Emergence of a Multiple-Sourcing Strategy in a Buyer-Supplier Network: Effects of different Quantity-Quality and Quantity-Price Trade-Offs

Kristian Strmenik, Christian Mitsch, Friederike Wall, Gernot Mödritscher
Department of Management Control and Strategic Management
University of Klagenfurt, Klagenfurt, Austria
Email: {first.last}@aau.at

***Abstract*—In this paper, a buyer-supplier network is considered, which consists of a buyer and several suppliers who differ from each other in terms of quality and price. A buyer who puts its focus solely on quality pursues a different strategy than a price sensitive buyer, and, hence, allocates the procurement volume of a product in a different way among the suppliers, which in turn affects the supplier structure. Besides the buyer's strategic considerations, the suppliers also try to act strategically to maintain their competitiveness. We apply an agent-based simulation to analyze how different procurement volumes and levels of precision of the buyer's quality measurement system affect the supplier structure when (1) the suppliers' qualities and prices are modeled by generalized logistic functions and log-linear models, respectively, (2) the buyer uses a proportional volume allocation rule to allocate its procurement volume among the suppliers, and (3) the buyer learns its own quality-price preference via temporal difference learning. In order to express the buyer's quality-price preference, we apply an additive weighted sum model. The results show that, for low (high) procurement volumes, the buyer learns that sourcing from suppliers who pursue a high-quality (low-cost) leadership strategy leads to a more profitable supplier structure. But if the buyer's precision of quality measurement system decreases, these suppliers are not able to continue their position in the market and, therefore, lose market shares to suppliers who focus on a different competitive strategy.**

***Keywords*—Multiple-sourcing strategy; Buyer-supplier network; Volume allocation; Temporal difference learning.**

I. Introduction

It is not uncommon for a firm to source goods and services from two or more suppliers at the same time. The sourcing strategy of a firm can be described by three essential criteria [3]: (1) a criterion for establishing a supplier base, (2) a criterion for selecting suppliers who receive an order from the firm, and (3) a criterion for allocating the quantity of goods among the suppliers. In this paper, we assume that a firm (hereafter referred to as the buyer) pursues a multiple-sourcing strategy, which is characterized by a proportional volume allocation rule. This means that each supplier receives at least a part of the total procurement volume. In addition, we assume that the criterion for allocating the procurement volume is only based on the supplier's product quality and the price that the supplier charges for the product. Moreover, the procurement volume needs to be allocated in a trustful way, so that the buyer's expectations and preferences regarding quality and price are met [2]. Within that, the preference whether quality or price is more important depends on different factors like, for example, the buyer's industry, the buyer's strategic positioning and business model, and the importance of the purchased product.

On the other hand, also the suppliers try to act strategically and position themselves well in the market to earn high rates of return, even if the industry structure is unfavourable. Following [10], the basis for this in the long-term is a competitive advantage of the supplier, which may either stem from differentiation or low cost. Both strategies require a fundamentally particular path, including the choice about the type of competitive advantage and the scope of the strategic target in which the supplier wants to achieve a competitive advantage. In a multiple-sourcing strategy the employed suppliers may differ from each other in terms of their objectives and competitive strategy, and, hence, the quality and price they are able to offer for a certain procurement volume.

The paper focuses on a possible quantity-quality trade-off, which might exist between the supplier's quality and the requested quantity. This trade-off indicates the responsiveness of quality to changes in volume and implies that with an increased volume the supplier is not able to maintain its level of quality, which will subsequently drop to a lower level. A quantity-quality trade-off might stem from technological reasons, for example, increasing the operating speed to produce more pieces or using less-skilled workers to meet the higher demand. Ultimately, this may result in a lower quality. Besides this trade-off, we assume that depending on the allocated procurement volume, suppliers may offer the buyer different price reductions.

The main objective of this paper is to investigate *how different procurement volumes affect the buyer's supplier structure when (1) the suppliers are heterogeneous with respect to the above-mentioned quantity-quality and quantity-price trade-offs, (2) the buyer pursues a multiple-sourcing strategy, and (3) the buyer learns its own quality-price preference based on its supplier environment.*

In certain situations, the observed quality of a product may not match the agreed quality because, for example, it might result from an imperfect buyer quality perception. Therefore, we extend our research question to *how the model reacts*

Copyright (c) IARIA, 2021. ISBN: 978-1-61208-898-3

when imperfect quality is imperfectly measured by the buyer's quality measurement system.

To answer our research questions, an agent-based simulation is set up, which captures a buyer and three heterogeneous (in terms of quality and price) suppliers. In particular, we describe the suppliers' qualities and prices by generalized logistic functions and log-linear models, respectively. To express the buyer's preference regarding quality and price, we apply an additive weighted sum model and, in order to model the buyer's multiple-sourcing strategy, a proportional volume allocation rule is used. Last but not least, we model the buyer's learning process via temporal difference learning.

The remainder of our paper is organized as follows. In Section 2, we review the literature regarding to sourcing strategies and volume allocations. In Section 3, we introduce our agent-based model, explain the model specifications, and introduce the buyer's learning method for learning its quality-price preference. The parameter settings for our simulation experiments are explained in Section 4. In Section 5, we present and discuss our results. Section 6 contains concluding remarks and suggests possible directions for future research.

II. Related Research

Several studies have been conducted on the topic of volume allocation discussing the benefits of certain sourcing strategies and suggesting using different criteria in case of a multiple-sourcing strategy. [12] presents a model to optimize the allocation of volumes among suppliers by considering different cost factors. The authors conclude that, if the reliability of the suppliers is low, the buyer should consider a multiple-sourcing strategy. [3] proposes a supplier selection and volume allocation model where minimum order quantities and supplier capacities are considered. The authors find out that, if suppliers are incapacitated, the preferred strategy of the buyer is to source the product from multiple suppliers. However, the largest part of the required volume should be allocated to the least cost supplier and only marginal quantities to all other suppliers. [8] introduces a model of quality selection in an imperfectly competitive market considering quantity-quality trade-offs with constant values and studies its implications.In his findings, the author shows that the stronger the relationship between these two factors is, the more sales are shifted from the high to the low quality supplier [8]. [13] set up an agent-based simulation and take the assumption of [8] work to extend the literature on volume allocation, taking into account the impact of a non-linear trade-off between quantity and quality on the buyer's supplier structure. With their simulation experiment, they find out that, in cases where the buyer has to allocate large procurement volumes, a proportional volume allocation mechanism that only considers the suppliers' qualities leads to stronger oscillations of the supplier volumes. To mitigate this phenomenon, the buyer should form its expectations not only based on short-term perception, but on a more sophisticated method by allowing adaptive expectations. In addition, the authors are also considering additional indicators such as prices in order to stabilize the behaviour of the buyer's volume allocation.

III. The Model

A. Overview

We consider a buyer-supplier network, which is characterized by a buyer who is ordering the same procurement volume of a certain product in every time period and different suppliers who are offering the demanded product. Each supplier is characterized by a non-linear quantity-quality and quantity-price trade-off. Table I gives an overview of the before-mentioned trade-off relationships of suppliers, which we investigate in our model.

TABLE I
TRADE-OFF RELATIONSHIPS OF DIFFERENT SUPPLIER TYPES.

competitive strategy	quantity-quality trade-off	quantity-price trade-off
quality leadership	high	low
'stuck in the middle'	medium	medium
cost leadership	low	high

Moreover, we take into account the possibility that the buyer puts more emphasis either on quality or on price, or to consider both as equally important while allocating the procurement volume among the suppliers. To get a relation between quantity and quality, we apply a generalized logistic function, which has an S-shaped form. This function type corresponds to our before-mentioned assumption that with an increase of the quantity, the supplier is not able to maintain its level of quality, which will subsequently drop to a lower quality level. S-shaped functions are very flexible and have essential properties so that they are often used, e.g., in neutral network learning methods as an activation function [9] or in biological growth models for animal sciences and forestry [6]. In our paper, we model the correlation between quantity and price with a log-linear model. A log-linear model can be used to describe, for example, the cost reduction in manufacturing, specifically, in areas with repetitive procedures such as production plants (e.g., [1]). Further we incorporate price reductions depending on the allocated procurement volume.

Finally, to express the buyer's preference regarding quality and price, we apply an additive weighted sum model, which is commonly used in multi-attribute decision making [4][15]. This type of model easily allows the buyer to put more weight either on quality or on price, or to consider both as equally important. The weight parameter in the additive weighted sum model is determined via temporal difference learning, a progressive learning process that comes from the area of reinforcement learning. With this type of learning, the buyer tries a variety of actions to find out, which of them seem to be the best. The big challenge of this learning approach is that the buyer has to exploit what it has already learned in order to receive high rewards, but it also has to explore in order to find better actions that may earn higher rewards in the future.

B. Model specifications

We suppose that the buyer initially, without placing an order, requests each supplier to submit an offer stating the quality and the price for the requested quantity. After receiving this initial information, the buyer allocates the procurement volume according to its quality-price preference among the suppliers. Following the delivery of the supplier volumes, the buyer imperfectly observes the quality of the suppliers and captures the price. In order to update the supplier volumes for the next order period, the buyer weights the observed quality and the captured price according to its quality-price preference. The sequence of events is sketched in Figure 1.

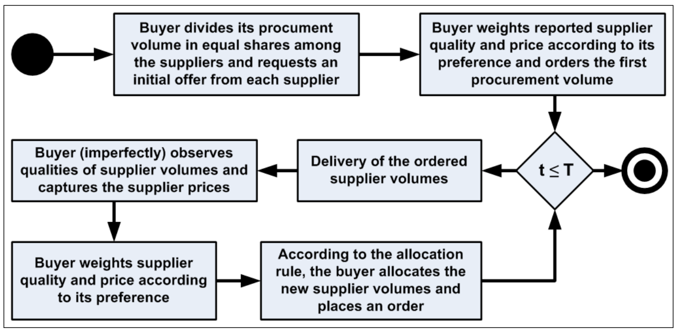

Figure 1. Sequence of events.

Buyer's procurement volume

We consider a buyer (abbreviated to B in formulas) who plans to allocate a constant procurement volume $X \in \mathbb{R}^+$ of a certain product among multiple suppliers in each period $t \in \{1, ..., T\} \subset \mathbb{N}$ of the entire observation time T. The buyer selects $m \in \mathbb{N}$ suppliers (abbreviated to S in formulas) for the delivery of the product, whereby the sum of the individual supplier volumes $x_{i,t}^S \in \mathbb{R}^+$ of m suppliers defines the buyer's procurement volume

$$X = \sum_{i=1}^{m} x_{i,t}^S \ . \tag{1}$$

Supplier quality

Each supplier is characterized by a quality function, which is described by a generalized logistic function, sometimes called Richards [11] curve. This S-shaped curve characteristic matches with our assumption that an increase in volume leads to a loss of quality and vice versa. Hence, each supplier's quality function is determined by

$$q_{i,t}^S(x_{i,t}^S) = H_i - \frac{H_i - G_i}{1 + C_i \cdot e^{-k_i \cdot x_{i,t}^S}} \tag{2}$$

where $q_{i,t}^S \in (0, 1)$ denotes the quality of the i'th supplier at time t. The quality parameters $(H_i, G_i, C_i, k_i) \in \mathbb{R}^4$ are set exogenously for each supplier at the very beginning of a simulation run. The parameter G_i (H_i) refers to the lower (upper) asymptote of the quality curve, while C_i is related to the quality in point $x_{i,t}^S = 0$, and k_i represents the logistic growth rate (or, in a negative sense, the logistic shrinkage factor). The quality parameters H_i, G_i, C_i, and k_i are purposefully designed so that the suppliers represent our 'typology' of suppliers. For the sake of simplicity, we suppose that all four parameters do not change over time.

Supplier price

We consider that each supplier sets a price $p_{i,t}^S \in \mathbb{R}^+$ for the quantity $x_{i,t}^S$ allocated by the buyer. Based on the aforementioned quantity-price trade-off, a monotonically decreasing price function is considered. Since we assume that doubling the quantity leads to a price reduction of a certain value, we use Wright's [16] log-linear model to describe the supplier price

$$p_{i,t}^S(x_{i,t}) = p^M \cdot (x_{i,t}^S + 1)^{\frac{log(1-L_i)}{log(2)}} \ . \tag{3}$$

While $p^M \in \mathbb{R}^+$ corresponds to the market price of one unit, $L_i \in (0, 1)$ denotes the supplier's relative price reduction. In regard to the experience curve effect, L reflects the proportion reduction in the unit cost with each doubling in the cumulative procurement volume (see, e.g., [1]). This means that with a doubling of the supplier volume, the supplier price deceases by $L_i \cdot 100\%$. Similar to the quality parameters, we also set the price parameters (p^M, L_i) exogenously at the very beginning of a simulation run.

Buyer's quality measurement

After the ordered supplier volumes are delivered, the buyer imperfectly observes the quality $q_{i,t}^B \in (0, 1)$ of the suppliers according to

$$q_{i,t}^B = q_{i,t}^S + Q_{i,t} \qquad \text{with } Q_{i,t} \overset{i.i.d.}{\sim} N(0, \sigma^2) \ . \tag{4}$$

We assume that there is a discrepancy between the actual quality $q_{i,t}^S$ and the observed quality $q_{i,t}^B$, since the observed quality is noise-afflicted captured in a normally distributed random variable $Q_{i,t}$ with mean 0 and standard deviation $\sigma \in \mathbb{R}^+$, which reflects the buyer's precision of quality measurement.

Buyer's quality-price preference

To express the buyer's preference regarding quality and price, we use the following additive weighted sum model

$$w_{i,t} = \alpha_t \cdot q_{i,t}^B + (1 - \alpha_t) \cdot \frac{p^M - p_{i,t}^S}{p^M} \tag{5}$$

where the individual weight of the i'th supplier is denoted by $w_{i,t} \in (0, 1)$. The term $(p^M - p_{i,t}^S)/p^M$ can be interpreted as a relative price saving on the part of the buyer. $\alpha_t \in [0, 1]$ and $(1 - \alpha_t)$ indicate the buyer's quality weight and buyer's price weight, respectively. Note, in our agent-based model, the parameter α_t is learned by the buyer using temporal difference learning (see Section III-C). A high (low) α_t indicates that the buyer puts more emphasis on quality (price) rather than on price (quality). Ultimately, a high individual weight $w_{i,t}$ implies that the buyer is generally content with the quality and price of the supplier.

Volume allocation

The procurement volume is allocated proportionately to all m suppliers depending on their individual weights. Thus, the buyer's volume allocation rule is given by

$$x^S_{i,t+1} = \frac{w_{i,t}}{\sum_{i=1}^{m} w_{i,t}} X \; . \tag{6}$$

Since the buyer might pursue different objectives regarding quality and price, a bigger share of the total procurement volume may be allocated to suppliers with a high individual weight.

Offer submission

At the beginning of a simulation run (hereafter abbreviated to t_1), we assume that the buyer splits the procurement volume in equal shares among the suppliers.

$$x^S_{i,t_1} = \frac{X}{m} \tag{7}$$

The suppliers are requested to submit an offer, stating quality q^S_{i,t_1} and price p^S_{i,t_1} for the requested supplier volume x^S_{i,t_1}. After the initial submission, the buyer orders its first delivery in consideration of (5) and (6). Whenever the supplier volumes are delivered (apart from t_1), the buyer observes the quality of each supplier according to (4). After the observe quality and the captured price have been determined, the buyer allocates the next procurement volume in accordance with (5) and (6). This procedure continues until $t > T$.

C. Learning method

In our agent-based model, the buyer has to specify the quality weight α_t in the additive weighted sum model (5). For this purpose, we use a temporal difference learning approach, which was invented by Sutton [14] because we suppose that the buyer does not have enough resources and, especially, no prior knowledge of its suppliers' structure to provide an adequate model of its multiple-sourcing environment.

Action-value function

The simplest temporal difference learning approach is given by the following update rule

$$V_{t+1}[\alpha_t] = (1 - \beta_t) * V_t[\alpha_t] + \beta_t * (\pi_t + \gamma * V_t[\alpha_{t+1}]) \tag{8}$$

where $V_t[\alpha_t] \in \mathbb{R}$ denotes the action-value function of action $\alpha_t \in A \subset [0, 1]$ with reward $\pi_t \in \mathbb{R}$, learning rate $\beta_t \in [0, 1]$, and discount factor $\gamma \in [0, 1)$. In the TD(0) method, the action α_t is a value from the discrete action space A that corresponds to all possible buyer's quality-price preferences and, besides that, the values of the action-value function are stored in a lookup table (labelled with square brackets) initialized to be zero for all actions, i.e., there is no information about the buyer's preference for quality and price when the simulation is started. Further, the action-value function can also be read as a long-term memory vector of length $|A|$ accumulating the discounted rewards over the time of a simulation run and, in the TD(0) algorithm, only one value of the action-value function is updated in each time step, while all other V-values remain unchanged.

Buyer's profit

The reward in the update rule (8) corresponds to the buyer's profit

$$\pi_t = \sum_{i=1}^{m} x^S_{i,t} \cdot min(q^B_{i,t}, q^S_{i,t}) \cdot p^R - x^S_{i,t} \cdot p^S_{i,t} \tag{9}$$

with retail price $p^R \in \mathbb{R}^+$. The buyer's profit, which is revenue minus costs depends only on how much the buyer produces in its firm because quality and price are functions of the delivered quantity $x^S_{i,t}$. For the sake of simplicity, we suppose that products of poor quality are sorted out.

Learning rate and learning time

A very key part in temporal difference learning is the speed of learning things. At the beginning of a simulation run, the learning rate β_t should be so high that any initial random fluctuations have only a minor impact and, on the other hand, β_t should decrease with time to assure that the buyer finds a local optimum of its action-value function [14]. Therefore, we use a variable learning rate β_t that decreases over time. The time or, more precisely, the number of time steps during which the buyer learns the parameter α_t is called the learning time $T_L \leq T$.

Action-selection policy

After the buyer has calculated the action-value function $V_t[.]$, the buyer tries to select an action α_{t+1} from its action space A in order to maximize the sum of its discounted rewards, which are received over time. Thus, the buyer is confronted with the trade-off between choosing the current action and choosing a varied action with the prospect of a higher reward in the future. An easy and common action-selection policy is the so-called ϵ-greedy policy, which means that with probability $(1 - \epsilon_t)$ the action with the highest $V_t[.]$ is chosen, while with probability $\epsilon_t \in [0, 1]$ a random action is selected [14].

$$\alpha_{t+1} = \begin{cases} \underset{\alpha \in A}{argmax}\; V_t[\alpha] & \text{with probability } (1 - \epsilon_t) \\ \sim Unif(A) & \text{with probability } \epsilon_t \end{cases} \tag{10}$$

Furthermore, in our model, ϵ_t is a decreasing function of the time with the two properties that, at the beginning of a simulation run, ϵ_{t_1} is one which indicates that the action-selection is total random (pure exploration) and, in the end, ϵ_{T_L} is zero (pure exploitation) and, thus with probability one, the final learned action α_{T_L} leads to the highest value of the action-value function learned during the simulation run.

Number of time steps after learning

After the buyer has learned which quality-price preference is a good choice, the buyer's supplier structure is analyzed. For this purpose, more time steps are simulated in which

Copyright (c) IARIA, 2021. ISBN: 978-1-61208-898-3

the parameter α_t is unchanged. In the case, the buyer's quality measurement system works perfectly, i.e., $\sigma = 0$, the buyer's supplier structure stabilizes within $T_S \in \mathbb{N}$ time steps after learning. Since we are also interested in how the model reacts when imperfectly quality is imperfectly measured by the buyer's quality measurement system, i.e., $\sigma > 0$, further $T_E \in \mathbb{N}$ time steps are required to evaluate the stochastic simulation outcomes. This results in a total number of $T_L + T_S + T_E$ time steps of which only the last T_E time steps are used to analyze the buyer's supplier structure.

IV. Parameter Settings

We conduct our simulation experiments in two steps: (1) we start with the 'perfect scenario' in which there is no discrepancy between the actual quality and the observed quality. (2) we investigate further scenarios, called 'imperfect scenarios', to find out how the buyer's supplier structure changes when imperfect quality is imperfectly measured by the buyer's quality measurement system.

TABLE II
Parameter settings.

Exogenous parameters	**Values/Types**		
Time steps to learn the parameter α_t	$T_L = 100$		
Time steps to stabilize the allocation	$T_S = 10$		
Time steps to evaluate the outcome	$T_E = 10$		
Number of sim. runs	$N = 1000$		
Number of suppliers	$m = 3$		
Market price	$p^M = 1$		
Retail price	$p^R = 1$		
Supplier type	Type 1	Type 2	Type 3
Supremum of $q^S_{i,t}$	$H_1 = 1.0$	$H_2 = 0.8$	$H_3 = 0.6$
Infimum of $q^S_{i,t}$	$G_1 = 0.0$	$G_2 = 0.0$	$G_3 = 0.0$
$q^S_{i,t}(x^S_{i,t} = 0)$ (in %)	$C_1 = 99$	$C_2 = 79$	$C_3 = 59$
Logistic growth rate	$k_1 = 0.23$	$k_2 = 0.109$	$k_3 = 0.068$
Inflection point	$x_1^{IP} = 20$	$x_2^{IP} = 40$	$x_3^{IP} = 60$
Relative price reduction	$L_1 = 0.05$	$L_2 = 0.10$	$L_3 = 0.15$
Action space	$A = \{0.0, 0.1, ..., 1.0\}$		
Discount factor	$\gamma = 0$		
Procurement volume	$X \in \{1, 2, ..., 100\}$		
Buyer's precision of quality measurement	$\sigma \in \{0, 0.01, 0.02, ..., 0.10\}$		

Supplier types

In our paper, we focus on a small buyer-supplier network and, therefore, we distinguish only between three different supplier types (hereafter type 1, type 2, and type 3), which represent fictitious companies pursuing different competitive strategies. Supplier type 1 captures a company, which pursues a high-quality leadership strategy and, thus, seeks to be unique regarding the high level of quality of its product it offers to the buyer and that the buyer rewards this with a premium price. Supplier type 3, on the other hand, captures a company, which pursues a low-cost leadership strategy where cost advantages are essential to gain a high return and long-term success. Supplier type 2 is considered to capture a company that failed to achieve one of the above-mentioned generic strategies and can be, to put it in the words of [10], labelled as 'stuck in the middle'. Similar to supplier type 1, also supplier type 2's quality deteriorates with a higher procurement volume and falls below the quality of supplier type 3.

Supplier quality

Figure 2 depicts the quality functions of our three supplier types (for the quality parameters see Table II). We set the quality parameters as follows: The quality parameter H_i corresponds to the quality that can be guaranteed with small volumes, while G_i is associated with the worst quality that can occur. In order to obtain the quality parameters C_i and k_i, we define the following two constraints: Parameter C_i refers to the quality, if nothing is produced, i.e., $C_i = -(q_i^S(0) - G_i)/(q_i^S(0) - H_i)$. For simplicity, we assume $q_i^S(0) = H_i - 0.01$. Parameter k_i can be obtained by solving $\frac{d^2}{dx_i^2} q_i^S(x_i) = 0$ or easier through $q_i^S(x_i^{IP}) = (H_i - G_i)/2$, which reflects the inflection point x_i^{IP} of the quality function, i.e., up to this point, the quantity-quality trade-off (which is nothing more than the first derivative of the quality function, i.e., $\frac{d}{dx_i} q_i^S(x_i)$) increases and, concurrently, the quality drops to half of its value. Both equations lead to the same solution, namely $k_i = ln(C_i)/x_i^{IP}$.

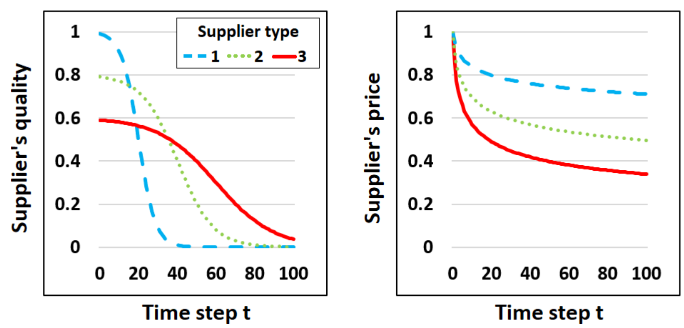

Figure 2. Suppliers' quality and price functions in our scenarios.

Incidentally, in the inflection point x_i^{IP}, the quantity-quality trade-off (hereafter abbreviated to τ_i^{IP}) finds its maximum because the first derivative is bell shaped with a peak at x_i^{IP}. For the sake of simplicity, we set the inflection points $x_1^{IP} = 20$, $x_2^{IP} = 40$, and $x_3^{IP} = 60$ or rather their quantity-quality trade-offs $\tau_1^{IP} = -0.0575$, $\tau_2^{IP} = -0.0218$, and $\tau_3^{IP} = -0.0102$ for our scenarios (see Table II).

Supplier price

The suppliers do not only distinguish from each other in terms of the quality, but also in terms of the price they charge. In regard to the supplier price considered in our model, Figure 2 depicts the price functions for our three supplier types. As mentioned before, a supplier price curve results, for instance, from a combination of various effects of learning, volume, and specialization. Therefore, we set the suppliers' price parameters in such a way that supplier type 2 can offer the buyer a relative price reduction that is twice as high as that of type 1, while supplier type 3 can provide a relative price reduction that corresponds to the sum of supplier type 1 and 2. This way of proceeding is also applied to the quality functions where the inflection points of the quality functions are determined. Note that our price (quality) function has a non-linear quantity-price (-quality) trade-off and, hence, cannot be expressed by a constant exogenous parameter. For the sake of simplicity, we suppose that the price of supplier type 1/2/3 decreases by $5\%/10\%/15\%$ each time the volume is doubled.

Action space

Next, we discuss the choice of the discrete action space A. If the number of possible quality-price preferences is small, the buyer's supplier structure cannot be investigated thoroughly, and, on the other hand, if the size of the action space is chosen too large then learning is slowed. Since we guess that the buyer can only differentiate between a limited number of quality-price preferences, we vary the parameter $\alpha_t \in A$ between 0 and 1 in steps of 0.1. Consequently, the buyer has eleven possible quality-price preferences.

Action-selection policy

For the action-selection policy in our model, we apply a monotonically decreasing ϵ-greedy policy (see Figure 3). To do this, we decide for a quadratic function ($\epsilon_t = at^2 + bt + c$).

$$\epsilon_t = -\frac{97}{970200} \cdot t^2 - \frac{1}{323400} \cdot t + \frac{9703}{9702} \tag{11}$$

To determine a, b, and c, we use the boundary conditions $\epsilon_{t=1} = 1$, $\epsilon_{t=T_L} = 0$, and $\epsilon_{t=T_L/2} = 0.75$, where the learning time T_L is set to 100. In order to achieve a higher level of exploration at the beginning than at the end, we set the third boundary condition to 0.75, which results in a degree of exploration of about 66% during a simulation run. Note that the area under the ϵ_t curve represents the degree of exploration.

Learning rate

With respect to the aforementioned learning rate properties, we use a linear function ($\beta_t = a + bt$).

$$\beta_t = \frac{111}{110} - \frac{1}{110} \cdot t \tag{12}$$

To solve a and b, we apply the boundary conditions $\beta_{t=1} = 1$ and $\beta_{t=T_L} = 0.1$, where the learning rate at the end of the learning time is fixed to 0.1 and this value is small enough so that, in all scenarios, the action-value function converges to a local optimum.

Discount factor

Another exogenous learning parameter in temporal difference learning is the discounting factor γ. Since we limit our research to a manageable number of scenarios, we set γ to zero, which means, that the buyer does not take future rewards into account.

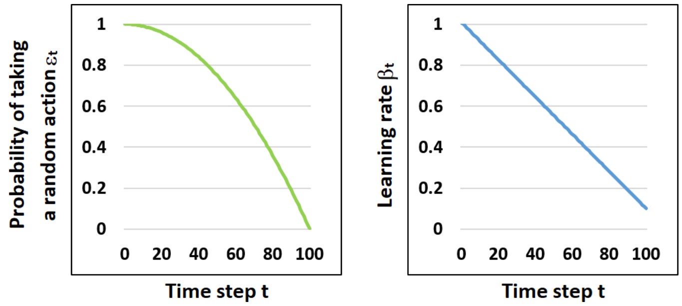

Figure 3. The left-hand-side graph shows the ϵ_t-greedy policy and the right-hand-side graph depicts the learning rate β_t.

Number of time steps and number of simulation runs

In our simulation experiment, the number of time steps to learn the parameter α_t in one simulation run is determined, on the one hand, based on the size of the action space A and, on the other hand, based on the degree of exploration. Our pre-generated simulations suggest a learning time of 100 to guarantee that the action-value function converges to a local optimum.

After the complexity of learning the parameter α_t has been determined by the cardinality of the action space and the degree of exploration, further T_S time steps are run through until the buyer's volume allocation converges. According to our pre-generated simulations, 10 time steps are enough to stabilize the behavior of the buyer's volume allocation in each scenario, hence $T_S = 10$. After the buyer's volume allocation has stabilized, only the buyer's quality measurement system has an impact and, therefore, we simulate further 10 time steps to take the stochastic fluctuations into account, hence $T_E = 10$. Note that, in all our examined scenarios, only the last T_E time steps are used to analyze the buyer's supplier structure.

Finally, we perform 1000 simulation runs for each scenario because, due to the coefficient of variance (ratio of standard deviation to the mean), 1000 simulation runs are sufficient to express the precision and repeatability of this simulation experiment.

Procurement volume and buyer's precision of quality measurement system

Besides the fixed exogenous parameters, we vary the procurement volume X and the buyer's precision of quality measurement system σ, which are also set exogenously at the very beginning of a simulation run. In particular, we model the procurement volume between 1 and 100 in steps of 1, because up to about 100 the buyer's profit in the scenarios is positive, and, in addition, we study a number of different values of σ ranging from 0 to 0.1 in steps of 0.01 in order to examine the effects of the buyer's suppler structure when the buyer's quality measurement system is not working perfect.

V. Results and Discussion

In this section, we analyze the results of our agent-based simulation in two main steps: (1) we present the results of our perfect scenarios in which the quality is perfect measured by the buyer's quality measurement system. (2) we analyze the imperfect scenarios when there is a discrepancy between the actual quality and the observed quality, i.e., the buyer's quality measurement system works imperfectly.

In each scenario, we start by analyzing the quality-price preference parameter α_t that the buyer learns during its learning time. Then, we study the buyer's supplier structure by comparing the suppliers' volumes $x^S_{i,t}$ relative to each other, because this allows for an easy interpretation and comparison, and this is also the main objective of this paper.

A. *Results of our perfect scenario with $\sigma = 0$*

Buyer's quality-price preference

We start by analyzing the buyer's quality-price preference α_t in our perfect scenario. Figure 4 depicts the means and standard deviations of α_t from 1000 simulation runs, whereby an extract for the means, standard deviations, and also the 95% percentiles of α_t from $X = 17$ to 23 is given in Table III. Interestingly, we identify a tipping point at around $X = 20$ where the buyer's preference changes from quality to price. Up to that tipping point, the buyer puts more emphasis on quality rather than on price and, from which onwards, the buyer's quality-price preference α_t gets lower, which means that the buyer prefers a supplier who charge a low price. In the tipping point, the quality-price preference is approximately 0.5, which implies that the buyer attaches equal emphasis on quality and price.

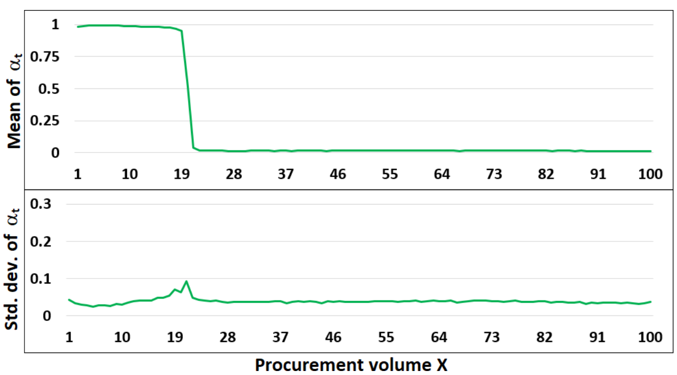

Figure 4. Means and standard deviations of the buyer's learned quality-price preferences α_t in our perfect scenario with $\sigma = 0$ from $X = 1$ to 100.

TABLE III
MEANS, STANDARD DEVIATIONS, AND THE 95% PERCENTILES OF THE BUYER'S QUALITY-PRICE PREFERENCES α_t IN OUR PERFECT SCENARIO WITH $\sigma = 0$ FROM $X = 17$ TO 23.

X	Mean of α_t	Std. dev. of α_t	95% percentile of α_t
17	0.98	0.05	[0.9 – 1.0]
18	0.97	0.05	[0.8 – 1.0]
19	0.95	0.07	[0.8 – 1.0]
20	0.52	0.06	[0.4 – 0.6]
21	0.04	0.09	[0.0 – 0.3]
22	0.02	0.05	[0.0 – 0.2]
23	0.02	0.04	[0.0 – 0.1]

In addition, it seems that the quality-price preference α_t slowly converges towards a value close to zero when the procurement volume increases. This would also be plausible because, if the procurement volume X becomes larger, the qualities of all three suppliers drop to almost zero, while the suppliers' prices still differ from each other. In such a case, the buyer prefers a supplier who charge a low price.

Buyer's supplier structure

In the next step, we analyze the buyer's supplier structure. For this purpose, we compare the suppliers' volumes $x^S_{i,t}$ relative to each other. Figure 5 reports the means of $x^S_{i,t}$. In the tipping point $X = 20$, every supplier type gets approximately a third of the procurement volume (specifically, $x^S_{1,t} = 35.58\%$, $x^S_{2,t} = 33.82\%$, and $x^S_{3,t} = 30.60\%$). This volume allocation appears plausible as far as the buyer attaches equal emphasis on quality and price. To the left of the tipping point, i.e., $X < 20$, supplier type 1 receives about 42% and supplier type 3 about 25% of the procurement volume and, for $X > 20$, the buyer's supplier structure turns over so that the relative supplier volume of supplier type 1 (type 3) slowly converges towards 16% (49%) when the procurement volume increases (cf. Figure 5). Only supplier type 2 always receives about 34% of the procurement volume.

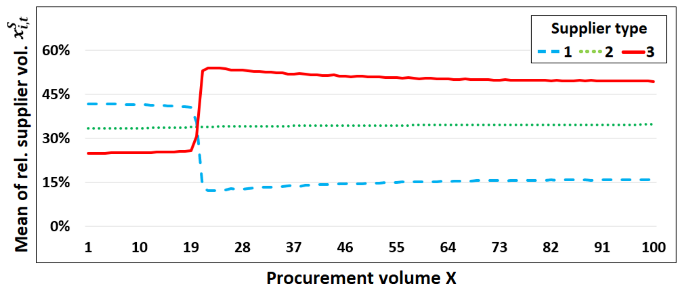

Figure 5. Means of the relative suppliers' volumes $x^S_{i,t}$ in our perfect scenario with $\sigma = 0$ from $X = 1$ to 100.

Simulation results over time

Conclusively, we investigate the quality-price preference and the supplier structure of the buyer for changes over time. In the first step, we look at the buyer's quality-price preference α_t. Figure 6 represents the buyer's learned quality-price preferences α_t over time. For procurement volumes lower than the tipping point, α_t slowly grows towards one, while for $X > 20$, α_t slowly approaches zero. Note that the buyer's quality-price preference remains unchanged after learning.

In the next step, we look at the buyer's supplier structure and, especially, how does the buyer's supplier structure shape the way an equilibrium-state is reached over time. For this purpose, we depict the means of the relative suppliers' volumes $x^S_{i,t}$. Figure 7 shows how an equilibrium-state is achieved when the procurement volume is given by $X = 10$, $X = 20$ (the tipping point), and $X = 30$. For lower procurement volumes ($X < 20$), the buyer needs approximately 1/3 of

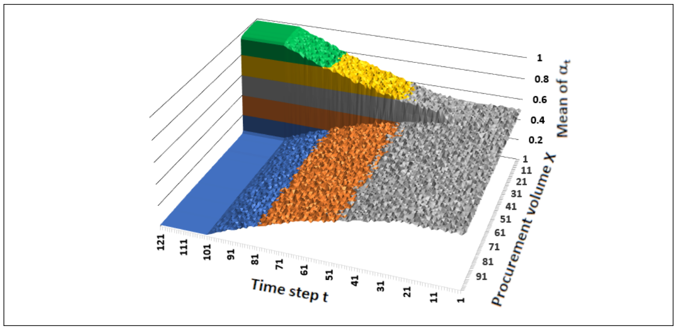

Figure 6. Means of the buyer's learned quality-price preferences α_t in our perfect scenario with $\sigma = 0$ from $X = 1$ to 100 and $t = 1$ to 120.

the learning time to separate the supplier types 1 and 3, while for larger procurement volumes ($X > 20$), the separation of suppliers proceeds much faster. In the tipping point, there is only a small separation, which means that all suppliers receive approximately the same proportion of the buyer's procurement volume.

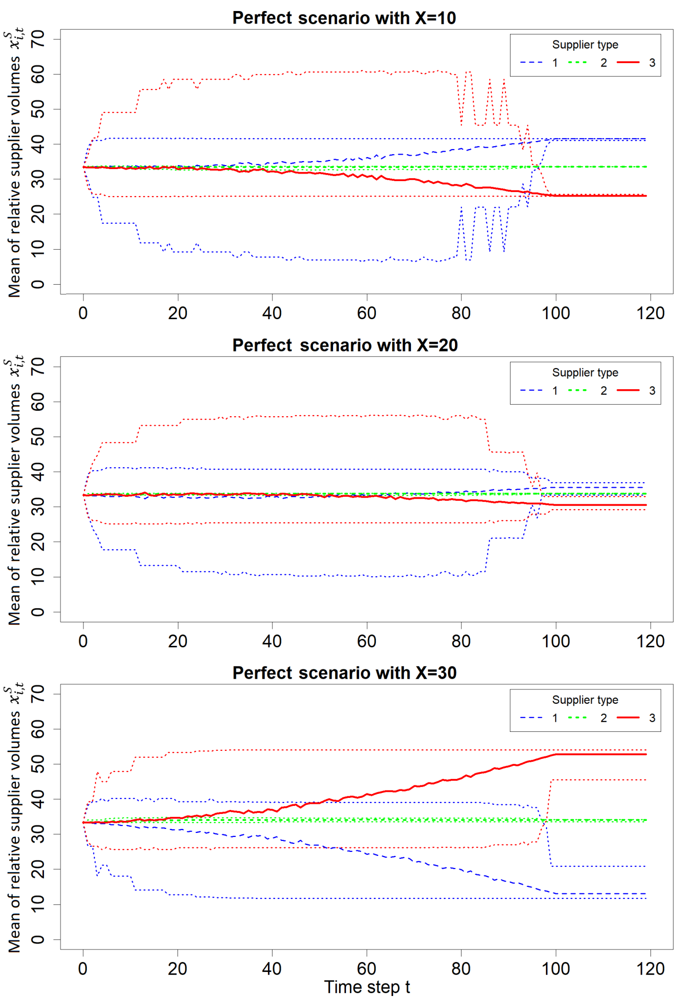

Figure 7. Means of the relative suppliers' volumes $x^S_{i,t}$ for selected procurement volumes with $\sigma = 0$ from $t = 1$ to 120. Means are represented by thick lines, while the 95% percentiles are displayed by thin lines.

B. *Results of our imperfect scenarios with $\sigma > 0$*

For the so far presented results, a perfect level of precision of the buyer's quality measurement was considered. In this section, we investigate scenarios in which the precision of the buyer's quality measurement is affected by noise and, hence, the quality is imperfectly measured.

Buyer's quality-price preference

Once again, we start with the buyer's quality-price preference. Figure 8 depicts the curve progressions of α_t from $\sigma = 0$ (minor differences on average between the actual quality and the observed quality) to $\sigma = 0.1$ (major differences).

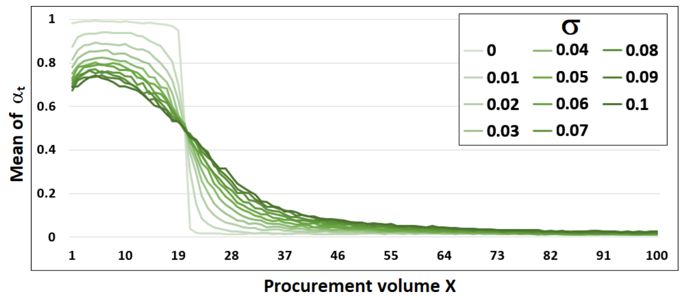

Figure 8. Means of the buyer's learned quality-price preferences α_t, whereby σ varies between 0 (light green) to 0.1 (dark green) from $X = 1$ to 100.

Buyer's supplier structure

Next, we plot the means of the relative suppliers' volumes $x^S_{i,t}$ from $\sigma = 0$ to $\sigma = 0.1$ (see Figure 9). Based on the curve progressions in Figure 9, the buyer's supplier structure becomes more stable because the value difference around the tipping point becomes smaller. In the case of $\sigma = 0.1$, the buyer's supplier structure is smooth enough so that there is no longer a jumping behavior around $X = 20$.

In comparison to the perfect scenario (see Table IV), supplier type 1 receives about 3.4% less (3.9% more) procurement volume to the left (right) of the tipping point and, thus, supplier type 1 loses (gains) market shares. Oppositely, supplier type 3 obtains about 3.4% more (3.8% less) procurement volume to the left (right) of the tipping point and, hence, supplier type 3 gains (loses) a greater share of the market than the other two suppliers. Note that, again, supplier type 2 receives around 34% of the total procurement volume and, on average, supplier type 2 is able to continue its position in the market. Summarized, if the buyer's precision of quality

TABLE IV
MEANS OF THE RELATIVE SUPPLIERS' VOLUMES $x^S_{i,t}$ TO THE LEFT AND RIGHT OF THE TIPPING POINT $X = 20$ WITH $\sigma = 0$ AND $\sigma = 0.1$.

	$X < 20$			$X > 20$		
	$\sigma = 0$	$\sigma = 0.1$	Diff.	$\sigma = 0$	$\sigma = 0.1$	Diff.
$x^S_{1,t}$	41.3%	37.9%	−3.4%	14.8%	18.7%	3.9%
$x^S_{2,t}$	33.5%	33.5%	0.0%	34.4%	34.3%	−0.1%
$x^S_{3,t}$	25.2%	28.6%	3.4%	50.8%	47.0%	−3.8%

Copyright (c) IARIA, 2021. ISBN: 978-1-61208-898-3

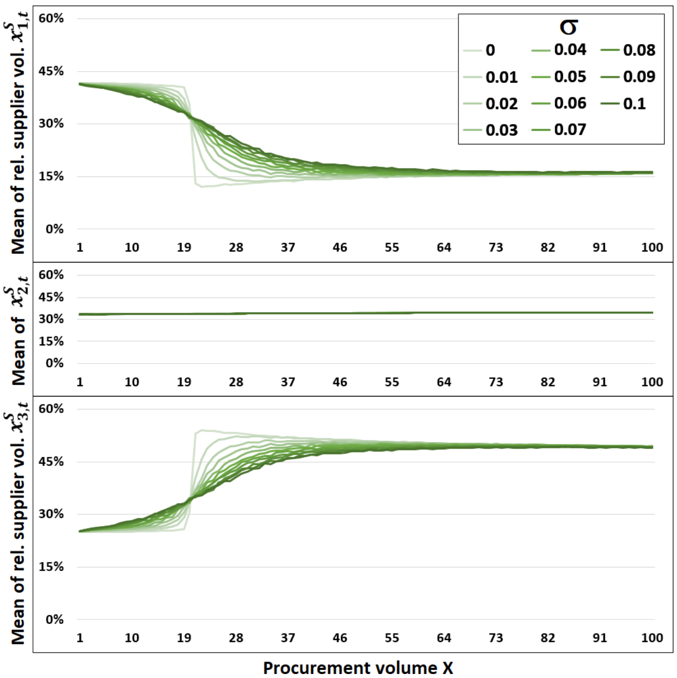

Figure 9. Means of the relative suppliers' volumes $x^S_{i,t}$ with $\sigma = 0$ to 0.1 from $X = 1$ to 100.

measurement system decreases, supplier type 1 benefits from large procurement volumes, whereas supplier type 3 profits from small procurement volumes.

VI. Conclusions

In our paper, we analyze the effects of different procurement volumes and levels of precision of the buyer's quality measurement system on the buyer's supplier structure in a multiple-sourcing environment using an agent-based simulation.

The results of our simulation experiment show that, in case the buyer has to learn its own quality-price preference via temporal difference learning, the buyer puts more emphasis on quality than on price when a small procurement volume is allocated among the existing suppliers. On the other hand, the higher the procurement volume allocated to the suppliers, the more important the price that is offered to the buyer. In such a case, the buyer is considered to pursue a cost leadership strategy searching for sources of cost advantage including, for instance, economies of scale, proprietary technology, preferential access to raw materials, and other factors.

Interestingly, we identify a tipping point at which the buyer's preference behaviour abruptly changes from quality to price. Up to that tipping point, the buyer puts more emphasis on quality rather than on price and, from which onwards, the buyer prefers a supplier who charges a low price. In the tipping point, the buyer attaches equal emphasis on quality and price.

Furthermore, we find that the poorer the precision of the quality measurement system and the lower the total procurement volume, the less the buyer orders from suppliers who focus on a high-quality leadership strategy. As a consequence, these suppliers lose part of their market shares to other suppliers. Such noises in the buyer's quality measurement system are at the expense of these suppliers, however the noise stabilizes the buyer's supplier structure, so that there is no longer a jumping behavior. Moreover, we find that the buyer separates the different supplier types much faster when the procurement volume is large and the quality measurement system works perfectly.

Finally, there are some limitations in this research: (1) for the sake of simplicity, we assume that the buyer only employs a limited number of suppliers and, in particular, that the suppliers' parameters are constant over time, which may not adequately represent the true market situation in a multiple-sourcing environment. (2) the buyer only makes decisions based on the suppliers' quality and price in order to allocate the procurement volume among the suppliers. (3) the suppliers do not communicate between each other nor they have the possibility to outsource part of their volumes to other suppliers. We believe this model provides some useful insights into the sourcing behaviour of a buyer who allocates different procurement volumes to suppliers who pursue different competitive strategies.

References

[1] M. J. Anzanello, and F. S. Fogliatto, "Learning curve models and applications: Literature review and research directions," International Journal of Industrial Ergonomics, vol. 41, no. 5, pp. 573–583, 2011.
[2] J. Asker, and E. Cantillon, "Procurement when price and quality matter," The Rand journal of economics, vol. 41, no. 1, pp. 1–34, 2010.
[3] G. J. Burke, J. E. Carrillo, and A. J. Vakharia, "Single versus multiple supplier sourcing strategies," European journal of operational research, vol. 182, no. 1, pp. 95–112, 2007.
[4] R. L. Keeney, and H. Raiffa, "Decisions with multiple objectives: preferences and value trade-offs," Cambridge university press, 1993.
[5] S. E. Kimes, "The basics of yield management," Cornell Hotel and Restaurant Administration Quarterly, vol. 30, no. 3, pp. 14–19, 1989.
[6] P. R. Koya, and A. T. Goshu, "Generalized mathematical model for biological growths," Open Journal of Modelling and Simulation, 2013.
[7] M. Lin, Jr. H. C. Lucas, and G. Shmueli, "Research commentary–too big to fail: large samples and the p-value problem," Information Systems Research, vol. 24, no. 4, pp. 906–917, 2013.
[8] B. C. McCannon, "The quality-quantity trade-off," Eastern Economic Journal, vol. 34, no. 1, pp. 95–100, 2008.
[9] T. M. Mitchell, "Machine learning," 1st ed. New York, NY: McGraw-Hill Education, 1997.
[10] M. E. Porter, "The Competitive Advantage: Creating and sustaining superior performance," NY: Free Press, 1985.
[11] F. Richards, "A flexible growth function for empirical use," Journal of experimental Botany, vol. 10, no. 2, pp. 290–301, 1959.
[12] A. J. Ruiz-Torres, and F. Mahmoodi, "A supplier allocation model considering delivery failure, maintenance and supplier cycle costs," International Journal of Production Economics, vol. 103, no. 2, pp. 755–766, 2006.
[13] K. Strmenik, F. Wall, C. Mitsch, and G. Mödritscher, "Volume allocation in multi-sourcing: effects of the quantity-quality trade-off," Central European Journal of Operations Research, 2020.
[14] R. S. Sutton, "Learning to predict by the methods of temporal differences," Machine learning, vol. 3, no. 1, pp. 9–44, 1988.
[15] J. Wallenius, J. S. Dyer, P. C. Fishburn, R. E. Steuer, S. Zionts, and K. Deb, "Multiple criteria decision making, multiattribute utility theory: Recent accomplishments and what lies ahead," Management science, vol. 54, no. 7, pp. 1336–1349, 2008.
[16] T. P. Wright, "Factors affecting the cost of airplanes," Journal of the aeronautical sciences, vol. 3, no. 4, pp. 122–128, 1936.

Copyright (c) IARIA, 2021. ISBN: 978-1-61208-898-3