Decision Patterns and Information Availability in the Beer Distribution Game

Jonas Kunze, Yalın Gündüz
Universität Karlsruhe(TH), Karlsruhe, Germany
Contact author: jonas.kunze@em.uni-karlsruhe.de

July 7, 2006

Abstract

In this article we present results of different experiments with the Beer Distribution Game focusing on decision patterns and the effect of varying information on the decision quality. Besides the known decision patterns such as the ordering heuristic presented by Sterman (1989) and the well investigated bullwhip effect in the Beer Distribution Game, we make two other observations. First, as an extension to available studies, we suggest that decision behavior could be explained by policies that change over time. Second, a non linear relationship between the anchor and the decision is presented and contrasted to the linear heuristic. Information seems to play an important role in the decision making process, but the effect is not necessarily positive. This could be explained by coordination problems. Overall, the complexity of the Beer Distribution Game raises various questions about the experimental design.

1 Introduction

In the late 1950s, in order to introduce the mechanics of System Dynamics in a role-playing setting, the Beer Distribution Game (BDG) was developed.This is a simulation game where each of the players is responsible for meeting the incoming orders in a supply chain. The BDG has proven to be an important tool in the spread of System Dynamics methodology, influencing many researchers and executives (Sterman (2000)). Especially for researchers, it provided an important experimental setup for understanding, analyzing and simulating dynamic systems (Sterman (1989), Sterman (1992); Senge (1994)). There have been alternative application areas of

BDG where it has been used as a complex problem in artificial intelligence (Geyer-Schulz (1998)) and as a basis for different experiments (Croson and Donohue (2002)).

The Beer Distribution Game has a simple and nice system dynamics model setup where the orders flow through the supply line. The inventories and the supply lines are the *stocks* of the system. As in all System Dynamics models, these stocks show a quantifiable level of either a physical or non-physical variable, such as inventory, cash, headcount, or customer satisfaction. In the BDG setting, the variable of interest for the players is the inventory and supply line stocks, which are ideally tried to be minimized because of inventory and backlog costs. The dynamic structure of the System Dynamics model is achieved through the *flows* in the system; while stocks represent the state of system elements, flows show how things are changing. In this sense flows revise and update the stock values: they can denote actions such as producing, hiring/firing, selling, or delivering. In the BDG setup, flows in the system refer basically to incoming order flow, and outgoing shipment flow[1]. The players therefore maintain system balancing behavior through control of the supply line and inventory (Sterman (2000), pp.684-694).

Overall, the System Dynamics Modeling of the real world systems provides a powerful way of representing a problem faced everyday: In a dynamic decision making setting, typically, decisions are taken in order to manage/regulate the level of stocks (i.e. inventory) through affecting the inflows and outflows to the stocks per unit time (i.e. producing, delivering). The desired level of inventory and the desired supply line are auxiliary variables in the system, which are individually decided by the players in a dynamic environment. In this study we focus on decision patterns, and the effects of information availability on the basis of the experimental BDG setup. The dynamic decision making tasks that the players face in managing the stock levels are examined in the light of previous literature. First addressed by Tversky and Kahneman (1974), anchoring and adjustment has been an important behavioral policy to estimate an unknown quantity by a reference point. We observe that the system balancing behavior that the players aim to achieve might result in a non-linear relationship between an anchor and the ordering decision. Further, we show that ordering policies might change in time. Information availability has importance in decision making, but the effect is not necessarily positive, which can be explained through coordination problems.

[1]For a stock-flow diagram of the BDG see Sterman (1989).

The study is organized as follows: In section 2 the BDG is described in detail. In section 3 the experiments and the obtained datasets are described, while section 4 focuses on the analysis of the datasets. Section 5 gives a short outlook.

2 The Beer Distribution Game

The BDG was developed as a board game. Usually, four players form a group and are members of a linear supply chain (retailer, wholesaler, distributor, and factory). There exists only a unique good (here: beer kegs), which is shipped downstream, i.e. from the factory, over the distributor and wholesaler, to the retailer. The number of goods shipped depends on the previous order that was placed by the downstream member, i.e. there exists an upstream flow of orders. All supply chain members have a stock (inventory) where goods are stored until they are used to fulfill an incoming order.

On the market side of the supply chain, there is a customer that orders a unique good from the retailer. These orders are a predefined function, which we will call the *demand function*. On the supply side of the supply chain, there is a source with a 100% service level which serves all orders that were placed by the factory. Sometimes, the factory has a production limit. This is put into the play by announcing that the factory is only able to order up to a fixed number of goods.

The transmission and processing of an order as well as the shipment of the ordered goods takes time. By using a discrete time structure, rounds are introduced into the game. Generally, each order needs two rounds to reach the supplier (next upstream member) and another two rounds are needed until the shipment arrives in the stock. Every round, one non-negative order decision must be taken by every player in the supply chain. For more details on the exact state transition, see Sterman (1989).

Up to now, the structure of the BDG has been described. But another question is essential for decision making: What is the information to which players have access? If the word "information" is taken in a broad sense, it includes information about the structure (number of team members, size of time lag, etc.), about the system state (number of goods in retailer's stock, number of goods ordered by the wholesaler two rounds ago, etc.), about the demand function (stochastic information, actual realization), and about the other players' decision making (communication allowance). Generally, structural information is provided. In the traditional board game, beer kegs were symbolized by little counters. If played on one table, shipment information became public knowledge. Orders, to the contrary, were symbolized by cards where a number was written. This information was hidden until the order reached the next upstream team member. Communication was usually forbidden.

Before using the BDG in an experimental setting for the current study, the structural and informational characteristics of a certain BDG was established. Furthermore, some terminology was changed to achieve a more neutral view on the decision task. First, the "beer" terminology was taken out. Instead of beer kegs, only goods were shipped. Second, the industrial production and distribution connotation was excluded by renaming retailer, wholesaler, distributor, and factory to dealer 1, dealer 2, dealer 3, and dealer 4.

When talking about the experiments, we use the following terminology (after Davis and Holt (1993)): A participant of the experiment is called *subject*. In a *session*, a group of subjects plays the game. Each session has a certain *treatment* describing a specific configuration of experimental variables, such as information and incentives.

3 Datasets

In this section we describe two different datasets that have been acquired in experiments with the BDG at the Universität Karlsruhe (TH).

The first dataset results from a preliminary experiment that was conducted in February and March 2004. 60 students participated in a total of 6 sessions. The decision task was to act in a supply chain with two players forming one group, and the demand function and the system initialization was as described in Sterman (1989). Two repetitions of the game were played with 30 rounds each.

The second dataset is from a lab experiment that was conducted in March 2005 and in which 87 students participated in 6 sessions. Some experimental design changes were made in order to achieve a more generalizable result. Each supply chain consisted of 3 subjects. Furthermore, the demand function was now a uniform distributed process between 0 and 10 ($U[0, 10]$) and stochastic information was provided to the subjects. After a training game of 8 rounds, the main game was played and lasted at least 30 rounds. Beginning with round 30, there was a 20% probability that the actual round was the last round. This method was chosen to eliminate end game effects. Therefore, the number of rounds in the sessions differs between

Dataset	1	2
From	Preliminary Experiment	Lab Experiment
Group size	2	3
Groups	30	29
Information	Local vs. Classic	Local vs. Classic
Rounds	2x 30 rounds	Training: 8 r. 30-37 rounds

Table 1: Characteristics of the datasets

30 and 37.

The research questions of the experimental studies focused on the informational availability (Table 1). The first dataset's experiment had two treatments which differed in the information about the system state. In one treatment ("local information") only information about the actual stock, the incoming orders and the incoming shipment were presented. The other treatment ("classic information") provided the information that could be gathered by players sitting at a table, where beer kegs are symbolized by little boxes that were moved around. Therefore, all shipments as well as the stock of all players were public knowledge. The experiment of dataset 2 was an extension of the preliminary experiment and also focused on the comparison between "local" and "classic" information.

It has to be noted that experimental control varied across the experimental studies. In the preliminary study, only small incentives were present and the end game effect was not addressed. In the lab experiment, significant incentives were present, intelligibility of the instructions and the screen layout were improved.

4 Analysis

In this section we first describe observed patterns in the decisions and discuss the use of an anchored linear order heuristic to describe these patterns. Then, we compare the results of the conducted experiments and discuss a coordination problem hypothesis to explain the results.

4.1 Patterns in Decisions

The existing studies focus on total transaction costs to compare the performance of different supply chains (Steckel et al. (2004), Kaminsky and Simchi-Levi (1998)) or they adopt an ordering heuristic to identify behavioral patterns (Sterman (1989), Croson and Donohue (2003), Croson et al. (2004)). In the following analysis of the dataset from the lab experiment we make two observations that call for an extension of the ordering heuristic presented in Sterman (1989).

The ordering heuristic by Sterman (1989) assumes that the decision maker orders the non-negative sum of three variables. These are namely, (i) the expected loss (expected incoming order), (ii) the difference between desired and actual stock, and (iii) the difference between desired and actual supply line. To describe the decision maker's behavior the author introduces some parameters to the ordering heuristic that are fitted to explain the decisions in the best possible way. Two inherent assumptions of this way to describe decision making are *consistency* and *linearity*[2]. Consistency denotes that the personal parameters in the heuristic are kept constant for all rounds, whereas the linearity assumption indicates that the order decision is the weighted sum of the three mentioned variables and that an increase in one of these would lead to a corresponding increase in the order decision.

4.1.1 Consistency in Decisions

Observation of subjects' decision making leads to the impression that subjects do not make their orders in a consistent way during the game. It rather seems that decision making could be well described by order policies that change from time to time. Although we will not formalize these thoughts sufficiently to make statistical significance tests, we present the decision making of 5 subjects from the lab experiment. In the following figures, for 29 rounds, the decision, the incoming order, and the decision time for each round is depicted. For the selected cases, these variables give a good impression about possible order policies that might have been employed by the subject.

Subject 7 In Figure 1 the decision making for subject 7 is depicted. From rounds 1 to 6 the decisions are always lower than the incoming orders. From round 8 on, the decisions are basically 5, except for rounds 13 and 14. Furthermore, the decision time is depicted and shows some outliers at rounds

[2] In the cases when the ordering heuristic predicts a positive number.

4 and 11. Interestingly, the noticed changes in decision behavior occur some rounds after these outliers in decision time.

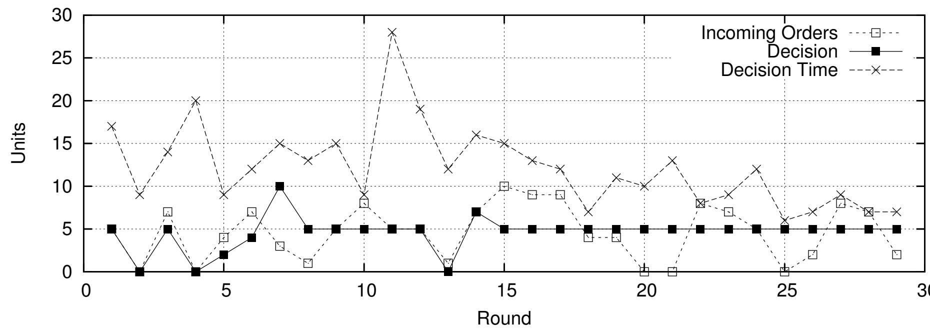

Figure 1: Decisions of Subject 7 (Lab Experiment)

Subject 17 The behavior of subject 17 is shown in Figure 2. It can be divided into two parts. In the first half of the rounds, a constant number of goods is ordered for several rounds. When this number changes, long decision times occur (round 3 and 8). From round 13 on, no clear order policy can be concluded from visual inspection of the presented variables. But if we include stock development, it can be seen that the over-ordering until round 13 resulted in an excessive stock of at least 30 goods from round 13 on. Therefore, the missing clear order policy could be explained by a lost interest in the game from round 13 on. The low decision times support this.

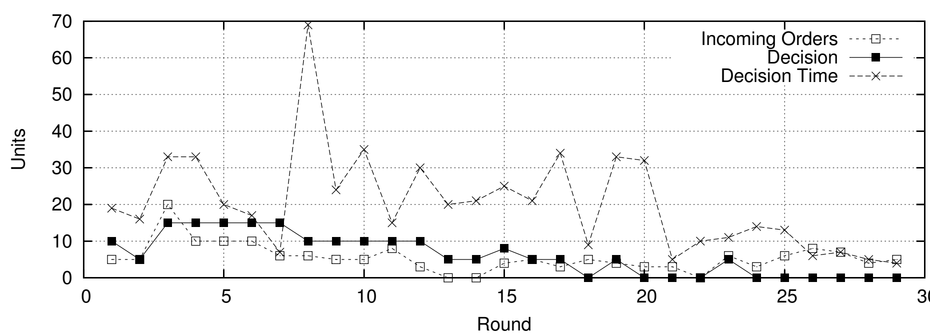

Figure 2: Decisions of Subject 17 (Lab Experiment)

Subject 18 In Figure 3 the time saving aspect of a simple strategy can be seen. In 18 out of 29 decisions, the subject orders exactly the number that has been ordered by his downstream group member (incoming orders). The decision time is extraordinarily low.

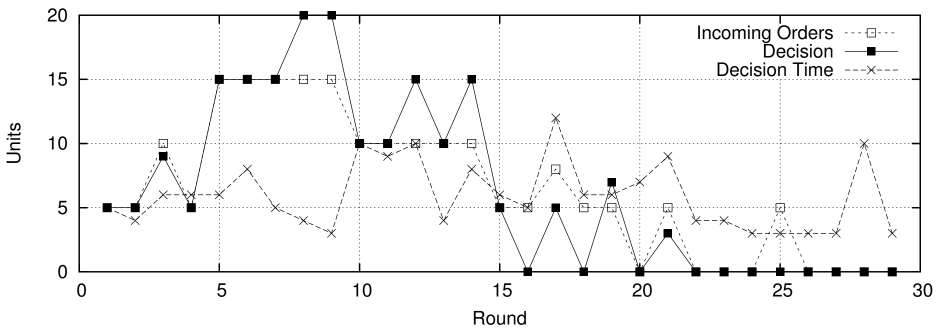

Figure 3: Decisions of Subject 18 (Lab Experiment)

Subject 44 In Figure 4, a clear scheme is observed for the first 15 rounds. Every time the decision taken differs from the number of incoming orders, there occurs a peak in the subjects decision time. From round 15 on, this pattern is lost.

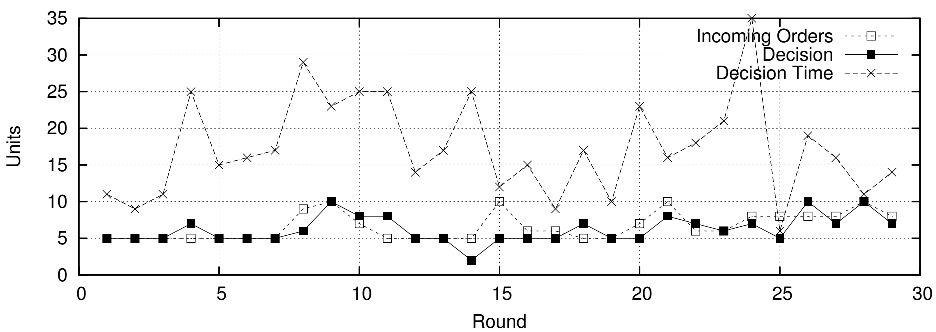

Figure 4: Decisions of Subject 44 (Lab Experiment)

Subject 82 For subject 82 (Figure 5) we have a similar finding as for subject 44. After a long planning phase (long decision times in the first 3 rounds), the subjects continuously orders 5. There is one exception in round 12 accompanied by a long decision time and some changes in the last rounds.

These changes might be motivated by the stock development, as the stock is larger than 15 units from round 26 on.

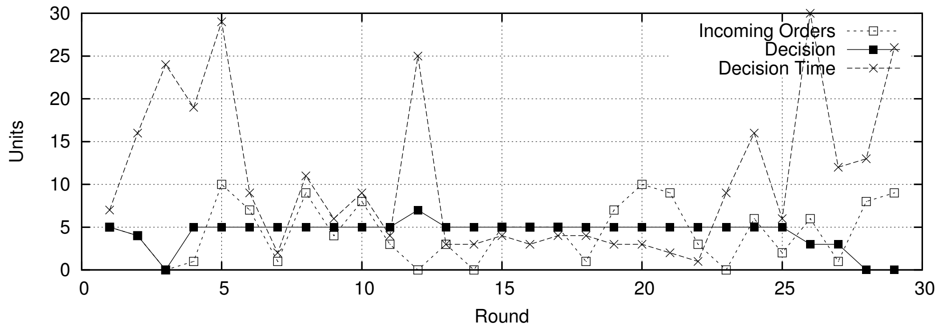

Figure 5: Decisions of Subject 82 (Lab Experiment)

4.1.2 Linearity Assumption

If linearity holds, then a linear relationship between the incoming orders (loss rate) and the decision is reasonable. To investigate this relationship we looked at dataset 2. First of all, we checked what is the best linear trend that could be found. We ran a linear regression with a fixed point in (0,0) and minimized squared errors. The gradient with minimum square errors is 0.825. This observation suggests that an incoming order of 5, would result in an average order decision of 5 x 0.825. The analysis below has been undertaken for both 0.825 and 1, as the perfect linear slope. However, as the results did not change significantly, only the section with the linear slope of 1 is presented.

We then counted the total amount of (incoming order,decision) pairs for every possible combination. For instance, if one subject in an arbitrary session had an incoming order of 6 and decided to order 8, we increased the count for the pair $(6, 8)$ by one. The results are the absolute frequencies of these pairs and they are depicted in Figure 6. In total there are 2,491 pairs.

Visual inspection shows that two squares of high frequency exists: A (0-5,0-5) square and a (5-10,5-10) square. In contrast to a linear relationship between incoming orders and decisions, we introduce the hypothesis that specific points (*anchor points*) have a special attractiveness in the decision making process. These anchor points are observed to be 0, 5, and 10.

To investigate our hypothesis we use the chi-square goodness of fit test. To apply the test, we need to define probability distributions that gives each

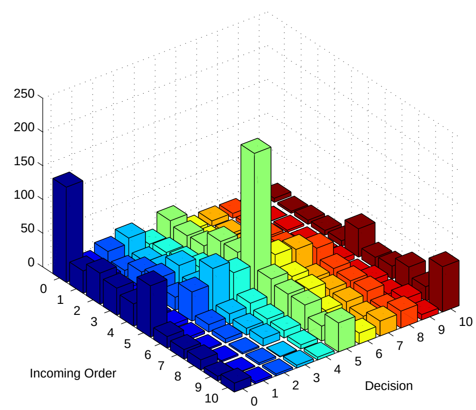

Figure 6: Frequency of Pairs of Incoming Orders - Decision Pairs

incoming order-decision pair a probability. As the decision is taken after having received the incoming order, we only need to define a conditional probability.

For the linear relationship, we assume that the decision is normally distributed. The mean value μ is equal to the incoming order and the standard deviation σ will be estimated from the sample. For each incoming order from 0 to 10 we have run a chi-square test. The results are depicted in Table 2 and show that the observed data does not fit the normal distribution. For any incoming order the normal distribution assumption has to be rejected, even if the confidence interval is 99.99%.

To test our hypothesis, we need to construct a probability function that puts special weight on the anchor points 0, 5, and 10. Therefore, we propose the following probability distribution to observe an anchor $a \in \{0, 5, 10\}$, given the incoming order to be x:

$$\text{prob}_{\text{pw}}(a|x) \begin{cases} \frac{5}{7} & \text{if } |a-x| \leq 1 \\ \frac{5}{7} - \frac{1}{7}(|a-x|-1) & \text{if } 1 < |a-x| \text{ and } |a-x| \leq 6 \\ 0 & \text{else} \end{cases}$$

The resulting distribution for the three anchor points is depicted in Figure 7. Using this probability distribution to explain the data would not work, as it only puts probability weights on the anchor points. Therefore, we use a linear combination of the normal distribution and the anchor points distribution. For each linear weight from 0 to 0.5 we have run chi-square

Incoming Orders	Number of observations	σ (est.)	Categories	χ^2 distributed test statistic
0	276	3.369	9	219.384
1	123	3.406	8	43.562
2	159	2.957	9	85.569
3	215	2.548	9	93.290
4	205	2.590	10	103.838
5	636	2.639	10	470.192
6	206	2.732	9	41.604
7	183	2.946	9	39.924
8	175	3.499	9	68.329
9	105	3.699	8	30.157
10	208	4.224	9	114.368

Table 2: Results of Chi-Square Test (Normal Distribution)

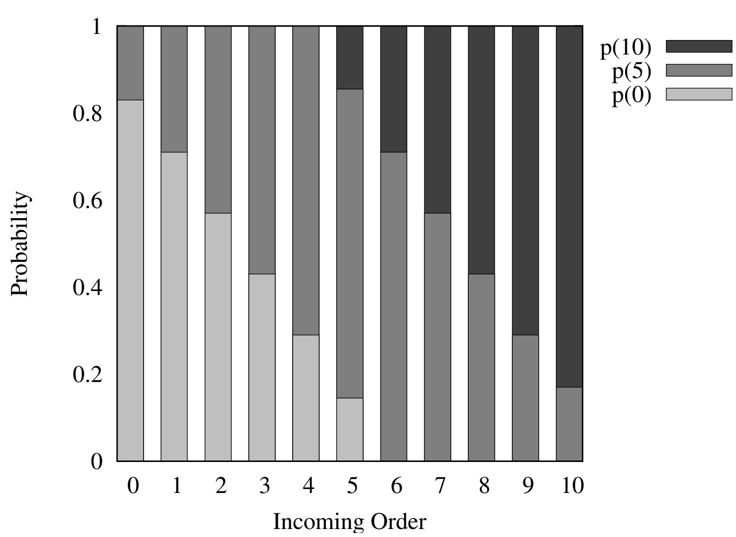

Figure 7: Probability to Observe an Anchor Point

tests for all possible incoming orders, and calculated the average p-values for the rejection of the hypothesis that the observations follow the mixed distribution.The results for the different linear weights are depicted in Figure 8. Note that the p-value for every chi-square test assuming the normal distribution is lower than 0.001%. The highest average p-value, is found if a

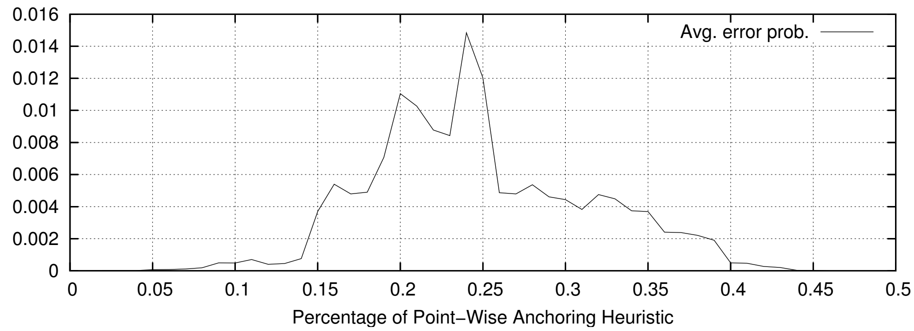

Figure 8: Average Error Probability for Different Linear Combinations

linear weight of 24% is given on the anchor points distribution. The results of the chi square test of this mixed distribution are presented in Table 3. Two of the chi-square test statistics(incoming order being 1 and 9) have a p-value that is higher than 5%, and another test statistic(incoming order equal to 4) has a p-value higher than 1%. These findings show that the fitting quality is increased by including the point-wise distribution that only puts probability weight on the anchors.

Another observation that puts emphasis on the importance of the anchor and multiples of it is found in dataset 1. The game initializes with 12 units in each stock and 4 units have been ordered in the past 4 rounds. The demand function rises from 4 to 8 units in the fifth round. Therefore, especially even numbers and multiples of 4 might have a special attractiveness for decision makers. In Figure 9 the frequencies of each decision taken is depicted.

To summarize this section, we investigated two characteristics of decision making in the BDG. First, decision making is not always consistent during the whole game. Changing situations may lead to a complete change in the order policy. Often, these changes are preceded by long decision time. Second, experimental setup suggests anchor points. These anchor points have a special attractiveness for decision makers.

Incoming Orders	Number of observations	σ (est.)	Categories	χ^2 distributed test statistic	p-value [%]
0	276	3.369	8	74.781	<0.0001
1	123	3.406	8	12.183	9.4705
2	159	2.957	8	27.864	0.0232
3	215	2.548	9	62.710	<0.0001
4	205	2.590	9	20.052	1.0140
5	636	2.639	10	95.238	<0.0001
6	206	2.732	9	48.916	<0.0001
7	183	2.946	8	35.022	0.0011
8	175	3.499	8	24.621	0.0885
9	105	3.699	7	12.219	5.7268
10	208	4.224	9	63.629	<0.0001

Table 3: Results of Chi-Square Test (Mixed Distribution)

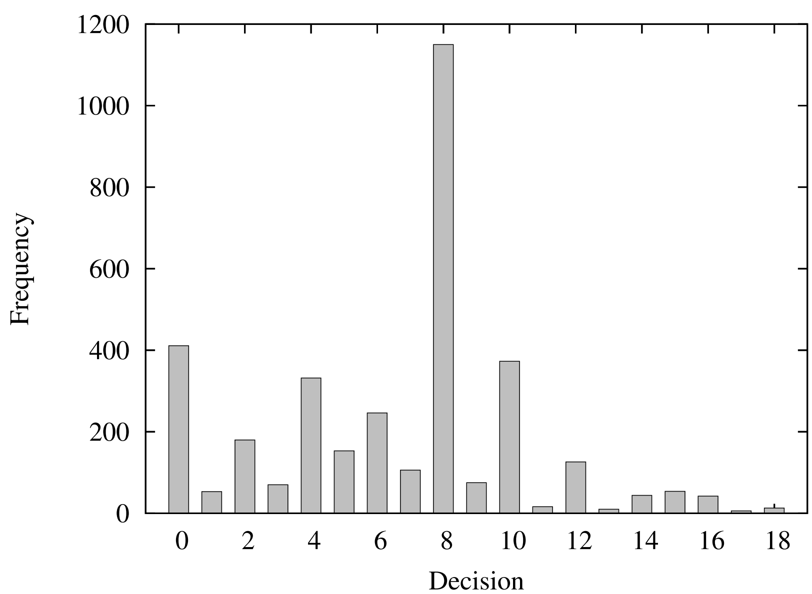

Figure 9: Frequencies of Decisions (Dataset 1)

4.2 Effect of Changes in Information Availability

To investigate the effect of the different information availability settings (*classic* vs. *local*), we make use of both datasets. In Figure 10 the distribution of the group costs for each treatment are depicted. These are the sum of group costs for the entire time horizon. It can be seen that the group costs of the treatment with local information are lower. For a more precise discussion of the results see Kunze (2004).

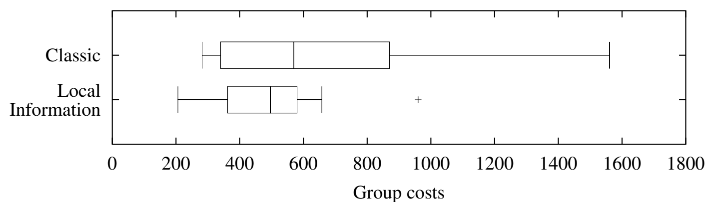

Figure 10: Group Costs in the Preliminary Experiment

The interesting finding that less information about the system state leads to better results, may be explained by an overall *coordination problem*. In order to achieve a supply chain with low transaction costs, participants need to coordinate their strategies. Coordination in the BDG is difficult to achieve as communication is prohibited. Therefore, decision makers look at the signals they receive from their group members (e.g. incoming orders from the downstream player, service level of the upstream player). With these signals decision makers choose a strategy that hopefully work well with the other players' strategies. In the *classical* treatment, a lot of shipment information is available. Therefore, the decision maker receives more signals and his strategy selection task gets more complex. In other words, more information allows more strategies and therefore makes the coordination problem more complex. On the other hand, more information will surely provide a better understanding about the system state and a more precise guess about the other group members' strategies.

To see whether this interesting finding holds for a supply chain with groups of size three, the lab experiment was conducted. For this experiment, the demand function was a uniform process between 0 and 10. In order to compare the two treatments, one realization of the random process was used two times, once in each of the two treatments. For six demand realizations, the group in the local information treatment produced lower costs than the group in the classic information treatment, while the opposite was true for the remaining eight demand realizations. Although there appears to be a

Treatment Session	1 1	1 3	1 5	2 2	2 4	2 6
Avg. round time [sec]	54.1	52.7	53.4	33.1	38.1	39.8
Std. deviation	21.8	15.7	15.7	10.2	9.4	11.1
Avg. dec. time [sec]	23.6	24.6	23.7	16.0	17.9	19.2
Std. deviation	15.8	14.1	17.5	10.1	11.3	11.4

Table 4: Round and decision times for the lab experiment

contrary trend to the results of the preliminary experiment, a pairwise test does not show significant differences between the treatments.

There are at two potential explanations to the fact that the groups in the classical treatment did not perform as bad as they did in the preliminary experiment. First, the *coordination problem* hypothesis could be wrong and there is actually no effect of the different informational patterns on decision making behavior. Second, the hypothesis could be right but the changes in the experimental setup counterbalance or interact with the coordination problem in an unforeseen way. A particularly strong possibility is that the different realizations of the demand function might have produced run length effects.

The first explanation is rendered unlikely by the fact that although group costs do not show a significant difference, the decision time of the subjects does. In Table 4 the round and decision times are depicted for dataset 2. Decision time is the time that a subject needs to type in his or her order decision and to verify it by clicking on the "Okay" button. Rounds were synchronized across teams, i.e. a round finished when all subjects of the session made their decisions. Therefore, round time is the time that all subjects had the actual system state information on their screen. Round time as well as for decision time are significantly higher in the classical treatment than in the local information treatment. This observation supports the intuition that the classical treatment demands more cognitive capacities as it offers more information to the decision maker.

The second possible explanation, which is the counterbalancing effect, seems more realistic, but two questions arise that have not been investigated: Is the effect of the treatment variable (information) independent of the effect of the realization of the demand function? Is the effect of the treatment variable independent of the group size?

It can be concluded that information about the system state likely has an influence on the decision making process. If more information is available, the decision making task is more complex and a higher decision time is needed. The unintuitive finding of the preliminary experiment, i.e. that groups with more knowledge about the system state perform worse than groups with less knowledge, was explained by the *coordination problem* hypothesis, but could not be replicated in the lab experiment. It remains unclear whether this is due to the changes in experimental setup.

5 Conclusion

Experiments with the BDG reveal interesting facts about the decision making process. Looking at overall patterns in the decision, we made two observations: First, people seem to change their decision making process from time to time. Second, if an anchor is present, certain anchor points have a special attractiveness for decision makers.

Two experiments were used to evaluate the influence of information about the system state on group performance. The results of the two experiments are contradictory, pointing to the importance of experimental setup. There might be a possible interaction between various structural design aspects (such as group size and demand function) and informational design aspects. Furthermore, the interdependencies between decisions within a game complicate the issues even more.

These findings point to a general problem when doing experiments in dynamic decision problems. The possible complex interactions between structural, informational and dynamic aspects make it difficult to test a certain research question with the BDG. These interactions should be subject for further research in the field of experiments with dynamic decision making problems.

References

CROSON, R. and DONOHUE, K. (2002): Experimental Economics and Supply Chain Management. *Interfaces*, *32*(5), 74–82.

CROSON, R. and DONOHUE, K. (2003): Impact of point of sale (POS) data sharing on supply-chain management: An experimental approach. *Productions and Operations Management*, *12*(1), 1–11.

CROSON, R. and DONOHUE, K. and KATOK, E. and STERMAN, J. (2004): Order Stability in Supply Chains: Coordination Risk and the Role of Coordination Stock. Tech. rep., Massachusetts Institute of Technology, Engineering Systems Division.

DAVIS, D. D. and HOLT, C. A. (1993): Experimental Economics. Princeton University Press.

GEYER-SCHULZ, A. (1998): Fuzzy Genetic Algorithms. In: H. T. Nguyen and M. Sugeno (eds.), *Fuzzy Systems: Modeling and Control*, The Handbooks of Fuzzy Sets Series, 403 – 460, Boston. Kluwer Academic Publishers.

KAMINSKY, P. and SIMCHI-LEVI, D. (1998): A new computerized beer game: A tool for teaching the value of integrated supply chain management. In: H. Lee and S. M. Ng (eds.), *Global Supply Chain and Technology Management*, vol. 1 of *POMS Series in Technology and Operations Management*, 216–225.

KUNZE, J. (2004): Einfluss von Informationsstrukturen auf das Entscheiderverhalten in einem Versorgungskettenexperiment. Master's thesis, Universität Karlsruhe (TH), Karlsruhe.

SENGE, P. (1994): The Fifth Discipline : The Art and Practice of the Learning Organization. Currency/Doubleday, New York.

STECKEL, J. H. and GUPTA, S. and BANERJI, A. (2004): Supply Chain Decision Making: Will Shorter Cycle Times and Shared Point-of-Sale Information Necessarily Help? *Management Science*, *50*(4), 458–464.

STERMAN, J. D. (1989): Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment. *Management Science*, *35*(3), 321–339.

STERMAN, J. D. (1992): Teaching Takes Off: Flight Simulators for Management Education. *OR/MS Today*.

STERMAN, J. D. (2000): Business Dynamics – Systems Thinking and Modeling for a Complex World. McGraw Hill.

TVERSKY, A. and KAHNEMAN, D. (1974): Judgement under Uncertainty: Heuristics and Biases. *Science, New Series*, *185*(4157), 1124–1131.