders. In the low-profit case, on the other hand, the observed (oversized) order is relatively closer to the optimum, but it leads to both decreased expected profit and increased risk level. Thus, managers should be particularly cautious about oversized orders in case of a low-profit product. We found some evidence that adopting an analytical approach (calculation or simulation) improves newsvendor decision making dramatically: some subjects reported that they used a simulation based decision tool that helped them perform (near) optimally in the experiment. Thus, the use of decision support can assist in mitigating the aforementioned behavioral biases.

6 Conclusions

We have analyzed the newsvendor problem under uncertain supply by using the stochastic yield model. The results of the analysis show that stochastic supply has a major impact on the newsvendor decision, and that interdependency between demand and supply can amplify this impact. We also tested how well human decision makers account for supply uncertainty. Our reference group replicated a well-known study by Schweitzer and Cachon (2000) with similar results. Our test group faced, in addition, supply yield randomness framed as quality problems. Comparison of these two groups revealed that it is difficult for humans to account for demand and supply uncertainties simultaneously.

There are some potential extensions related to our work and modeling the newsvendor problem. First, we modeled both uncertainties with uniform distributions. This is convenient for both theoretical and experimental purposes, but it would be interesting to study other distributions and a robust demand-supply framework, too. Second, the stochastic yield model implies that the supply yield is independent of the order quantity. In many settings, the yield is not exogenous, but large orders can imply better yields or vice versa. Especially the proposed model for interdependency between demand and supply yield is not entirely realistic: it would be interesting to analyze the case where the supply yield is dependent on the order quantities, not on the end demand. Thus, other forms of yields and interdependencies should be explored in further research.

We also suggest two potential future research topics related to newsvendor experiments: First, the behavior appears heterogeneous. In all our experiments, there were those who clearly understood what the problem was about (calculating the critical ratio based on cost asymmetry) and perhaps just could not do the calculations; and then there were those who chased demand or arbitrarily adjusted their orders based on intuition. Perhaps a more structured study could help identify different types of behavior, instead of focusing on the average behavior to reveal mean behavioral effects such as pull-to-center. Second, we would like to propose a tentative yes to Gavirneni and Isen (2010), who ask: “If the subjects were provided a decision support system would the subjects be able to reach the optimal decision?". Systematic experimental research that would investigate the use of decision support systems in newsvendor setting would help to confirm this proposal, and would have much value for practitioners.

Supply uncertainty appears in many industries, and it should be accounted for in newsvendor type decisions related to inventory management, procurement, and sourcing. But its impacts are difficult to assess, as our results demonstrate. For example, under supply uncertainty, it can be optimal to order more than the maximum demand, and such decisions can be very difficult to justify to higher management. However, simple decision tools based on spreadsheet solutions can be very effective. Also, our experiment suggests that subjects can learn the dynamics of the problem. This indicates that both decision support tools and experiments combined with management training can be used to overcome decision biases in the newsvendor setting.

References

Avci, B., Loutfi, Z., Mihm, J., Belavina, E., and Keck, S. (2012). Comparison as incentive: Newsvendor decisions in a social context. *Production and Operations Management (forthcoming).*

Balakrishnan, N. and Lai, C. (2009). *Continuous Bivariate Distributions*, chapter 2: Distributions expressed as copulas. Springer New York.

Benzion, U., Cohen, Y., Peled, R., and Shavit, T. (2008). Decision-making and the newsvendor problem: An experimental study. *The Journal of the Operational Research Society*, 59(9):1281–1287.

Bolton, G. E. and Katok, E. (2008). Learning by doing in the newsvendor problem: A laboratory investigation of the role of experience and feedback. *Manufacturing and Service Operations Management*, 10(3):519–538.

Bolton, G. E., Ockenfels, A., and Thonemann, U. W. (2012). Managers and students as newsvendors. *Management Science*, 58(12):2225–2233.

Bostian, A. A., Holt, C. A., and Smith, A. M. (2008). Newsvendor "Pull-to-Center" effect: Adaptive learning in a laboratory experiment. *Manufacturing and Service Operations Management*, 10(4):590–608.

Burke, G. J., Carrillo, J. E., and Vakharia, A. J. (2009). Sourcing decisions with stochastic supplier reliability and stochastic demand. *Production and Operations Management*, 18(4):475–484.

Cachon, G. P. and Lariviere, M. A. (2005). Supply chain coordination with revenue-sharing contracts: Strengths and limitations. *Management Science*, 51(1):30–44.

Chen, Y., Xu, M., and Zhang, Z. G. (2009). A risk-averse newsvendor model under the cvar criterion. *Operations Research*, 57(4):1040–1044.

Dada, M., Petruzzi, N. C., and Schwarz, L. B. (2007). A newsvendor's procurement problem when suppliers are unreliable. *Manufacturing and Service Operations Management*, 9(1):9–32.

De Véricourt, F., Jain, K., Bearden, J. N., and Filipowicz, A. (2013). Sex, risk and the newsvendor. *Journal of Operations Management*, 31(1–2):86–92.

Gavirneni, S. and Isen, A. M. (2010). Anatomy of a newsvendor decision: Observations from a verbal protocol analysis. *Production and Operations Management*, 19(4):453–462.

Gerchak, Y. and Grosfeld-Nir, A. (2004). Multiple lotsizing in production to order with random yields: Review of recent advances. *Annals of Operations Research*, 126(1–4):43–69.

Gumbel, E. J. (1960). Bivariate exponential distributions. *Journal of the American Statistical Association*, 55(292):698–707.

Henig, M. and Gerchak, Y. (1990). The structure of periodic review policies in the presence of random yield. *Operations Research*, 38(4):634–643.

Inderfurth, K. (2004). Analytical solution for a single-period production inventory problem with uniformly distributed yield and demand. *Central European Journal of Operations Research*, 12(2):117–127.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291.

Katok, E. (2011). Using laboratory experiments to build better operations management models. *Foundations and Trends in Technology, Information and Operations Management*, 5(1):1–86.

Keren, B. (2009). The single-period inventory problem: Extension to random yield from the perspective of the supply chain. *Omega*, 37(4):801–810.

Keren, B. and Pliskin, J. S. (2006). A benchmark solution for the risk-averse newsvendor problem. *European Journal of Operational Research*, 174(3):1643–1650.

Khouja, M. (1999). The single-period (news-vendor) problem: literature review and suggestions for future research. *Omega*, 27(5):537–553.

Kremer, M., Minner, S., and Van Wassenhove, L. N. (2010). Do random errors explain newsvendor behavior? *Manufacturing and Service Operations Management*, 12(4):673–681.

Lau, N., Hasija, S., and Bearden, J. N. (2013). Newsvendor pull-to-center reconsidered. *Decision Support Systems*. Available at: http://dx.doi.org/10.1016/j.dss.2012.12.041.

Lurie, N. H. and Swaminathan, J. M. (2009). Is timely information always better? the effect of feedback frequency on decision making. *Organizational Behavior and Human Decision Processes*, 108(2):315–329.

Petruzzi, N. C. and Dada, M. (1999). Pricing and the newsvendor problem: A review with extensions. *Operations Research*, 47(2):183–194.

Porteus, E. L. (1990). Stochastic inventory theory. In Heyman, D. and Sobel, M., editors, *Handbooks in OR & MS, Vol. 2*. Elsevier Science Publishers, North-Holland, Amsterdam, The Netherlands.

Qin, Y., Wang, R., Vakharia, A. J., Chen, Y., and Seref, M. M. (2011). The newsvendor problem: Review and directions for future research. *European Journal of Operational Research*, 2(1):361–374.

Rockafellar, R. T. and Uryasev, S. (2000). Optimization of Conditional Calue-at-Risk solution. *Journal of Risk*, 2(3):21–41.

Schweitzer, M. E. and Cachon, G. P. (2000). Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science*, 46(3):404–420.

Su, X. (2008). Bounded rationality in newsvendor models. *Manufacturing and Service Operations Management*, 10(4):566–589.

Tomlin, B. (2006). On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science*, 52(5):639–657.

Yang, S., Yang, J., and Abdel-Malek, L. (2007). Sourcing with random yields and stochastic demand: A newsvendor approach. *Computers & Operations Research*, 34(4):3682–3690.

Yano, C. A. and Lee, H. L. (1995). Lot sizing with random yields: A review. *Operations Research*, 43(2):311–334.

A Appendix

A.1 Derivation of (3.2)

$$\pi(q) := \frac{\partial \mathbb{E}[\Pi(q)]}{\partial q} = \frac{1}{2}\mu_Z(p+s-2c) - \frac{1}{2}(p-s)\frac{\partial}{\partial q}\mathbb{E}[|Z\cdot q - D|]\Big), \tag{A.1}$$

where

$$\mathbb{E}[|Z\cdot q-D|] = \begin{cases} \int_{\underline{Z}}^{1}\Big(\int_0^{z\cdot q}(z\cdot q-d)f(z,d)\mathrm{d}d + \int_{z\cdot q}^{\overline{D}}(d-z\cdot q)f(z,d)\mathrm{d}d\Big)\mathrm{dz}. & \text{if } q\in[0,\overline{D}], \\ \int_{\underline{Z}}^{\overline{D}/q}\Big(\int_0^{z\cdot q}(z\cdot q-d)f(z,d)\mathrm{d}d + \int_{z\cdot q}^{\overline{D}}(d-z\cdot q)f(z,d)\mathrm{d}d\Big)\mathrm{dz} & \\ \qquad + \int_{\overline{D}/q}^{1}\Big(\int_0^{\overline{D}}(z\cdot q-d)f(z,d)\mathrm{d}d\Big)\mathrm{dz} & \text{if } q\in[\overline{D},\overline{D}/\overline{Z}] \end{cases}$$

Note that it is also possible that $\overline{D} < \underline{Z}\cdot q$, which would imply that the order is *always* larger than maximum demand. For a profit seeking newsvendor, however, this is not an option because one could always improve profit by ordering a little less. Thus, the optimal quantity must lie in the interval $[0,\overline{D}/\overline{Z}]$.

If $q\in[0,\overline{D}]$ then

$$\begin{aligned} \frac{\partial}{\partial q}\mathbb{E}\big[|Z\cdot q-D|\big] &= \int_{\underline{Z}}^{1} z\Big(\int_0^{z\cdot q} f(z,d)\mathrm{d}d - \int_{z\cdot q}^{\overline{D}} f(z,d)\mathrm{d}d\Big)\mathrm{dz} \\ &= \int_{\underline{Z}}^{1}\int_0^{\overline{D}} z f(z,d)\mathrm{d}d\mathrm{dz} - 2\cdot\int_{\underline{Z}}^{1}\int_{z\cdot q}^{\overline{D}} z f(z,d)\mathrm{d}d\mathrm{dz} \\ &= \mu_Z - 2\cdot\int_{\underline{Z}}^{1} z\int_{z\cdot q}^{\overline{D}} f(z,d)\mathrm{d}d\mathrm{dz} \end{aligned} \tag{A.2}$$

and if $q\in[\overline{D},\overline{D}/\overline{Z}]$ the derivative is

$$\begin{aligned} \frac{\partial}{\partial q}\mathbb{E}\big[|Z\cdot q-D|\big] &= \int_{\underline{Z}}^{\overline{D}/q} z\Big(\int_0^{z\cdot q} f(z,d)\mathrm{d}d - \int_{z\cdot q}^{\overline{D}} f(z,d)\mathrm{d}d\Big)\mathrm{dz} \\ &+ \int_{\overline{D}/q}^{1} z\Big(\int_0^{\overline{D}} f(z,d)\mathrm{d}d\Big)\mathrm{dz} \\ &= \int_{\underline{Z}}^{1}\int_0^{\overline{D}} z f(z,d)\mathrm{d}d\mathrm{dz} - 2\cdot\int_{\underline{Z}}^{\overline{D}/q}\int_{z\cdot q}^{\overline{D}} z f(z,d)\mathrm{d}d\mathrm{dz} \\ &= \mu_Z - 2\cdot\int_{\underline{Z}}^{\overline{D}/q} z\int_{z\cdot q}^{\overline{D}} f(z,d)\mathrm{d}d\mathrm{dz}. \end{aligned} \tag{A.3}$$

Thus, the derivative of the expected profit can be combined from (A.2) and (A.3):

$$\pi(q) = (s-c)\mu_Z + (p-s)\int_{\underline{Z}}^{\min\{1,\overline{D}/q\}} z \int_{z\cdot q}^{\overline{D}} f(z,d)\mathrm{d}d\mathrm{d}z. \tag{A.4}$$

A.2 Calculation of profit distributions in Figure 2

Let Y be the profit for given q, and $Q = Z\cdot q$ the received order. Further, we assume $s = 0$, i.e., no salvage cost. Now, the distribution of Y can be calculated with:

$$Pr(Y \leq y) = \underbrace{\int_{\underline{Z}\cdot q}^{300} f_D \int_{\underline{Z}\cdot q}^{g(d)} f_Q \,\mathrm{d}t\mathrm{d}d}_{D>Q} + \underbrace{\int_{\underline{Z}\cdot q}^{q} f_Q \int_{0}^{h(t)} f_D \,\mathrm{d}d\mathrm{d}t}_{Q>D}, \tag{A.5}$$

where where f_D is the demand distribution, f_Q is the distribution of receiver order, i.e., $f_Q = (q-\underline{Z}\cdot q)^{-1}$, and $h(t)$ is the upper limit for q in case $Q > D$ and $g(d)$ the upper limit for d in case $D > Q$ - see Figure 14 for illustration. Note that A_1 and A_2 are just examples for particular q and y values; the complete limits are piecewise defined with $g(d) = \max\{\min\{q, \frac{y}{p-c}, d\}, \underline{Z}\cdot q\}$ and $h(t) = \max\{\min\{t, \frac{y+ct}{p}, \overline{D}\}, 0\}$.

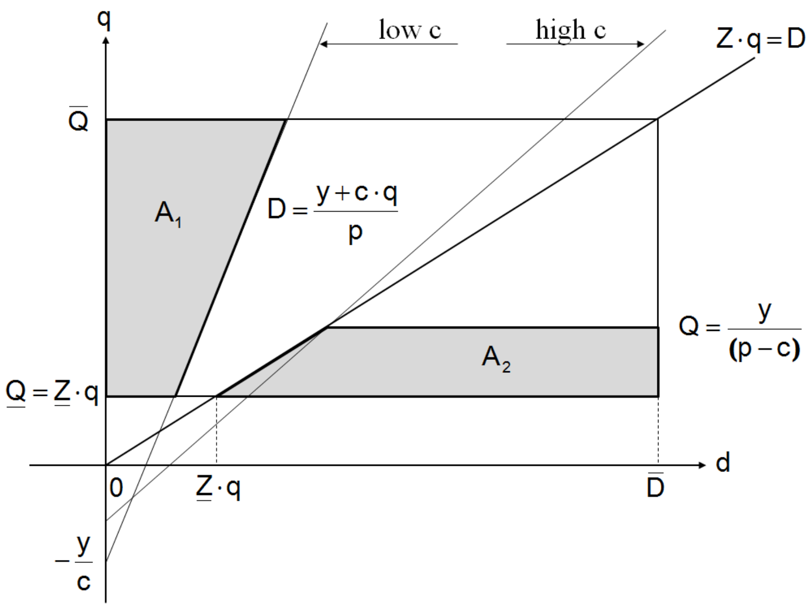

Figure 14: The integration limits illustration for (A.5).

As an example, setting $p = 12$, $c = 3$, $\overline{D} = 300$, and $\underline{Z} = 0.4$, and $q = 303$ gives:

$$Pr(Y \leq y) = \begin{cases} 1 & \text{if } y \geq 2700, \\ -\frac{697}{4848} + \frac{37962 - y^2}{2208870} & \text{if } 2691 < y < 2700, \\ -\frac{73}{160} + \frac{37962y - 7y^2}{35341920} & \text{if } \frac{5454}{5} < y \leq 2691, \\ \frac{(6363 + 10y)}{36000} & \text{if } -\frac{1818}{5} < y \leq \frac{5454}{5}, \\ \frac{(909 + y)^2}{3926880} & \text{if } -909 < y \leq -\frac{1818}{5}, \\ 0 & \text{if } y \leq -909. \end{cases} \tag{A.6}$$

First order differentiation of (A.6) gives the probability distribution function.

A.3 Expected profit under interdependent uniformly distributed demand and supply

For a high-profit product, the expected profit is:

$$\mathbb{E}[\Pi(q)] = \begin{cases} 0 & q = 0 \\ q\left(\frac{63}{10} - \frac{13}{1250}q + \frac{7}{2500}\theta q - \frac{127}{18750000}\theta q^2\right) & 0 < q \leq 300 \\ 3000 - \frac{2000}{3}\theta - \frac{37}{10}q + \frac{4}{5625}q^2 + \frac{16}{16875}\theta q^2 - & \\ \frac{19}{31640625}\theta q^3 - \frac{3 \cdot 10^5}{q} + \frac{3.5 \cdot 10^5}{27q}\theta - \frac{4.5 \cdot 10^6}{q^2}\theta & 300 < q \leq 750 \\ 1800 - \frac{21q}{10} & q > 750 \end{cases} \tag{A.7}$$

When $q \leq 300$, rearranging terms gives $\mathbb{E}[\Pi(q)] = X + \theta(Yq - Zq^2)$. In the given range, $(Yq - Zq^2)$ is always positive and thus increasing θ increases profit. For $300 < q \leq 750$, we can express profit with $\mathbb{E}[\Pi(q)] = X + \theta \cdot \Phi(q)$. Because $\Phi(300) > 0$ and $\Phi(750) = 0$, it is sufficient to show that $\Phi(q)$ is strictly decreasing to show that it is positive. Differentiation gives $\partial\Phi/\partial q = \frac{(q-750)^2}{10546875q^3}(1687500000 - 2062500q - 8500q^2 - 19q^3)$, which is always negative in the given interval. Thus $\Phi(q)$ is always positive and, again, θ increases the expected newsvendor profit. Note that the low-profit condition yields exactly same form of expected profit with one exception: all multipliers of q are more negative. In fact, one gets the expected profit in the low profit case by adding -21/5 to each multiplier of q. Thus, the impact of dependency remains the same.