Behavior in single-echelon dual sourcing settings: An experiment

Tamás Csermely[†], Stefan Minner[††]

Abstract

We conducted an experiment with a single-echelon dual sourcing setting including a slow, cheap and a fast, expensive supply mode. We show that subjects have a preference for flexibility and thus overuse the fast and expensive supplier compared to the normatively optimal behavior. We claim that this behavioral pattern is a result of subjects' conscious preference for less on-hand inventory and that it is driven by myopic optimization. Moreover, we prove that the observed behavior is also consistent with ex post inventory error minimizing preferences, which can be explained by the psychological cost model described by Ho, Lim, and Cui (2010). We see that, once we take these characteristics into account, the relative overusage of the fast supply mode is not prevalent and that subjects use the fast and the slow option according to the above specified extension of normative theory. Furthermore, we document significant cost savings in all dual sourcing treatments compared to the single-sourcing setting even if the fast supply mode is very pricy, but even more so if it is relatively less expensive. Furthermore, we observe that the gap between the normatively achievable cost and the actual one decreases considerably by introducing an expensive dual sourcing option, whereas we see no such improvement compared to the normative costs if the fast mode is relatively inexpensive. We show that the slow supply mode is used significantly more often ceteris paribus if it is framed as an internal partner compared to when it is described as an external supplier. This can be explained by extending the psychological cost model to external suppliers which implies that the choice for internal versus external suppliers are in fact driven by behavioral causes as well.

Introduction

The task to match supply with demand at a reasonable cost level is extremely challenging especially if multiple supply modes exist. In order to bridge this gap, theoretical research was conducted on the implications of multiple sourcing already before the millennium. Most importantly from this stream of the literature, Ramasesh, Ord, and Hayya (1993) confirmed that substantial cost savings could be achieved if the lead times and the cost structure of the supply modes are not the same. Burke, Carrillo, and Vakharia (2007) considered a single-product dual sourcing setting with demand uncertainty allowing different characteristics for their suppliers. They found that single-sourcing is only optimal in particular cases when no diversification benefits can be obtained. Veeraraghavan and Scheller-Wolf (2008) tested the dual index sourcing policy and claim that it costs only 1-2% more than the policy obtained by stochastic dynamic programming, however, it leads to cost reductions of up to 50% compared to single-sourcing. Their results were validated in various environments including non-stationary demand, supply disruptions and demand spikes. Lyon (2006) documented that a dual-sourcing policy had lead to cost decreases for the US government. Klosterhalfen, Kiesmuller, and Minner (2011) tested various dual sourcing policies in a single-product periodic review system with backorders. They showed that following a constant-order policy means a good trade-off in terms of simplicity and results. A review on multiple sourcing in supply chains is provided by Minner (2003).

In addition, it was assumed that sourcing decisions could be motivated by ownership structure (Lee and Billington, 1993). This issue symbolizes a relevant aspect for supply chain design and strategy: It implies that non-financial motivations exist for sourcing from internal or external partners, which leads to the broader make-or-buy discussion. Experimental research started to boom in the field of operations management, however, the previously described theoretical findings have not been tested empirically. Thus, our goal is to test the above mentioned issues compared to normative theory. Boudreau (2004); Bendoly, Donohue, and Schultz (2006); Bendoly, Croson, Goncalves, and Schultz (2010) provide extensive reviews on the behavioral operations management literature. We

† Research and teaching associate in Prae-Doc, Vienna University of Economics and Business, Institute for Public Sector Economics, Welthandelsplatz 1., 1020 Vienna, Austria. E-mail: tamas.csermely@wu.ac.at
† Ph. D. candidate, University of Vienna, Doctoral School of Logistics and Operations Management
†† Full Professor, Logistics and Supply Chain Management, TUM School of Management, Technische Universität München, Room 1573, Arcisstraße 21, 80333 Munich, Germany. E-mail: stefan.minner@tum.de

are interested in whether the overusage of the fast mode can be explained by psychological costs (Ho et al., 2010), ex post inventory error minimization (Kremer, Minner, and Van Wassenhove, 2013), or other preferences such as risk attitude (de Vericourt, Jain, Bearden, and Filipowicz, 2013), loss aversion (Schweitzer and Cachon, 2000), regret aversion (Kremer and Minner, 2008) or cognitive reasons (Moritz, Hill, and Donohue, 2013).

The experimental design

We ran an experiment at the Vienna Center for Experimental Economics with 86 subjects who earned €25.56 on average. Subjects participated in an inventory task, an incentivized preference elicitation tasks and a questionnaire. In the inventory task, subjects were told that they owned a store that supplies spare parts and that they had to make decisions on how many units to order in 48 periods. Subjects could order from two options. Orders placed to the more expensive fast option were delivered instantly, whereas the slow option was less expensive, but had a lead time of one period. Both options had no capacity constraints, so subjects could have ordered any given integer combination from these options. Feedback was provided after each period: Subjects saw the demand, their ordered quantities from both options, their inventory level, the amount in transit with the slow option, the corresponding costs in that specific period and also their total cost until that particular period. Initially, subjects had 6 units on stock and 4 items in transit from the slow supplier. In each period, the demand for spare parts was randomly determined between 0 and 8 following a uniform distribution. We used the same uncorrelated demand sequence for all subjects to make sure that our results cannot be driven by the different realization of random numbers. We denote I as the on-hand inventory level, h as the unit holding cost, b as the unit backlog cost. S and F are the respective quantities ordered from the slow and the fast option with their corresponding prices p_S and p_F. The total cost C for subject i in period t can be calculated as follows:

$$C_{it} = \sum_{t=1}^{t}[h * \max\{I_{it}, 0\} - b * \min\{I_{it}, 0\} + p_S * S_i + p_F * F_i] \quad (1)$$

We used a relative rewarding mechanism based on subjects' total cost across the 48 periods in the inventory task. This practice is in line with other experimental studies (Croson, Donohue, Katok, and Sterman, 2013). Subjects' payment P was determined by the following formula:

$$P_i = €15 * \frac{\max(C) - C_i}{\max(C) - \min(C)} \quad (2)$$

The parameters h=2, b=8 and p_S=6 were chosen to resemble the ratio of holding, backorder and unit costs in the real world and they were kept constant in all the treatments. We used a between-subject design with 4 different treatments and N subjects per treatment. Subjects could order from only the slow option in the single-sourcing ST treatment (N=24). We had three treatments with both order options present. In treatment D10T (N=22), we used the parameter p_F=10 whereas p_F was 7 in D7T (N=18). In the former one, it is not optimal to use the fast supply mode according to normative theory, whereas it is optimal under particular circumstances to do so in the latter one. Every order option was framed as an external partner in all these previously mentioned treatments, whereas we framed the fast option as an external partner and the slow option as an internal one in treatment D7TF (N=22) where p_F had a value of 7 as in D7T.

We elicited cognitive reflection (Frederick, 2005), regret aversion (Raeva, Mittone, and Schwarzbach, 2010), risk attitude (Drichoutis and Lusk, 2012) and loss aversion (Gächter, Johnson, and Herrmann, 2010) only after the inventory task to avoid potential framing effects. In the end, subjects received an additional €5.00 for filling out a questionnaire which surveyed socioeconomic characteristics. There is evidence (Croson et al., 2013) that indirect effects stemming from strategic uncertainty could play a role in several economics-related settings. Our design rules this issue out since each of our task involves only individual decision-making. To make sure that subjects understand the nature of the tasks they were supposed to do, we included several control questions for each part. In addition to the inventory task, one of the preference elicitation tasks was chosen to be payoff-relevant in line with the random lottery incentive mechanism (Cubitt, Starmer, and Sugden, 1998).

Analysis

We calculated the normatively optimal usage ratios, the optimal ordering policy and also the expected total cost with stochastic dynamic programming over a horizon of 500 periods and evaluated it over 48 periods truncating last periods to avoid endgame effects in line with Whittemore and Saunders (1977)[2].

We compared the total costs in all the treatments to see whether introducing a more expensive second sourcing option can lead to cost reductions. The total costs were[3] 2009.0 (448.4) in ST, 1704.8 (224.0) in D7T, 1750.2 (191.9) in D7TF and 1752.8 (138.4) in D10T. The Mann-Whitney-Wilcoxon test (henceforth MWW) signals that the total cost is significantly higher in ST than in D7T ($p=0.03$), D7TF ($p<0.01$) and D10T ($p<0.01$), whereas total costs in D7T, D7TF and D10T could not be distinguished from each other statistically. Moreover, the normative solution predicted a total cost of 1647.4 for treatments ST and D10T and estimated 1598.9 for D7T and D7TF. We documented that the average total cost of 1704.8 (224.0) in D7T was 6.6% higher – although not statistically different from – the theoretically predicted 1598.9 (one-sample Wilcoxon-test, $p=0.31$), whereas the total cost of 1750.2 (191.9) in D7TF was 9.5% higher ($p<0.01$). This suggests that the internal/external framing is responsible for this difference in total costs[4]. In treatments ST and D10T, the actual total costs of 2009.0 (448.4) and 1752.8 (138.4) were significantly higher than the normatively achievable 1647.4 (one-sample Wilcoxon-tests, $p<0.01$). This implies that significant cost reductions (-12.8%) can be achieved just by adding a faster sourcing opportunity even if it is very expensive, which is in line with theoretical predictions. In addition, we saw that the gap between the model-predicted cost and the actual one was similar if the fast sourcing mode was relatively cheap (105.9; 6.6%) or expensive (105.4; 6.4%), whereas it was relatively large in ST (361.6; 21.9%). All the implications of the above mentioned results hold if we truncate the first ten and the last five periods from our analysis[5].

In treatments ST and D10T, the ordering policy of (3;11) is optimal, whereas it is (2;10) in treatments D7T and D7TF. In theory, the model predicts a slightly higher average inventory level for the former ones before the demand is realized. The reason for this is that it is not optimal/not possible to order from the fast option in ST and D10T. We investigated whether subjects could have followed any systematic ordering policy during the experiment. We found no evidence for having followed a particular order up to policy, since the inventory levels before demand realization were heterogeneous even if analyzing them individually or aggregating them by treatments. They were 6.90 (3.70) in ST, 7.02 (2.33) in D7T, in 8.04 (1.86) D10T and 8.35 (2.52) in D7TF. Moreover, the relatively large coefficients of variation (CoV) for the slow (fast) option suggest that they could not have exercised a constant order policy either: CoVs were 0.81 (-) in ST, 0.93 (1.78) in D7T, 0.86 (1.93) in D7TF and 0.65 (3.45) in D10T. These values were in almost all cases significantly higher than the CoVs that the normatively optimal ordering policy predicted. Moreover, we found no evidence that our subjects could have followed a dual-index policy. First of all, this would have been very complicated cognitively. Moreover, we observed significant negative correlations between the on-hand inventory level and orders from the slow, the fast and the total units ordered per period at the individual level[6]. We did not observe any significant correlations between the ordered amounts and the demand pattern. It seems plausible to conclude that subjects determined their order quantities based on their inventory balance – underweighting the supply line (namely the units in transit from the slow option) as in other settings (Croson and Donohue, 2006) – combined with random errors (Su, 2008). This would explain the standard deviations of the on-hand inventory levels before demand and also the negative correlations between ordered units and the inventory balance.

We saw that subjects' on-hand inventory levels before demand realization were heterogeneous. In addition, they were significantly lower than the predicted 10 or 11 in each treatment (a series of one-sample Wilcoxon tests,

[2] We also computed optimal values with the value iteration algorithm to check our results: This optimization was performed over an infinite horizon with the convergence criterion of 0.00001 on the total cost. The value iteration algorythm was used in the literature (e.g. Munos, R., & Szepesvari, C. (2008). Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, *9*, 815-857.) to optimize finite horizon problems The recorded differences between the two approaches never exceeded 0.6%.

[3] The following values are means with respective standard deviations in parentheses.

[4] Later we will see that this is driven by the different usage of the fast and slow option.

[5] Note that we conducted this type of analysis to claim that the differences we report cannot be attributed to learning or end game effects. We conducted this robustness check for all the upcoming analyses as well, but the tests yield the same results and conclusions in each and every case as without truncating these periods.

[6] Following a dual-index policy can be ruled out with the argument that negative the correlations between ordered amounts and the actual on-hand inventory was higher and significant, whereas they were lower and insignificant after taking the units into account that are in transit from the slow option in line with dual index policy.

p<0.02), however, they were not statistically distinguishable from each other (a series of MWW-tests where p>0.164). It is important to rule out that subjects had a lower inventory in general since they did not recognize that it is optimal to order more units and keep more inventory than the potential maximum demand. Out of our 86 subjects 38 had a mean inventory level above 8 across the 48 periods before the demand was revealed. Moreover, there was no single person who did not have an inventory of 8 at least for a single instance before the demand in that particular period was revealed. All in all, it seems plausible that subjects had a conscious preference for holding less items on stock. One explanation for this behavioral pattern could be that subjects ignored or misjudged the future and only did a myopic optimization as in other studies (Thaler, Tversky, Kahneman, and Schwartz, 1997; Appel and Musshoff, 2011) taking the current situation and one single period into account. An alternative explanation for this is the widely-documented ex-post inventory error minimizing behavior (Kremer et al., 2013), which we will explain by psychological costs later. These explanations seem even more pronounced since we saw no significant correlation between subjects' inventory level before demand realization and their elicited risk attitude, loss aversion parameter, cognitive ability or regret preferences.

Next, we examined whether the fast and the slow options had been overused by our subjects. Table 1 contains the theoretically optimal and the actual usage ratios for the slow option. Looking at the columns 'Actual usage' and 'Model 1' at first glance, we see that the slow option is underused – thus the fast one is overused – in the treatments which have the same framing for both supply modes compared to what the dynamic programming model predicts (a series of one-sample Wilcoxon-tests, p<0.02). However, if we take subjects pronounced preference for lower on-hand inventory into account, a different picture emerges. We calculated the optimal ratios taking the actual on-hand inventory of subjects into account as we saw that subjects had a general preference to have less inventory. This means that they were more often in a state when they had less inventory on-hand compared to what Model 1 assumes. Thus, Model 2 takes this preference into account and yields to higher usage for the fast option. The lower the inventory level, the more frequently subjects should order from the fast mode.

Table 1: Actual and model-predicted usage proportions of the slow option

Treatment	Actual usage	Model 1	Model 2	Model 3	Model 4
ST	100.00% (0.00%)	100.00%	100.00%	100.00%	-
D7T	59.37% (20.72%)	72.22%	55.80%	57.74%	-
D10T	89.82% (11.65%)	100%	92.72%	83.67%	-
D7TF	74.12% (13.83%)	-	-	-	72.22%

Notes: This table provides data on the usage ratio of the slow option relative to all units ordered. Model 1 denotes results based on the normative dynamic programming model. Model 2 means the same model taking the actual inventory levels and thus the different on-hand inventory levels into account. Model 3 abbreviates the results obtained by the dynamic programming model taking additional psychological costs into account. The column 'Actual usage' contains the actual mean ratio and its standard deviations in parentheses. Model 4 incorporates the same psycholoigcal costs as Model 3 and also additional ones for external suppliers when internal options are also present.

We see that there is no significant difference between the actual and the theoretically optimal usage (a series of one-sample Wilcoxon-tests, p>0.41) if we take these corresponding subgame-perfect strategies in 'Model 2' into account. Moreover, there is a plausible explanation for this behavioral pattern: Subjects follow reference-dependent preferences and attribute different additional psychological costs to leftovers and backorders (Ho et al., 2010). 'Model 3' incorporates exactly the same psychological cost parameters[7] as they estimated in their model. We see that there is no statistically significant difference (a series of one-sample Wilcoxon-tests, p>0.23) between the actual and the predicted usage in the treatments where both options are framed similarly. Moreover, accounting for these different psychological costs is in line with the exhibited ex-post inventory error minimizing behavior, since

[7] Note that we used exactly the same parameters as Ho et al. (2010) estimated since our study used the same shortage and leftover costs as they did. This means a cost of 9.96 per unit for leftovers and 6.52 per unit for shortages. We do a complete parametric analysis in the full version of our paper with a two-dimensional graph. We vary these psychological cost parameters between 0% and 200% of the ones reported by Ho et al. (2010). Generally, we see that the ratio of ordered amounts from the slow option is relatively stable between broad boundaries and increases/decreases significantly stepwise after particular thresholds. This phenomenon is mostly attributed to the fact that the optimal ordering policy changes at these thresholds.

the optimal policy changes to (0;8) for treatments ST and D10T and to (-2;6) for treatments D7T and D7TF, implying that there is no significant difference between the theoretically optimal order up to levels and the ones we actually observed (a series of one-sample Wilcoxon-tests, $p>0.15$). Finally, we see that framing matters a lot, since the usage proportions for the slow option in D7T and D7TF are significantly different from each other (MWW-test, $p=0.02$). This bias can be explained by the extension of the psychological cost model. In addition to the psychological costs for shortages and leftovers, we assume that subjects attach an additional psychological cost to the external supplier if both an external and an internal vendor is available. We denote this as 'Model 4' which predicts a usage proportion of 72.22% for the slow option as long as this additional psychological cost falls between 25.70% and 41.42% of the original procurement price[8]. The optimal ordering policy changes outside this range which leads to higher and lower usage proportions. The actual usage in our treatment was statistically indistinguishable from what 'Model 4' estimated: 74.12% (one-sample Wilcoxon-test, $p=0.41$). Note that the range we specified for the additional psychological cost is in line with the total cost of ownership approach – namely that assuming the same unit procurement cost for both options, sourcing from an external supplier has higher additional costs attached to it than an internal one.

We conducted a robustness check after calculating the average usage ratios for each treatment. We calculated the normatively optimal usage ratio at the individual level for the slow and fast option for each subject taking their actual on-hand inventories and the corresponding frequencies into account. In addition, we computed the difference between the actual usage ratio and the theoretically optimal individual one. We could not document any correlations between either of these values and the preferences we elicited during the experiment, which further supports our arguments.

Conclusion

This paper documents that subjects have a preference for overusing the fast and more expensive order option in a dual sourcing setting. Most importantly, we saw that adding a faster supply mode to the supply network leads to significant cost reductions even if this additional option costs 66.67% more than the slightly slower supplier. We showed that the overusage of the fast option is in line with the corresponding subgame-perfect ordering strategy and that it can be explained by the preference for ex post inventory error minimization. Additionally, we saw that the preference for less on-hand inventory coincides with the reference-dependent model that incorporates psychological costs. Furthermore, we document that ownership structure has a non-negligible effect on supplier usage as there is a general biased tendency towards using an internal supplier instead of an external partner. In the end, this leads us to the conclusion that sourcing strategy as well in/outsourcing decisions are susceptible to behavioral causes, not only to cost considerations.

References

Appel, F., & Musshoff, O. (2011). How appropriate are myopic optimization models to predict decision behaviour - a comparison between agent-based models and business management games. *EAAE 2011 Congress Paper*.

Bendoly, E., Croson, R., Goncalves, P., & Schultz, K. (2010). Bodies of knowledge for research in behavioral operations. *Production and Operations Management*, *19*(4), 434-452.

Bendoly, E., Donohue, K., & Schultz, K. L. (2006). Behavior in operations management: Assessing recent findings and revisiting old assumptions. *Journal of Operations Management*, *24*(6), 737-752.

Boudreau, J. W. (2004). Organizational behavior, strategy, performance, and design in management science. *Management Science*, *50*(11), 1463-1476.

Burke, G. J., Carrillo, J. E., & Vakharia, A. J. (2007). Single versus multiple supplier sourcing strategies. *European Journal of Operational Research*, *182*(1), 95-112.

Croson, R., & Donohue, K. (2006). Behavioral causes of the bullwhip effect and the observed value of inventory information. *Management Science*, *52*(3), 323-336.

[8] As we mentioned before a comprehensive parametric analysis was also conducted and reported in the full version of the paper.

Croson, R., Donohue, K., Katok, E., & Sterman, J. (2013). Order stability in supply chains: Coordination risk and the role of coordination stock. *Production and Operations Management*, *forthcoming*.

Cubitt, R. P., Starmer, C., & Sugden, R. (1998). On the validity of the random lottery incentive system. *Experimental Economics*, *1*(2), 115-131.

de Vericourt, F., Jain, K., Bearden, J. N., & Filipowicz, A. (2013). Sex, risk and the newsvendor. *Journal of Operations Management*, *31*(1-2), 86-92.

Drichoutis, A., & Lusk, J. (2012). *Risk preference elicitation without the confounding effect of probability weighting.* MPRA Paper, No. 37776.

Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, *19*(4), 25-42.

Gächter, S., Johnson, E. J., & Herrmann, A. (2010). *Individual-level loss aversion in riskless and risky choices.* IZA Discussion Papers 2961, Institute for the Study of Labor (IZA).

Ho, T. H., Lim, N., & Cui, T. H. (2010). Reference dependence in multilocation newsvendor models: A structural analysis. *Management Science*, *56*(11), 1891-1910.

Klosterhalfen, S., Kiesmuller, G., & Minner, S. (2011). A comparison of the constant-order and dual-index policy for dual sourcing. *International Journal of Production Economics*, *133*(1), 302-311.

Kremer, M., & Minner, S. (2008). The human element in inventory decision making under uncertainty: A review of experimental evidence in the newsvendor model. *Zeitschrift für Betriebswirtschaft*, *4*, 83-97.

Kremer, M., Minner, S., & Van Wassenhove, L. N. (2013). On the preference to avoid ex-post inventory errors. *Production and Operations Management*, *forthcoming*.

Lee, H. L., & Billington, C. (1993). Material management in decentralized supply chains. *Operations Research*, *41*(5), 835-847.

Lyon, T. P. (2006). Does dual sourcing lower procurement costs? *Journal of Industrial Economics*, *54*(2), 223-252.

Minner, S. (2003). Multi-supplier inventory models in supply chains: A review. *International Journal of Production Economics*, *81-82*, 265-279.

Moritz, B. B., Hill, A. V., & Donohue, K. L. (2013). Individual differences in the newsvendor problem: Behavior and cognitive reflection. *Journal of Operations Management*, *31*(1-2), 72-85.

Munos, R., & Szepesvari, C. (2008). Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, *9*, 815-857.

Raeva, D., Mittone, L., & Schwarzbach, J. (2010). Regret now, take it now: On the role of experienced regret on intertemporal choice. *Journal of Economic Psychology*, *31*(4), 634-642.

Ramasesh, R. V., Ord, J. K., & Hayya, J. C. (1993). Note: Dual sourcing with nonidentical suppliers. *Naval Research Logistics*, *40*(2), 279-288.

Schweitzer, M. E., & 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. M. (2008). Bounded rationality in newsvendor models. *Manufacturing & Service Operations Management*, *10*(4), 566-589.

Thaler, R. H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *Quarterly Journal of Economics*, *112*(2), 647-661.

Veeraraghavan, S., & Scheller-Wolf, A. (2008). Now or later: A simple policy for effective dual sourcing in capacitated systems. *Operations Research*, *56*(4), 850-864.

Whittemore, A. S., & Saunders, S. C. (1977). Optimal inventory under stochastic demand with 2 supply options. *Siam Journal on Applied Mathematics*, *32*(2), 293-305.