Strategic Learning by e-Service Suppliers in Service Value Networks

Christian Haas, Steven O. Kimbrough and Clemens van Dinther

Journal of Service Research published online 21 December 2012
DOI: 10.1177/1094670512462130

Journal of Service Research
00(0) 1-18

DOI: 10.1177/1094670512462130
http://jsr.sagepub.com

Strategic Learning by e-Service Suppliers in Service Value Networks

Christian Haas[1], Steven O. Kimbrough[2], and Clemens van Dinther[1]

Abstract

This article examines strategic behavior of e-Service suppliers that offer electronic services in complex Service Value Networks (SVNs). In SVNs, consumers request a bundle of e-Services, and the SVN acts as an aggregator of single service instances, automatically configuring services from different e-Service suppliers into a complex service bundle, which is then offered to the consumers. In this context, e-Service suppliers who want to maximize their business success need to configure their services according to the (for them unknown) preferences of the consumers. Current literature, however, does not offer much guidance on how to find a fitting service offer in this situation, especially when the suppliers are faced with a changing consumer base. For this reason, we study two specific learning regimes that are able to capture and deal with the inherent complexity of the corresponding strategy space. Besides finding beneficial service offers, e-Service suppliers might also learn collusive behavior if it aligns with individual incentives. The potential occurrence of such behavior is the second aim of our work. Our results show that even with relatively simple learning regimes, e-Service suppliers are able to find beneficial offers and learn to collude tacitly (and presumably legally), which increases their profits.

Keywords

multiagent systems, simulation, strategic learning, service value networks

Introduction

Imagine, realistically, that the start-up online retailer "Easy & Clever Store" (ECS) finds that it needs a variety of information and communications technology (ICT) services, such as address validation, invoice generation, online payment, storage, and backup in order to provide the comprehensive business processes its customers demand. Instead of implementing these services itself, which might not be in its area of core expertise, ECS wants to source online payment and storage services as electronic, ICT services from a third-party vendor. ECS discovers the Service Value Network (SVN) platform "Easy SVN" (eSVN) which configures service bundles out of individual service components (normally software and data modules) in response to individual requests. On the platform eSVN various e-Service suppliers are offering their service modules, for example, several online payment service applications (*Easy Payment*, *Clever Payment*, and *Smart Payment*), and storage service systems (*Storage Online*, *Database Online*, and *Simple Data Storage*). If ECS becomes a new customer of eSVN and submits a request for a bundle of payment and storage services, eSVN then acts as a broker and assesses combinations of these six services. (In the present example, there are 3×3 possible configurations.) According to the customer's (ECS's) announced preferences, eSVN calculates the best fitting service bundle. Finally, ECS integrates the service bundle into its online store in such a way that its customers cannot see that online payment is provided through third party e-Service suppliers.

This example serves to introduce the general scenario discussed in this article. While the example is notional, the general story is established and realized in practice. As a result of advances in ICT, new organizational forms have evolved and firms create value by acting jointly over networks, "in which [arrangement] each actor contributes incremental value to the overall offering" (Basole and Rouse 2008, p. 57). Corroborating the point, Rust (2001) emphasizes that the "rapid expansion of the information economy and electronic networks" is an important long-term trend and an important driver for the rise of e-Services.

SVNs are components of an emerging e-Service form and have been studied recently by Blau, Krämer, et al. (2009) and Krämer et al. (2010). SVNs focus on ICT (or electronic) services, since these can be easily combined through plug-and-play functionality, which is harder for non-electronic services (based on the lack of standardization and/or communication between single service components).

In an SVN, there are four principal kinds of agents. (see Figure 1.) (1) *Consumers* (or customers, ECS in the example)

[1] Karlsruhe Institute of Technology, Institute of Information Systems and Management, Karlsruhe, Germany
[2] The Wharton School, University of Pennsylvania, Philadelphia, PA, USA

Corresponding Author:
Christian Haas, Institute of Information Systems and Management, Karlsruhe Institute of Technology, Karlsruhe 76128, Germany
Email: ch.haas@kit.edu

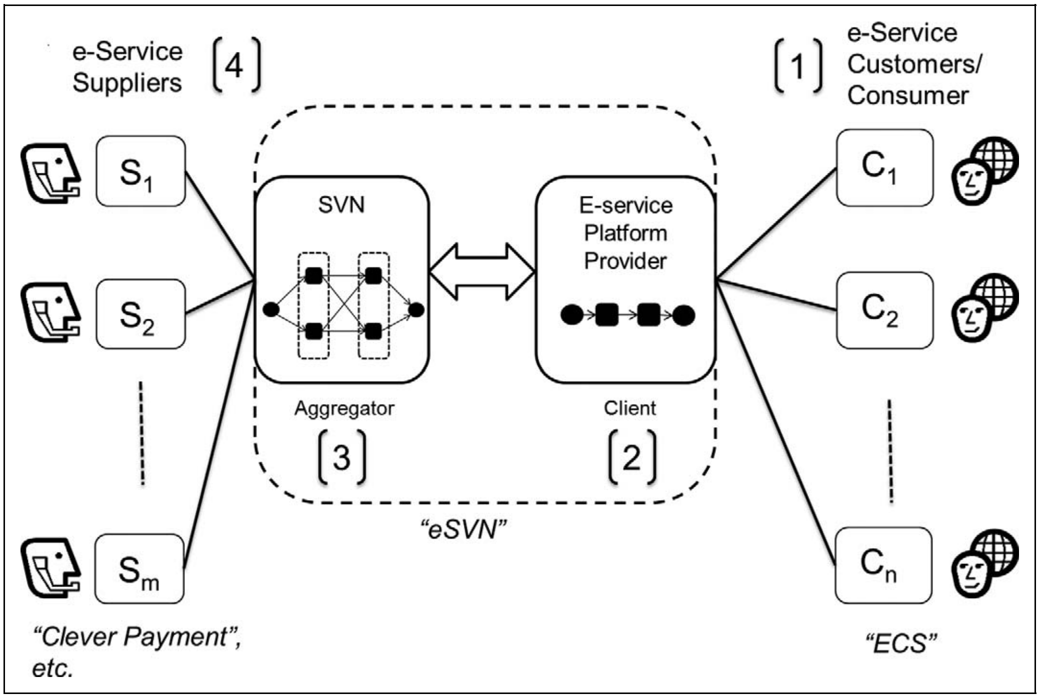

Figure 1. Scenario terminology.

request a complex service bundle, consisting of several services, which is offered by (2) an *e-Service platform provider* (eSVN in the example) on an e-commerce platform. The e-Service platform provider in turn is a client of an automated IT platform, (3) an *SVN* (also part of eSVN in the example). The role of the SVN is *automatically* to aggregate services provided by (4) e-Service suppliers (*Clever Payment*, *Database Online*, etc. in the example), seamlessly integrating the single services from potentially many different e-Service suppliers into a packaged product to the client. In this kind of setup, the e-Service platform provider, who sells directly to the consumer, is a retailer. The SVN is an aggregator, assembler, and distributor, and the e-Suppliers to the SVN are software OEMs (original equipment manufacturers). Typically (as in our example), the e-Service platform provider owns and operates the SVN, as hinted by the enclosing box in Figure 1. This intermediary (2+3 in Figure 1) can combine single service instances into a complex service bundle (e.g., representing a business process), and seamlessly provide it to end consumers. Salesforce.com and its associated services marketplace, AppExchange, is an example of such an integrated e-Service platform provider, combining both SVN service (3) and platform provision (2) into a single entity. (We note that Figure 1 is very close in form to the model for e-Services originally proposed by Rust and Kannan (2002); see their Figure 1.1. Rust and colleagues (e.g., Rust and Lemon 2001) are working with a model that is a generalization of the SVN model, and their comments apply directly to the present case.)

Most research on e-Services and the resulting platforms, such as SVNs, focuses on the development, characteristics, and organizational aspects of the described components and systems. For example, recent work covers a general organizational adaptability model for companies in networks similar to SVNs (Busquets, Rodon, and Wareham 2009), value network analysis with respect to organizational performance (Allee 2009), a theoretical model of how companies can successfully integrate in a value network based on a service-dominant logic (Lusch, Vargo, and Tanniru 2010), as well as rewarding companies that participate in an SVN (Conte et al. 2011). When it comes to specific managerial implications about how to act in such a new form of value network, especially the question of service offer configuration, current literature does not provide much guidance. Yet, as the configuring of service offers determines the business success of the company, how e-Service suppliers can learn to make advantageous offers in the complex strategy space of an SVN needs to be explored. The challenge of learning is exacerbated by the fact that consumers typically will have heterogeneous valuations for particular service attributes,[1] which are in general not known to the supplier (since it acts through the intermediary, and does not have direct contact with the end customer). Heim and Sinha (2001) point out that consumers' needs can be dynamic and have to be evaluated and fulfilled dynamically.

An SVN is generally believed to be a suitable way to cope with such dynamic preferences. Recalling the previous example, suppose that there is a recurring demand for a certain combination of services, for example, payment as well as storage services, and an e-Service supplier of payment services wants to offer its service via an SVN. The aim of the e-Service supplier is to choose a service configuration such that its service is selected for the integrated product. The main objective for each service supplier is to decide on a configuration of the service without *ex ante* knowing the clients' or consumers' preferences. If consumers' preferences are heterogeneous, it is unlikely that a single service configuration matches all clients' requests, hence the assessment of offer performance should be done over a certain period of time.

If consumers' preferences can be clustered in consumer segments with similar characteristics, service suppliers have to decide on which segment to focus. Of course, the efficacy of this decision depends on the decisions of competitors. Therefore, it would be beneficial for the service suppliers to coordinate their decisions. Put bluntly, service suppliers can either share the market or compete on price. Sharing the market aims at reducing competition in such a way that each supplier serves, mainly or even solely, a certain market segment. Direct negotiation is normally not possible, due to rules on restraint of trade.

This leads to our first research question or cluster of questions: Are e-Service suppliers able to learn profitable service configurations even though consumers' preferences are unknown to them? Are the suppliers able to coordinate their market decisions without direct communication? How does the applied learning strategy affect the results?

Since the ability to establish (tacitly) collusive behavior depends on learning by the suppliers in determining their service configurations, we compare two alternative learning approaches for this research question: a genetic algorithm (GA) and a form of reinforcement learning.

Instead of coordinating their decisions regarding service configurations (and the consumer segments respectively), service suppliers might also coordinate prices. For example, if several suppliers compete for a consumer segment, it might be beneficial for them to collectively raise prices. This would be indicated by higher prices than expected in a market with perfect competition.

This raises our second research question: Are service suppliers able to coordinate prices tacitly, that is, without direct communication?

In order to address these questions, this article is structured as follows. In the next section, we discuss related work on SVNs and learning in strategic contexts. The third section describes the model and the design of the simulation experiments. We evaluate the simulation results in the fourth section for the GA and next section for Probe and Adjust (P&A). We summarize the main findings and give an outlook on future work in the concluding section.

Related Research

Service Networks

Ostrom et al. (2010) point out the importance of research in the field of service networks. The evident trend toward flexible and adaptable networks among specializing partners is confirmed by a survey of senior executives and corporate decision makers performed by the Economist Intelligence Unit (Franklin 2005). Executives perceive ICT as a competitive tool that is critical to a firm's ability to adapt business models and strategies. ICT penetration in business processes will increase. In turn, the need for flexibility, specialization, and modularization requires company-spanning ICT support for business processes. Heuser, Lacher, and Perlmann (2007) remark that such flexible ICT infrastructure can be achieved through the modularization of business processes and the encapsulation of the underlying logic within software components that can be invoked as a web service on the Internet. "Traditional I[C]T infrastructures in which infrastructure and applications were managed and owned by one enterprise are giving way to networks of applications owned and managed by many business partners. Standards and the pervasiveness of network technologies provide the technology support for this trend" (Curbera et al. 2003, p. 28). This approach of constructing modular and interchangeable building blocks of software by encapsulating application logic into services and making them publicly available is known as Service-Oriented Architecture (SOA; MacKenziee et al. 2006; McAfee 2005; Papazoglou 2008; Papazoglou and Dubray 2004; Schroth and Janner 2007).

The recent trend of modeling and assembling SOA components is expected to continue and, according to a Gartner report, "through 2014, the act of composition will be a stronger opportunity to deliver value from software than the act of development" (Gartner-Inc. 2010). Firms are now able to build process-based SOA compositions for organization-specific purposes. The service components are offered by different service providers, implement specific protocols, and as such, "provide a distributed computing infrastructure for both, intra-and cross-enterprise application integration and collaboration" (Papazoglou and Georgakopoulos 2003, p. 27). Provisioning and composition of services is organized in SVN (Blau, Krämer, et al. 2009). Krämer et al. (2010) argue that SVNs are a special case of Smart Business Networks (SBNs; Heck and Vervest 2007). In SVNs, the coordination and orchestration of services is performed automatically by a universally accessible network orchestration platform. Thus, Krämer et al. (2010) define SVN as follows: "Service Value Networks are Smart Business Networks that provide business value by performing automated on-demand composition of complex services from a steady, but open pool of complementary as well as substitutive standardized service modules through a universally accessible network orchestration platform." Two parts of this definition are especially important for present purposes: first, service modules need to be standardized and accessible through a network; second, the composition needs to be automated. Therefore, because ICT-related (electronic) services are well defined and (compared to non-electronic services) easily standardized, SVNs target networks of electronic, ICT-related services.

Service Value Networks

The underlying hypothesis of cooperation among competitors—a form of "co-opetition" (Brandenburger and Nalebuff 1997)—is to combine the advantages of markets such as adaptability, flexibility, and efficiency with those of hierarchies such as control and protection of knowledge, skills, and competencies (Hamel, Doz, and Prahalad 1989; Heck and Vervest 2007; Hoetker 2006; Miles and Snow 1986). Flexibility and adaptability is achieved by building combinable modules (Hoetker 2006; Hoogeweegen and Vervest 2005). Proprietary knowledge is protected within modules and cooperation is needed to achieve interoperability of the partners' modules. We use the term *business network* (Anderson, Håkansson, and Johanson 1994, p. 1) as an umbrella term to include a collection of connected firms (Astley and Fombrun 1983; Miles and Snow 1992) or business relationships (Cook and Emerson 1978; Håkansson and Johanson 1993). The term network indicates an extension of theory beyond commonly studied dyadic forms of exchange (Emerson 1976, p. 357). Cooperation is economically motivated, with the participants expecting to realize mutual benefits and to cocreate business value (Blankenburg Holm, Eriksson, and Johanson 1996). Bengtsson and Kock (2000, p. 424) have pointed out that horizontal relationships can be purely competitive, purely cooperative, or a combination of the two.

Definitions of SVNs provided by Basole and Rouse (2008) and by Spohrer et al. (2008) emphasize collaboration of complementing and competing firms to create mutual benefit (as in SBNs). In an effort to distinguish SVNs, Krämer et al. (2010) propose a definition as described earlier that emphasizes the automated, dynamic creation of composite services in order to create value. Besides detailing the components of the definition, Krämer et al. (2010) argue that such dynamic SVNs with their inherent ability to combine modular services in various ways can successfully address long tail phenomena observed in many industries (Anderson 2007).

On the important question of how to price services, Roth, Woratschek, and Pastowski (2006) discuss settings in which it is best for service suppliers to select their price based on a negotiation or on a posted price. An economic model for SVNs was originally proposed by Blau, Conte, and van Dinther (2010) who develop an auction-based market mechanism that computes the best combination of modular services in determining a composite service based on consumer requirements. In addition, a few questions regarding SVNs have been studied from an economic perspective, for example, how to automatically select combinations out of a set of possible alternatives and to determine prices for such combined services (Blau, Conte, and van Dinther 2010; Blau, van Dinther, et al. 2009; Conte et al. 2009; van Dinther 2010). The proposed auction mechanism can be understood as a special form of negotiation.

Applying a mechanism design approach, the proposed mechanism of Blau, Conte, and van Dinther (2010) exhibits a number of desirable properties which ensure that single suppliers have no incentive to unilaterally submit untruthful service characteristics to the SVN. Because this incentive compatibility assumes that only one service supplier deviates, Blau, van Dinther, et al. 2009 study simple collusion strategies for service suppliers which could emerge due to the structure of the auction and its pricing mechanism. They find that service suppliers mutually benefit from collusion under certain circumstances, for example, if the competitiveness of the market is low. However, their simulations did not take into account the ability of service suppliers to select and adjust their service offers on other attributes than price.

Besides collusion, SVNs have been studied from the viewpoint of the network operator. Conte et al. (2009) introduce a contribution-based payment scheme that rewards service suppliers not only if their services are allocated to customers, but also for their mere presence in the network. They show that their scheme might be especially beneficial in the formation phase of the network, since it helps to foster essential objectives such as network variety and incentives to join the network.

Learning in Strategic Contexts

In SVNs, suppliers have to learn to adapt their service offers to changing consumer preferences, while at the same time taking competing offers from other companies into account. From a strategic point of view, this setting has been investigated both by the comprehensive field of behavioral game theory (see Camerer 2003 for a comprehensive review) and by agent-based computational economics (ACE, see e.g., Tesfatsion 2002). Looking at how companies can best react to their competitors' actions, game theory tries to find equilibrium outcomes (and their associated strategies). ACE, on the other hand, aims at simulating and understanding the dynamic properties of the system by modeling the involved parties as a set of intelligent agents, each of whom act based upon certain policies. In this work, e-Service suppliers want to learn advantageous service offers. This task can be of considerable complexity when all potential alternatives are taken into account. Due to this complexity, we model suppliers as agents with learning capabilities. A purely game-theoretic approach would be infeasible and cannot provide much guidance for choosing among the potentially many equilibria. Furthermore, learning conceived as some form of reinforcement learning is well accepted in strategic contexts (e.g., Brenner and Witt 2003; Erev and Roth 1998; Roth and Erev 1995, and many others).

One of the key characteristics of agents in ACE models is their ability to act upon perceptions of the environment and engage in goal-directed behavior (Wooldridge and Jennings 1995). In order to adapt to new situations in changing environments, agents often implement learning mechanisms that allow them to explore their environment and to assess the effects of their actions. Literature on learning usually categorizes learning models into two distinct classes. On one hand, reinforcement learning models build on the law of effect (Thorndike 1927), which basically states that actions that are tied to positive outcomes are chosen more frequently in the future, whereas actions with negative outcome are chosen less often. Such learning mechanisms normally assume that agents do not have models of their environment and do not actively reason and reflect on their actions and the consequences that result from these actions. On the other hand, cognitive learning models assume that agents have some kind of model of their environment and actively reason over their actions' potential effects (Brenner 1999, 2006).

The work and the results we report in this article rely on implementation of two learning procedures, one a GA and the other an algorithm called P&A. Both our learning procedures can be categorized broadly as reinforcement learning. The advantage of this type of procedure is that the agents do not require information about the strategies of other agents. Although there are some learning models that explicitly consider simultaneous learning (i.e., taking the potential decisions of other agents into account; e.g., Hu and Wellman 1998), these do not allow for finding new actions (in our case service offer configurations), and an evaluation of all possible service configurations would most likely require an infeasible amount of computation and time.

GAs (see Goldberg 1989; Holland 1992) are a widely used class of learning mechanisms, which have also been used extensively for economic applications (Arifovic 1994). A key advantage of GAs is that new strategies or solutions, which were not available initially, are developed over time. Hence, GAs are appropriate for sampling large strategy spaces because solutions are not restricted to the initially specified solutions (Duffy 2006). Evolutionary computation (the general category in which GAs fall, along with evolution programming, genetic programming, etc.) has a venerable history with regard to learning in games (e.g., Fogel 2002). Other related examples of using GAs are product family design to reduce manufacturing costs (D'Souza and Simpson 2003), and resource distribution in grid systems to increase service reliability (Dai and Wang 2006).

The second algorithm, P&A, is also a form of reinforcement learning. It differs substantially from the reinforcement learning methods widely employed in the behavioral game theory literature (e.g., Camerer 2003; Erev and Roth 1998 for a comprehensive review) in being designed to learn a continuous quantity, rather than a discrete option among a short list of options. This is necessitated by the fact that our agents have a large number of choices to make at each decision point, and so the standard behavioral models are not applicable. For example, the number of pure strategies in Erev and Roth's (1998) study was never more than five, which is usually much smaller than the number of service offer configurations that a company has to consider. P&A has been initially proposed by Kimbrough and Murphy (2009), who study this learning strategy in the context of Cournot oligopoly markets. It has since been applied to various other settings as well. For example, Kimbrough (2012) presents P&A models for Cournot, Bertrand, and stepped supply curve competition, while Skyrms (2010) applies it to agents on networks.

Experimental Setup

Modeling an SVN with a procedural model (simulation) requires specification of three main roles: e-Service suppliers (4 in Figure 1), consumers (1 in Figure 1), and broker aggregators (the SVN platform itself, 2 and 3 in Figure 1). As our model focuses on the e-Service suppliers, we assume that the broker (SVN) acts on behalf of the client (2, the e-Service platform provider). Building on the formal SVN model developed by Blau, van Dinther, et al. (2009), we now describe these components.

E-Service Consumers

Consumers represent the demand side of the market. They request complex services according to their preferences and with specific attribute levels, which are to be provided as a packaged service bundle by the e-Service platform provider. We assume that consumers switch suppliers according to their valuation of service attributes in the bundles they receive. Switching costs are low since offered services must satisfy standards of interoperability, which supports the plug-and-play-based exchange of service components. Nevertheless, switching costs such as soft factors like customer loyalty are important components as Bansal and Taylor (1999) remark. In our scenario, consumers are interested in the service bundle they receive from the e-Service platform provider and do not need to know the suppliers of the individual service components. Hence, factors like customer loyalty are of less importance in this case. However, such factors can easily be represented in our model as service characterizing attributes. For ease of understanding, we have concentrated on attributes commonly used in the Web Service literature.[2] Services in general can be described by attributes as discussed by Menascé (2002) and Zeng et al. (2003), among others. Furthermore, service quality models specify service attributes whose values describe the quality characteristics of services. Examples of service quality models include Fassnacht and Koese (2006) and Collier and Bienstock (2006), who design several dimensions of (e-)service quality. However, as their proposed dimensions are rather qualitative in nature, we consider only service attributes that can be quantified easily. In order to determine the relative importance of certain attributes, user surveys can be conducted. For example, Holloway and Beatty (2008) use a critical incident analysis to study the importance of goods and service attributes. As their attributes are rather coarse-grained and general in nature, we will consider the importance of e-Service attributes instead.

In our procedural model, services are characterized by four attributes: price and three quality attributes (performance, security, and availability level). These attributes were chosen due to their importance for cloud service customers (Hosting 2009) and because they afford modeling of differentiated service configurations, which is of interest for SVNs generally. Besides ease of understanding, these attributes can be replaced by other attributes that can describe the services and that the consumers have preferences for.

Consumers have individual preferences with respect to the three quality attributes. These preferences are expressed through preference weights λ_i with $\sum \lambda_i = 1, i \in \{1, 2, 3\}$. Consumers have a certain willingness to pay, α, for a service that

Table 1. Consumer Segments and Preferences

	Performance Segment	Security Segment	Availability Segment
Willingness to pay	[32.5,37.5]	[32.5,37.5]	[32.5,37.5]
Preference weights [P, S, A]	{0.8, 0.1, 0.1}	{0.1, 0.8, 0.1}	{0.1, 0.1, 0.8}
Performance bounds	{[0.5, 0.6], [0.9, 1.0]}	{[0, 0], [0.4, 0.6]}	{[0, 0], [0.4, 0.6]}
Security bounds	{[0, 0], [0.4, 0.6]}	{[0.5, 0.6], [0.9, 1.0]}	{[0, 0], [0.4, 0.6]}
Availability bounds	{[0, 0], [0.4, 0.6]}	{[0, 0], [0.4, 0.6]}	{[0.5, 0.6], [0.9, 1.0]}

Note. Preference weights are displayed for attributes performance (P), security (S), and availability (A).

yields a perfect score and that fulfills their preferences, whereas for subperfect services (with scores less than 1), they pay a fraction of α according to the allocation mechanism described below. In the simulation model, three types of consumers, or consumer segments, are modeled so that requests of these segments occupy different places in the demand/service space. Table 1 shows the consumer segments and respective attribute preferences used in the simulation model. Requests from consumer segments are uniformly drawn from segment-specific parameter intervals. To illustrate, for a consumer in the availability segment the availability requirement associated with a request will be uniformly drawn from an interval whose lower bound is itself uniformly drawn from the discrete set $\{0.5, 0.6\}$ and whose upper bound is uniformly drawn from $\{0.9, 1.0\}$.

Note that consumers do not need to have perfect information about all the offerings, as only their preferences are an input to the client and broker, which performs the matching automatically.[3]

SVN Platform

The SVN platform intermediates between consumers (1 in Figure 1) and (e-Service) suppliers (4 in Figure 1) by matching demand and supply. The platform contains a list of service offers posted by suppliers and classifies them into candidate pools of services with similar functionality. In the simulation model, it is assumed that there are two candidate pools of the same size and that a complex service consists of the combination of two services out of these candidate pools. As in our initial example, the services of the first candidate pool can be the previously described payment services, whereas the second candidate pool can be the storage services. The attribute values A_f^l of the aggregated complex service are calculated as the average of attribute levels a_j^l for the quality attribute type l of modular service j, $A_f^l = \frac{l}{k}\sum_{j=1}^{k} a_j^l$, and the sum of individual service prices for the complex service price (where k is the number of service modules the complex service is composed of; in the previous example we have $k = 2$). Upon arrival of a consumer request, the e-Service platform provider (in conjunction with the SVN) uses its list of service offers to calculate the welfare maximizing allocation, where welfare is the sum of consumer utility and supplier surplus. The allocation mechanism is based on a scoring auction that calculates the utility for each feasible combination of services that fulfills the requirements of the consumer request, and selects the service combination maximizing this utility. Hence, the allocation mechanism is given by

$$\Theta := \operatorname{argmax} U_f = \operatorname{argmax}\left(\alpha \times S\left(A_f\right) - P_f\right),$$

where f is a service combination out of the set of feasible service combinations F, U_f is the utility of f, P_f is the sum of service prices of all services in f, and $S(A_f)$ is the score of the service combination according to the Scoring function $S\left(A_f\right) = \left(\sum_{l=1}^{L} \lambda_l \times \left\|A_f^l\right\|\right)$, where

$$\left\|A_f^l\right\| = \begin{cases} 0, if\ A_f^l \leq \gamma_{lb} \\ \frac{A_f^l - \gamma_{lb}^l}{\gamma_{ub}^l - \gamma_{lb}^l}, if\ \gamma_{lb} < A_f^l < \gamma_{ub} \\ 1, if\ A_f^l \geq \gamma_{ub} \end{cases}.$$

is the relative score of attribute level A_f^l of service combination f for attribute l, $l \in \{1, 2, 3\}$. γ_{lb} is the lower bound of the consumer's preference interval, γ_{ub} the upper bound. As mentioned previously, the term $\alpha \times S(A_f)$ denotes the actual payment of a consumer based on the score of service combination f. Hence, the consumer's willingness to pay is a function of the distance of service combination f to the consumer's announced preferences. The model assumes that if the utility of the best service combination is negative, no service offer will be allocated and the consumer request is not satisfied.

After calculating the allocation, all service suppliers who own an offer that is part of the best service combination f^* receive a transfer payment $t_s = \sum_{j \in \sigma(s)} p_j + (U^* - U_{-s}^*)$, where $\sigma(s)$ denotes the set of allocated service offers in f^* owned by supplier s, p_j is the posted price for offer j, U^* denotes the maximum utility of f^* and U_{-s}^* is the (second-best) utility of the best service combination without the service offers owned by s.

Thus, allocated suppliers receive a discount $(U^* - U_{-s}^*)$ in addition to the posted price for their offers. This payment scheme is an example of payment schemes used in the class of Vickrey-Clarke-Groves mechanisms (see Clarke 1971; Groves 1973; Vickrey 1961). Further, as Blau, van Dinther, et al. (2009) show, it is a weakly dominant strategy for service suppliers to report their true service characteristics to the SVN (and not to overstate or understate them). Notice that the addition of these discounts can, in fact, lead to negative overall utility for consumers, as their utility without discounts might be smaller than the sum of discounts to allocated suppliers.

E-Service Suppliers

E-Service suppliers (4 in Figure 1) are described by their objective function, their service offers, and their learning mechanisms. The main goal of suppliers is to optimize their objective function measuring the success of service offer o, by selecting appropriate service offer configurations. In the

Table 2. Service Supplier Types and Cost Structures

	Balanced Type	Performance Type	Security Type	Availability Type
Variable Costs $[b_{pf}, b_{sl}, b_{av}]$	10, 10, 10	5, 15, 15	15, 5, 15	15, 15, 5
Performance Bounds	[0, 1]	[0.5, 1]	[0, 1]	[0, 1]
Security Bounds	[0, 1]	[0, 1]	[0.5, 1]	[0, 1]
Availability Bounds	[0, 1]	[0, 1]	[0, 1]	[0.5, 1]

simulation model, profit maximization is used as the objective. Suppliers offer services by specifying their price and attribute levels (constituting an offer configuration). Thus, we can describe each service offer (or bid, B) by a 4-tuple:

$$B(o) = (p(o), a_{pf}(o), a_{sl}(o), a_{av}(o)).$$

Providing a service o with quality levels $a_{pf}(o), a_{sl}(o), a_{av}(o)$ induces costs $c(o)$, where $c(o) = f(a(o))$ is a function of the attribute values. We assume that the cost function is nonlinear. This is in accordance with theoretical and empirical observations that the cost of providing highly dependable services drastically increases as system dependability approaches 100% (Sommerville 2007, pp. 47-50). The generalized cost function in the simulation model is given as: $c(o) = b_{pf}\left(\frac{a_{pf}-pf_{lb}}{pf_{ub}-pf_{lb}}\right)^2 + b_{sl}\left(\frac{a_{sl}(o)-sl_{lb}}{sl_{ub}-sl_{lb}}\right)^2 + b_{av}\left(\frac{a_{av}(o)-av_{lb}}{av_{ub}-av_{lb}}\right)^2.$

Here, b_i denotes supplier-specific variable attribute costs (i.e., how expensive it is for a supplier to offer higher quality levels of that attribute), and a_{lb}, a_{ub} are the supplier's lower and upper bounds of the service attribute levels. Suppliers set the price according to $p(o) = c(o) + \delta$, where $\delta \in \{0, 0.2, \ldots, 3.8, 4.0\}$ is the mark-up (MU) on the actual service costs. Each attribute level a_i is chosen from 21 uniformly distributed values in a specific range, for example, $a_{pf} \in \{0, 0.05, \ldots, 0.95, 1.0\}$ for the range [0, 1].

There are four supplier types modeled, as specified in Table 2. The balanced supplier type is used to model homogeneous suppliers where no single supplier has any cost advantages. In contrast, the three remaining supplier types each have cost advantages for one service attribute, yet have higher variable costs for the other two attributes. This represents the fact that variable costs for suppliers might depend on multiple factors, such as supplier size, technology, and so on.

As described above, the main goal of service suppliers is to learn to select successful service offers. Each supplier uses an instance of a learning mechanism to evaluate and adjust its current service offers. The learning mechanisms studied in this article are a GA and reinforcement learning in the form of P&A. Details about the learning mechanisms will be discussed in subsequent sections.

Simulation Process

Having established the main components of the procedural model, we now describe the flow of action during its execution. Figure 2 shows a flowchart of the simulation. At the beginning of the simulation run, initial service offers (positions) of suppliers are chosen randomly (Step 1) and the counter of simulation rounds is increased (Step 2). As long as the maximum number of episodes per simulation round, denoted by maxEpisodes, is not reached (Step 3), the episode counter is increased and the simulation enters the inner loop (Step 4). Steps 3–9 comprise this inner loop in which service offers and requests are matched. In Step 5, a consumer request is drawn from a scenario-specific consumer distribution. Suppliers select a service offer in Step 6 and submit their offers to the SVN. How the suppliers select an offer depends on the learning mechanism they use in the scenario. In the case of a GA, they select an offer out of their current population, whereas in the case of P&A they randomly vary the levels of their current service offer on one or more attributes.

After the consumer request has been drawn and service offers are submitted, the SVN calculates the resulting allocation in Step 7. Subsequently, in Step 8, suppliers receive their payoffs, and several performance measures such as suppliers' payments and consumer utility are recorded. If the inner loop is completed and the stopping criterion is not fulfilled (Step 9), suppliers use their learning mechanisms to update their current service offers (Step 10). The actual implementation of the updating procedure depends on the learning mechanism the suppliers use. In the case of the GA, they calculate a new population of service offers, and for P&A they change their current service attribute levels in the direction with higher average reward. After a certain number of simulation rounds the stopping criterion is fulfilled, the simulation records are saved (Step 11), and the simulation run ends.

Simulation rounds are repeated until a stop condition (Step 9) is reached. For all scenarios, convergence of total supplier profits is used as the stopping criterion. Suppliers' positions and payoffs often converged after a certain number of simulation rounds. In a few cases, however, convergence was not particularly clear, with the results suggesting steady but slow improvement of supplier positions and payoffs. In these cases, the simulation was stopped when the performance of the suppliers over a large number of simulation rounds exhibited only very minor improvements. This was only observed when suppliers used the GA.

The previous description of a simulation run is fairly generic and many possible instantiations are possible. For the simulation runs in this study, the calculation of allocations and payments (Step 7) is fixed for all scenarios. In contrast, the consumer segment distribution (Step 5), the implementation of supplier learning (Step 10), and the number of episodes between service offer updates (Step 3) varied across scenarios. Usually, suppliers assess the payoff of their current offers over 100 episodes per simulation round; yet in the GA case, the number of episodes is sometimes increased to 500 in order to study the impact of service offer evaluation time on GA performance.

The initial service configurations of suppliers are drawn randomly from the allowed range of service offers. Because of

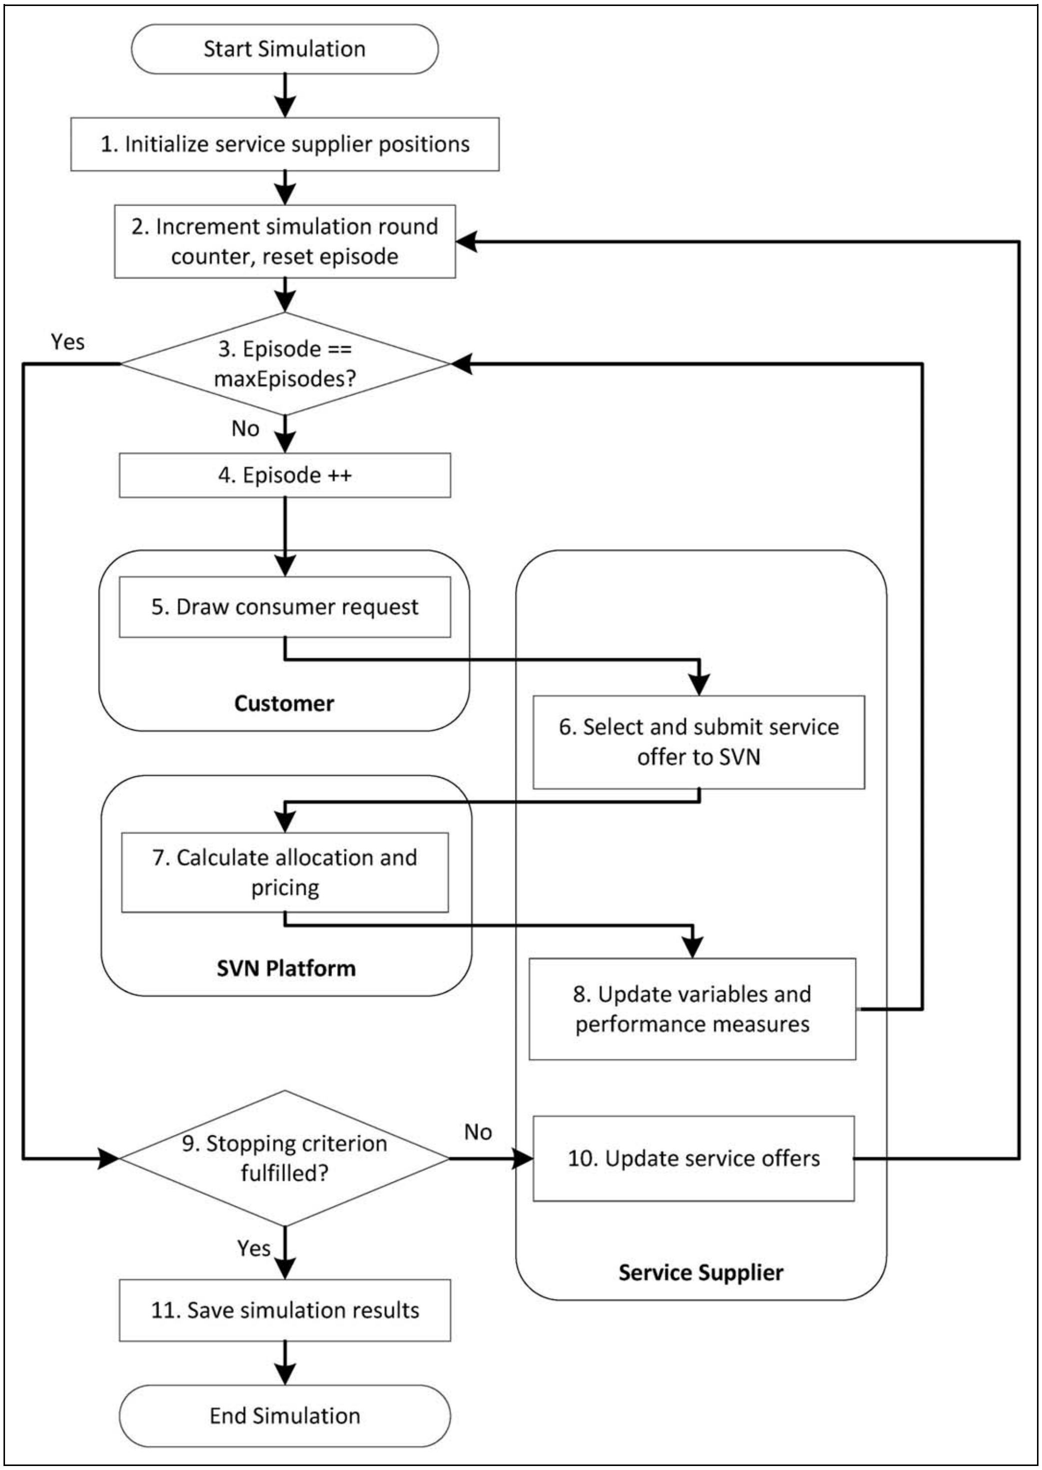

Figure 2. Simulation flowchart.

this, profits of individual suppliers depend highly on these initial service configurations, and the simulation needs a certain amount of time to converge to stable supplier configurations. In order to alleviate this startup problem (Law 2007, 508ff.), the approach of initial-data deletion or replication-deletion is used. Specifically, in all scenarios, the method of Welch (1983) is used to determine convergence.

The Welch procedure is applied on the sum of suppliers' profits. This is done because profits depend on the relative position of suppliers' offers in the service space, and a convergence in profits indicates that suppliers have settled on their service offer positions. Alternatively, the allocation percentages, that is, the number of successfully served consumer requests, could be used for convergence estimation. However, a comparison of these two convergence metrics revealed that convergence occurs faster for allocation percentages because, despite being able to serve a consumer request, suppliers might be able to increase their profit through an appropriate adjustment of the service offer configurations. Hence, convergence tends to occur earlier in the case of allocation percentages. In all subsequently considered scenarios, convergence with respect to service supplier profits can be observed after a certain number of simulation rounds. Depending on the scenario, 5,000 or 10,000 simulation rounds are run and the last 2,000 rounds are used for parameter estimation.

Example

Consider the following scenario in which a consumer requests a complex service composed of the two basic services mentioned in our beginning example, a payment service and a storage service. The consumer, drawn from the performance-oriented consumer segment, specifies that her preference weights for the attributes performance, security and availability are $\Lambda = [0.8, 0.1, 0.1]$. Further, she specifies the boundaries of the attribute levels as {[0.5, 1.0], [0.0, 0.5], [0.0, 0.5]} for the aforementioned attributes performance, security, and availability. Her willingness to pay, α, for a service that perfectly fulfills her request is 35. Each basic service is offered by two homogeneous e-Service suppliers, where suppliers s_1 and s_2 offer the payment service and suppliers s_3 and s_4 offer the storage service. The service offer levels are as follows: s_1 offers service levels of [1.0, 0.2, 0.2] at the price of $p_1 = 11.8$, s_2 offers [0.5, 0.5, 0.5] at $p_2 = 8.5$, s_3 offers [0.5, 0.5, 0.5] at $p_3 = 8.5$, and s_4 offers [1.0, 0.2, 0.2] at $p_4 = 11.8$. It is also assumed that the suppliers' costs for the service are $c_i = p_i - 1$, that is, suppliers charge a MU for service provision.

The SVN uses these offers to calculate the resulting allocation. There are four possible service combinations, and for each of these combinations the consumer utility according to $U = \alpha \times S(A_f) - P_f$, where $S(A_f)$ is the score and P_f the price of the complex service f. The basic service attribute levels are aggregated by calculating the average of the basic service levels. Thus, complex service $f_1 = s_1 \oplus s_3$, which is the combination of basic services s_1 and s_3, has the attribute levels [0.75, 0.35, 0.35] and costs $c_1 = 20.3$. Similarly, we have $f_2 = s_1 \oplus s_4$ with attribute values [1.00, 0.5, 0.5] and costs $c_2 = 23.6$, $f_3 = s_2 \oplus s_3$ with attribute values [0.5, 0.5, 0.5] and costs $c_3 = 17.0$, and $f_4 = s_2 \oplus s_4$ with attribute values [0.75, 0.35, 0.35] and costs $c_4 = 20.3$. The score $Sc_1 = S_1(Af)$ for f_1 is calculated as:

$$S_{c_1} = 0.8 \times \left(\frac{0.75-0.5}{1.0-0.5}\right) + 0.1 \times \left(\frac{0.35-0.0}{0.5-0.0}\right) + 0.1 \times \left(\frac{0.35-0.0}{0.5-0.0}\right) = 0.54,$$

similarly $S_{c2} = 1.0$, $S_{c3} = 0.2$, and $S_{c4} = 0.54$. Hence, consumer utility U_i for service i, $i \in \{1, 2, 3, 4\}$, is $U_1 = 0.54 \times 35 - 20.3 = -1.4$, $U_2 = 1.0 \times 35 - 23.6 = 11.4$, $U_3 = 0.2 \times 35 - 17.0 = -10.0$, and $U_4 = -1.4$. The allocation mechanism then selects complex service f_2 as it yields the highest utility $U^* = 11.4$ for the consumer, and service suppliers s_1 and s_4 are allocated. For the calculation of discounts, the maximum utility without the allocated suppliers is calculated. Without service supplier s_1, maximum utility U^*_{-s1} is -1.4, hence s_1 receives an additional discount of $U^* - U_{-}{}^{*}{}_{s1} = 12.8$. In this example, supplier s_4 receives the same discount of 12.8. Thus, the allocation yields following performance measures. Service suppliers s_1 and s_4 receive payments of 24.6 each (price plus discount) and have a surplus of $24.6 - 10.8 = 13.8$ each, through the discounts consumer surplus is $U^* - 2 \times 12.8 = -14.2$, and overall welfare is $W^* = -14.2 + 2 \times 13.8 = 13.4$.

Learning via Genetic Algorithm

As the first learning strategy, service suppliers use a GA to learn service configurations. GAs have been widely used for and successfully applied to many optimization problems. The core idea of the metaheuristic is to start with a number of potential solutions, the population, and evolve this population by applying certain genetic operators to its members. The goal of the GA is to find solutions that yield a high performance value, called the fitness. Since the operators ensure that solutions with better fitness values have a higher chance of being selected for subsequent populations, the fitness of the population is likely to improve over time. However, as with all heuristic procedures, the GA is not guaranteed to find globally optimal solutions (Goldberg 1989).

Typically, a GA implements three genetic operators, which we also use in our simulation model. The selection operator selects solutions from the old population based on their fitness values. The other two operators, crossover and mutation, are only applied to these selected solutions. During crossover, pairs of the selected solutions are chosen and parts of the encoded solution strings are swapped between the solutions, in analogy to the biological recombination of chromosomes. Finally, the mutation operator randomly changes the values of parts of the encoded solution according to a prescribed probability.

Of the many possible parameter settings to specify the genetic operators of a GA, parameter settings according to the work of De Jong and Spears (1991), who showed that these settings outperform other settings for many optimization problems, are used.

Each supplier has a population of 50 service offers, which is initialized randomly. Figure 3 shows that each service offer is encoded as a chromosome consisting of four genes (alleles) representing the four attributes of a service offer. Each gene has values in the range of [0, 20] as described in earlier sections. The current population of service offers is valid for one simulation round, after which the supplier generates a new population using the genetic operators.

In each episode, the supplier selects one service offer out of the population, submits it to the SVN, and evaluates its performance. After each simulation round (of 100 or 500 episodes in our runs), the selection procedure is performed by a tournament selector of size two which performs 50 tournaments to select an intermediate population suitable for crossover. In every tournament, two service offers of the old population are randomly chosen and the offer with the higher fitness is selected for the intermediate population. Afterward, pairs of chromosomes out

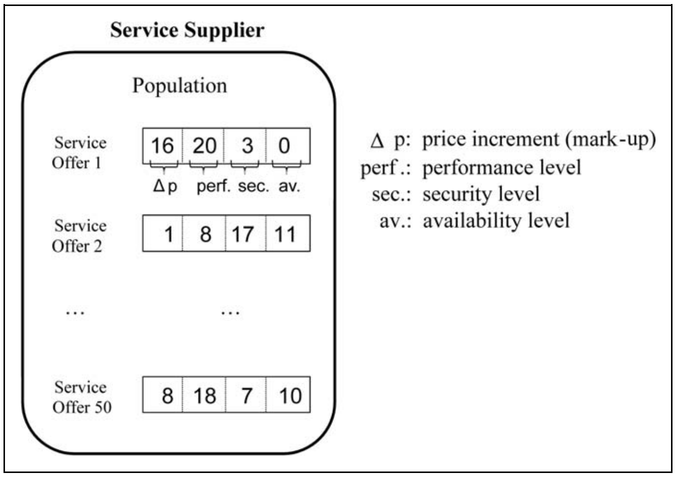

Figure 3. Encoding of service offers in the simulation model.

of the intermediate population are randomly chosen and the crossover operation is performed on these pairs with probability .6. If crossover is not performed, the chromosomes are left unchanged. Finally, mutation is applied on the new population of service offers, where the probability that a gene (an attribute of a service offer) , is mutated equals .005. In the case of mutation, the level of the attribute is randomly chosen in the range [0, 20].

The potential performance of the GA crucially depends on its ability to correctly assess the fitness of individual service offers. Because the payoffs of service offers are affected by the simultaneous choices of competitive suppliers, it is beneficial for suppliers to try each service offer of the population as often as possible in order to properly estimate its fitness in various settings. Hence, increasing the number of episodes per simulation round should increase the performance of the GA with respect to the success of suppliers as they are able to test their offers more often before updating the population.

We study the effect of supplier learning in SVNs and occurring collusion in several simulation scenarios, representing different market environments and consumer settings.

Therefore, we formulate the following hypotheses:

Hypothesis 1: Allocation percentages of suppliers are higher if consumers are segmented rather than randomly distributed.

Hypothesis 2: Average supplier profits are larger in segmented markets than for random consumer requests.

Hypothesis 3a: Supplier heterogeneity incentivizes suppliers to divide the market and thereby increases the percentage of allocated services.

Hypothesis 4a: Supplier heterogeneity increases average supplier profits.

Hypothesis 5a: Service suppliers are able to earn higher profits if they can charge additional mark-ups (MU) on their marginal costs (MC).

The results from these scenarios are discussed below. In these discussions, the term allocation percentage denotes the percentage of episodes of a simulation round in which the consumer request was successfully served.

Scenario 1: Tacit Market Division Among Suppliers Using GA

In the first scenario, there are three e-Service suppliers per candidate pool competing for consumer requests (delivered via the client e-Service provider). Two potential demand structures are considered, with the intention of studying the effectiveness of the proposed learning scheme. First, consumer preferences for service offer configurations can be randomly drawn out of the space of possible preference values, and second, the demand can be structured into three equally large client classes with distinct preferences for service configurations as shown in Table 1. In the case of structured consumer demand, it is optimal for all service suppliers (from their perspective) to specialize in one of the consumer segments, thereby avoiding competition and achieving an equal share of the entire market. Hence, in the optimal case, each supplier offers a service configuration for a specific segment and receives equal profit due to the symmetric consumer segment design.

As described earlier, all homogeneous suppliers have the same cost function, which means that none of the suppliers has an incentive to specialize on a particular segment. Table 3 presents the results of the first scenario and allows for several interesting observations. First, both in markets with random consumer requests and in markets where demand is clustered into consumer classes, service suppliers are not able to serve all consumers requests. In other words, in segmented markets suppliers do not learn the optimal solution to each specialize on a distinct segment in all replications. Second, both the percentage of successfully allocated service offers and the average profit of suppliers are higher in markets with distinct consumer segments. Because the data are not normally distributed, we used a Wilcoxon's Rank-Sum test (Wilcoxon 1945) to test whether the differences are statistically significant. At least for supplier profits they are significant, with p values of $<.001$. For allocation percentage, the average is higher for segmented markets; yet, the p values do not indicate a significant difference at the 5% level. This refers to Hypotheses 1 and 2. This finding can be explained by the very nature of the SVN mechanism itself. For segmented markets, consumer requests are centered on certain areas in the service configuration space, and individual requests are repeatedly drawn from these areas. This continuity enables suppliers to better learn promising service configurations as compared to random consumer requests, which explains the higher allocation percentage. In this special setting, it might not appear surprising that when the number of (specialized) suppliers equals the numbers of consumer segments, they are able to divide the market. However, this rather simple scenario serves a different purpose: to study the applicability of the learning mechanism itself. Only if the proposed mechanisms are working in rather simple settings, is it worthwhile to study their performance in more complex settings. Furthermore, allocated suppliers receive a discount based on their contribution to the SVN. Due to the nature of demand in segmented markets, a

Table 3. Results for Random (R) and Segmented (S) Markets for Three Homogeneous Suppliers Using GA With 100 or 500 Episodes per Simulation Round

	Allocation Percentage				Average Supplier Profit			
	Average	*SD*	Min	Max	Average	*SD*	Min	Max
R, 100	0.694	0.016	0.663	0.719	1.093	0.364	0.958	1.211
S, 100	0.777	0.154	0.639	0.979	4.300	2.324	2.592	7.242
R, 500	0.693	0.013	0.673	0.728	1.092	0.306	0.956	1.233
S, 500	0,858	0,151	0,650	0,988	5,251	1,654	2,615	7,448

p Values	Allocation Percentage	Average Supplier Profit
3P, 100, R vs. S	.856 (l)	<.001 (l)
3P, 500, R vs. S	.054 (l)	<.001 (l)
3P, S, 100 vs. 500	<.001 (l)	<.001 (l)

Note. (l) indicates one-sided tests (*less than*).

Table 4. Results for Heterogeneous Suppliers in Segmented Markets Using GA With 100 or 500 Episodes per Simulation Round

	Allocation Percentage				Average Supplier Profit			
	Average	*SD*	Min	Max	Average	*SD*	Min	Max
100 ep.	0.992	0.003	0.979	0.994	6.311	0.448	4.631	7.565
500 ep.	0.994	0.001	0.993	0.995	6.598	0.361	5.502	7.661

p Values	Allocation Percentage	Average Supplier Profit
Ho, 100 vs. He, 100	<.001 (l)	.003 (l)
Ho, 500 vs. He, 500	.151 (l)	.005 (l)

Note. GA = genetic algorithm.
(l) indicates one-sided tests.

single supplier contributes to a large degree to a request from the consumer segment she specializes on, especially if other suppliers' offers are not comparable and target distinct consumer segments. Hence, in the optimal solution, each supplier focuses on a distinct segment and receives comparably high discounts in addition to pure profit margins for the service offer.

A third observation concerns the solutions in the respective replications. A detailed analysis of the simulation results shows that while in 8 of the 20 replications service suppliers indeed do learn the optimal solution (confirmed by high profits and a rate of over 99% of allocated service requests), in the remaining 12 replications suppliers get stuck in local optima, in which two suppliers compete for the same segment. In these cases, only about 67% of consumer requests get successfully served because the competition for one segment leaves the third segment unserved. Obviously, in these replications, suppliers get stuck in local optima and their GAs are not able to overcome this inefficient solution. Moreover, the results indicate that an increase in the number of episodes per simulation round increases the ability of suppliers to learn the optimal division of the market. The difference in allocated service offers and supplier profit is also statistically significant as the *p* values of <.001 for an increase in allocated offers and an increase in supplier profits show.

Alternative to homogeneous suppliers with equal cost functions, we introduce heterogeneity among suppliers by varying cost functions that give suppliers cost advantages on certain attributes. This means that these suppliers can offer services with high levels on a certain attribute at lower prices than their potential competitors.

Table 4 shows the simulation results of heterogeneous service suppliers. Compared to the previous case, the results show that in all replications service suppliers learn the global optimum of tacit market division. The results clearly indicate that heterogeneity among suppliers significantly improves their ability to divide the market among themselves and to avoid competition, although they, as before, start with random service configurations. Both the allocation percentages and average supplier profits are significantly higher for heterogeneous suppliers as confirmed by a Wilcoxon's test which yields *p* values of <.005 (except for the difference in allocation percentages in the 500 episodes case, where the *p* value is not significant). Hence, heterogeneity of service suppliers facilitates their ability to learn a tacit division of the market. This refers to Hypotheses 3a and 4a.

Scenario 2: Collusion in Competitive Environments

In the previous scenario, the suppliers were able to set an increment on the price of their service offer, which enables them to influence their profits. Another form of collusion occurs when

Table 5. Results for Three (3P) and Six (6P) Heterogeneous Suppliers per Candidate Pool, Using GA or P&A With MU or With MC

	Allocation Percentage				Average Supplier Profit			
	Average	*SD*	Min	Max	Average	*SD*	Min	Max
3P, GA, MU	0.992	0.003	0.979	0.994	6.311	0.448	4.631	7.565
3P, GA, MC	0.997	0.000	0.996	0.997	5.988	0.471	4.890	6.909
6P, GA, MU	1.000	0.000	0.999	1.000	0.479	0.398	0.312	1.512
6P, GA, MC	1.000	0.000	1.000	1.000	0.165	0.123	0.129	0.950

p-Values	Allocation Percentage	Average Supplier Profit
3P, GA, MU vs. MC	<.001	.003 (g)
6P, GA, MU vs. MC	<.001	<.001 (g)

Note. GA = genetic algorithm; MC = marginal costs; MU = mark-ups. (g) indicates one-sided tests (*greater than*).

suppliers agree (explicitly or tacitly) to set prices at other than the competitive level. If more than one supplier competes for a consumer segment, suppliers have a myopic incentive to decrease the price of their offers below that of their competitors in order to draw the consumer requests to their offer. Thus, in competitive settings, prices should equal the MC of service suppliers, because if one service supplier sets her price above costs, the competitor would have the incentive to cut the supplier's price in order to gain consumers. Alternatively, suppliers could tacitly agree on price levels strictly above MC, which would increase their average profits. This simulation scenario thus compares competitive settings where suppliers can charge MU on their costs with settings where suppliers have to offer their services at MC. In one setting, three suppliers compete for the consumers, in the other setting the competition is increased by three additional suppliers.

As shown in Table 5, average supplier profits are significantly higher (p values of .003 and <.001 for three and six suppliers per candidate pool, respectively) in cases where suppliers can charge MC on their costs. These results indicate that suppliers tacitly agree on strictly positive price increments because this leads to a bilateral increase of average profits for all competing suppliers if all suppliers adhere to the collusive level of price increments. This is in line with Hypotheses 5a.

Learning via P&A

Originally applied in Cournot (quantity competition) contexts, we adopted P&A for our supplier learning setting in SVNs. Using a random search in the vicinity of the current service offer, P&A evaluates the payoffs of (allocated) service offers and moves its anchor levels in the direction of the highest payoff.

Figure 4 in the Appendix shows pseudocode for P&A. Basically, suppliers have a current service offer, which they randomly perturb in order to learn if attribute level adjustments to the current offer are profitable (and so probe the system for a response). The algorithm keeps two payoff records for each attribute that is perturbed, depending on whether the attribute was increased or decreased (Lines 3–5). In each episode, (Lines 8-26) all suppliers perturb a certain number of attributes of their current service offer within a specific search range and submit the resulting offer to the SVN (Lines 11–15). After the SVN calculates the resulting allocation and payoffs, suppliers record the payoff of the perturbed service offer depending on the direction of change (Lines 17–25).

After each simulation round, the algorithm calculates the mean payoffs of service offers with increased and decreased attribute levels and adjusts the current level of the service offer in the direction of higher average payoffs. The direction of attribute adjustment is chosen randomly in the event of a tie (Lines 27–37). Afterward, the payoff records are reset and a new simulation round begins (Lines 38–39).

In order to test the sensitivity of P&A on algorithm parameters, a full factorial design with two factors is used where each factor has two possible settings. The first factor is the search range of the algorithm which defines the maximum possible attribute adjustments based on the current attribute levels. The two settings for the search range are $[-3, 3]$ and $[-5, 5]$. For example, if the current level of a service attribute, which itself is encoded as an integer in the range of [0, 20], is x, the search range of $[-3, 3]$ means that the supplier can select adjusted service offers with attribute levels in $[x - 3, x + 3]$ (provided the result is confined to the permitted range, [0, 20]). The second factor to be studied is the number of attributes which are simultaneously adjusted during the simulation. In the first setting, one of the four service attributes is randomly chosen and adjusted each simulation round, whereas the second setting adjusts all attributes simultaneously.

In general, both aforementioned factors are assumed to affect the performance of P&A. On one hand, increasing the search range enables suppliers to better sample the service space and potentially find better service offers. On the other hand, adjusting all service attributes simultaneously increases the flexibility of suppliers and potentially leads to faster convergence. In order to analyze P&A, we formulate the following research hypotheses:

Hypothesis 3b: Supplier heterogeneity incentivizes suppliers to divide the market and thereby increases the percentage of allocated services.

Hypothesis 4b: Supplier heterogeneity increases average supplier profits.

Hypothesis 5b: Service suppliers are able to earn higher profits if they can charge additional MU on their MC.

Hypothesis 6: For P&A, larger search ranges yield higher average profits in low competition settings.

Hypothesis 7: P&A yields higher average profits if all attributes are adjusted simultaneously.

Scenario 3: Tacit Market Division Among Suppliers Using P&A

As in simulation Scenario 1, all service suppliers in this scenario have identical cost functions and thus no preference/advantage for particular consumer segments. Suppliers only differ in the applied learning mechanism, thus enabling a comparison of the results with the first scenario.

Table 6. Results for Homogeneous Suppliers in Segmented Markets, Using P&A With Search Range [−3,3] (SR3) or [−5,5] (SR5) for Each Attribute, Adjusting One (One) or All (All) Attributes per Simulation Round

	Allocation Percentage				Average Supplier Profit			
	Average	SD	Min	Max	Average	SD	Min	Max
One, SR3	0.577	0.181	0.313	0.936	3.356	2.577	0.453	6.988
One, SR5	0.606	0.102	0.314	0.929	2.973	2.786	0.458	6.837
All, SR3	0.580	0.114	0.303	0.832	2.858	2.192	0.451	5.773
All, SR5	0.494	0.065	0.437	0.612	2.482	1.462	1.704	4.077

Note. GA = genetic algorithm; MC = marginal costs; MU = mark-ups; SR = search range.

Table 6 shows the simulation results for varying parameter settings of P&A. We used a 2×2 design to study the effect of the search range and the number of simultaneously adjusted attributes. Interestingly, contradicting to Hypothesis 6 a larger search range yields lower average profits than smaller search ranges do, although it should enable suppliers to sample a larger fraction of potential service configurations. Adjusting all attributes simultaneously in each episode also decreases the ability of service suppliers to learn the optimal, collusive outcome, which also contradicts Hypothesis 7. This is also a surprising result, as changing the number of adjusted attributes should, at the most, alter the speed of potential convergence as the suppliers are able to sample the service configuration space faster.

Comparing the results with the GA scenarios, we see that although allocation percentages are significantly lower (p value of <.001) in the P&A case, the Wilcoxon's test cannot confirm that GA also yields higher average supplier profits (p value of .495). Thus, due to the higher variance of results in P&A scenarios, P&A seems to be much more volatile with respect to the random starting locations of service suppliers. In some cases, it is not able to overcome the worst local optimum where all three suppliers specialize on the same segment (resulting in allocation percentages below 33%).

Like Scenario 1, heterogeneous suppliers were studied with respect to their ability to learn the optimal market division. Table 7 presents the results of the scenario runs. Compared to the previous scenario with homogeneous suppliers, the results clearly indicate that heterogeneity greatly facilitates supplier learning, which we also observed when the suppliers used a GA. Compared to homogeneous suppliers, both allocation percentages and average supplier profits significantly increase as the p values of <.001 suggest. Hence, incentivized by their individual strengths and cost advantages, heterogeneous suppliers learn to effectively and tacitly divide the market among themselves when using P&A, confirming Hypotheses 3b and 4b.

Scenario 4: Detecting Collusion Between Competing Suppliers

Studying the occurrence of tacit collusion is also interesting if a pure market division among suppliers is not possible due to

Table 7. Results for Heterogeneous Suppliers in Segmented Markets, Using P&A With Search Range [−5,5] (SR5) for Each Attribute, Adjusting All Attributes per Simulation Round

	Allocation Percentage				Average Supplier Profit			
	Average	SD	Min	Max	Average	SD	Min	Max
All, SR5	0.926	0.129	0.704	1.000	6.149	0.281	2.591	8.495

p Values	Allocation Percentage	Average Supplier Profit
All, SR5, Ho vs. He	<0.001	<0.001

competition. As described in Scenario 2, in this case we examine collusion among suppliers by comparing supplier profits if they can charge price increments on their costs with scenarios where they have to offer their services at their MC level.

Table 8 shows that average supplier profits in a competitive scenario are significantly higher (p value <.001) if they can charge price increments despite competition for the consumer segments, which is in accordance with Hypothesis 5b. Even the average profit of the worst-performing service supplier in the MU case is higher than the profit of the best-performing supplier in the MC case. This indicates that suppliers specializing in certain consumer segments tacitly agree upon a strictly positive price level, which contradicts classical game theory and economics competition argumentation which indicates that in these scenarios price should equal MC.

Conclusion

Due to the predicted importance of SVNs,[4] it is important to understand the potential behavior of the market participants and their mutual interactions. Therefore, in this article, we studied strategic learning of e-Service suppliers and the potential for collusion among them in SVNs. SVNs aggregate single electronic (ICT) services into a complex service bundle, based on the consumer's preferences. Collusion in an SVN can occur in a variety of ways, of which we studied two particular forms of tacit collusion without direct communication between SVN participants offering the same electronic service, such as payment or storage.

In the first case, few service suppliers compete for a certain number of consumer segments, and suppliers can mutually benefit by dividing the market among themselves. As the simulation results for both learning strategies show, service suppliers, in most cases, indeed learn that such a tacit agreement is beneficial in terms of increased profit. Because the individual objective of profit maximization is aligned with the global goal to serve all consumer requests, service suppliers are incentivized to learn the division of the market, which is confirmed by our simulation results. In virtue of using two very different learning regimes, we can count this as a robust finding. A direct comparison of the two learning strategies shows that the GA performs slightly better if there are few homogenous providers, that is, in narrow oligopoly markets. However, when several

Table 8. Results for Eight Heterogeneous Suppliers per Candidate Pool, Using P&A With MU or With MC

	Allocation Percentage				Average Supplier Profit			
	Average	SD	Min	Max	Average	SD	Min	Max
MU	1.000	0.000	1.000	1.000	0.595	0.208	0.285	0.903
MC	1.000	0.000	1.000	1.000	0.160	0.093	0.115	0.201

p Values	Allocation Percentage	Average Supplier Profit
MU vs. MC	<.001	<.001

Note. GA = genetic algorithm; MC = marginal costs; MU = mark-ups.

heterogeneous suppliers have to compete for consumer segments, the difference between the strategies is much smaller.

The second case studies potential collusion in more competitive environments with a larger number of service suppliers. In this case, competition for the consumer segments arises and service suppliers are not able to avoid intense competition by each specializing in a distinct segment. As our simulation results show, service suppliers tacitly agree on a certain level of price fixing under both learning mechanisms. By mutually setting price levels for services above their costs, allocated service suppliers are able to increase their profits as compared to purely competitive scenarios where they are forced to offer their services at the level of their respective MC. Comparing the two learning strategies, we see that in both cases the suppliers learn to raise prices above the competitive level, which indicates a robust finding.

These results are especially interesting as they are relevant for SVNs where several suppliers compete for certain consumer segments. Although service suppliers only rely on the success of their own service offers, they are able to learn that refraining from excessive price competition yields higher profits. Interestingly, in contrast to game theoretic and economic reasoning which postulates that each service supplier participating in the price fixing scheme has short-term incentives to cut the collusive price levels, the simulation results show that the tacit collusive agreement continues to exist even in the long term after a large number of incoming consumer requests. This shows that service suppliers as modeled in our simulation prefer long-term, beneficial collusive agreements to competitive scenarios with potential short-term benefits. Again, in virtue of using two very different learning regimes, we can count this as a robust finding.

Moreover, our simulation results show that, although service suppliers start with random service configurations, they are able to learn profitable, tacit collusive schemes by applying their learning strategies based on the success of their current offers. Although the learning mechanisms we studied are rather simple and do not rely on specific information about the market, collusive behavior among suppliers nevertheless emerges. Hence, if cognitively simple suppliers can learn to collude, more sophisticated suppliers should be able to perform at least as well. In fact, the occurrence of collusion seems to be inherent to the market scenario. The results indicate that the level of collusion depends on the learning mechanisms, which behave differently under certain settings and have their specific strengths and weaknesses. The GA is less likely to get stuck in local optima due to its ability to efficiently sample the large solution space. In contrast, the effectiveness of P&A heavily depends on the random starting configurations of service suppliers, yet is more flexible in the incremental search for optimal configurations.

Finally, what do our findings mean for practice? Clearly, no single study can be dispositive on such matters. Much more research, building on what we have so far, will have to be done. We can, however, consider what is suggested by our findings and treat them as hypotheses having a non-negligible degree of support. The main hypothesis supported by the findings is that in an SVN it is not only possible, but rather easy for suppliers to achieve tacit collusion at the expense of the consumers (Hypothesis 5). An ancillary hypothesis is that they can do this by segmenting the market and becoming niche monopolists (Hypotheses 1 and 2). What this means for suppliers is that opportunities may exist to achieve economic rents in these situations. What this means for consumers and for e-Service platform providers is that they will need to observe carefully whether this is happening and if it is, they will need to develop strategies to counter the collusion. For the SVN provider, this entails a delicate balancing act between making the consumers happy (with the benefits of a competitive market) and making the suppliers happy (with the benefits of a less than fully competitive market).

All of these issues are promising avenues for future research, for there are still open questions and the need for further investigations. One possibility in proceeding is to increase the depth of the simulation model by capturing additional features. For example, more complex collusion mechanisms can be modeled where suppliers explicitly decide on the participation in cartels and bidding rings. Likewise, the learning mechanisms of service suppliers can be upgraded to capture additional information about competitors and the market instead of just evaluating payments, or new types of learning mechanisms can be added to the model such as classifier systems, which are able to learn conditional strategies based on different market settings.[5]

It is quite possible that several SVNs might evolve concurrently and compete for consumer attention and demand. This leads to the interesting situation in which not only do suppliers compete with other suppliers in the same SVN, but they also have to decide which services to offer in which SVNs. Additionally, SVNs themselves might engage in competition for both service supplier offerings and consumer requests, as both are important for the success of an SVN due to substantial network effects. Finally, a comparison of the results obtained in the work at hand with empirically observed data from real SVNs would be informative and allow for the validation and calibration of the simulation model.

We think that our results provide valuable insights into potential collusive service supplier behavior in SVNs. They can help understand and improve both the strategies of potential service suppliers who consider offering services in SVNs and the very design of SVNs and corresponding implementation issues in practice.

Appendix

A. Pseudocode for Probe and Adjust

```
randomly select service offer levels for each supplier
initialize simulation parameters
forall the attributes l do
    initialize record lists returnsUp^l = returnsDown^l = []
end
while simulation stopping criterion is not met do
    simulation round ++
    while episode < maxEpisodes do
        episode ++
        draw consumer request
        forall the suppliers do
            select attribute(s) to adjust
            draw adjusted service offer within search range
            submit service offer to SVN
        end
        SVN: calculate allocation and pricing
        forall the suppliers do
            forall the attributes l that were adjusted do
                if attributed adjusted down then
                    append returns to returnsDown^l
                else
                    append returns to returnsUp^l
                end
            end
        end
    end
    forall the suppliers do
        forall the attributes l that were adjusted do
            if mean(returnsUp^l) > mean(returnsDown^l) then
                increase attribute l
            else if mean(returnsDown^l) > mean(returnsUp^l) then
                decrease attribute l
            else
                randomly decrease or increase attribute l
            end
        end
    end
    reset record lists for all attributes l
    reset calculation parameters
end
```

Figure 4. Pseudocode for probe and adjust.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

Notes

1. See for example, Danaher (1998). The elicitation of these preferences, however, is difficult. It is possible to explicitly design experiments in order to collect data on preferences, for example, Verma, Thompson, and Louviere (1999) set up an experiment to collect such data for pizza delivery services and Gustafsson and Johnson (2004) compare two methods of how to derive customers' valuation of service attributes.
2. Considering the relevance of switching costs in our simulations, we studied the simulation results and found the following interesting effect: Because e-Service suppliers receive a discount based on the relevance of their service to the consumers' preferences, the ratio of discount versus price of the service combination can be seen as a proxy for the added value of the e-Service supplier. In other words, the higher the discount relative to the prices, the more valuable the e-Service supplier is to the consumer. In the results, we found that the ratio discount/prices is in the range of [0.13, 2.7], which means that even in the most competitive scenario, the allocated e-Service supplier could increase its price by 13% and still be allocated. In this light, we do not further discuss switching costs in this article.
3. We thank an anonymous reviewer for pointing out the importance of this issue, as the assumption of perfect information would question the practical applicability of this approach.
4. In 2008, Gartner predicted that within the next years one third of all business applications will be delivered through on-demand service subscription instead of purchasing product licenses, see http://www.gartner.com/it/page.jsp?id=593207.
5. Furthermore, the current simulation model is based on the nonbudget-balanced mechanism described in the third section, where budget balance means that the internal payoffs within the market should sum up to zero, which assures that the SVN does not need external subsidization. For a successful and profitable application of SVNs, Blau (2009) developed an extension for the mechanism that assures budget balance. As a result, the mechanism no longer is incentive compatible, that is, service suppliers do not necessarily have the incentive to submit their service offers truthfully to the SVN. Hence, it would be interesting to study the resulting supplier strategies in these nonincentive compatible SVNs. We leave this for future research.

References

Allee, Verna (2009), "Value-Creating Networks: Organizational Issues and Challenges," *Learning Organization*, 16 (6), 427-442.

Anderson, Chris (2007), *The Long Tail : How Endless Choice is Creating Unlimited Demand*. London, UK: Random House Business Books.

Anderson, James C., Håkan Håkansson, and Jan Johanson (1994), "Dyadic Business Relationships within a Business Network Context," *Journal of Marketing*, 58 (4), 1-15.

Arifovic, Jasmina (1994), "Genetic Algorithm Learning and the Cobweb Model," *Journal of Economic Dynamics and Control*, 18 (1), 3-28.

Astley, Graham W. and Charles J. Fombrun (1983), "Collective Strategy: Social Ecology of Organizational Environments," *Academy of Management Review*, 8 (4), 576-587.

Bansal, H.S. and S.F. Taylor (1999), "The Service Provider Switching Model (SPSM)," *Journal of Service Research*, 2 (2), 200-218.

Basole, R.C. and W. B. Rouse (2008), "Complexity of Service Value Networks: Conceptualization and Empirical Investigation," *IBM Systems Journal*, 47 (1), 53-70.

Bengtsson, M. and S. Kock (2000), "Coopetition in Business Networks to Cooperate and Compete Simultaneously," *Industrial Marketing Management*, 29 (5), 411-426.

Holm, Blankenburg, K. Eriksson Desirée, and J. Johanson (1996), "Business Networks and Cooperation in International Business Relationships," *Journal of International Business Studies*, 27 (5), 1033-1053.

Blau, Benjamin, Tobias Conte, and Clemens van Dinther (2010), "A Multidimensional Procurement Auction for Trading Composite Services," *Electronic Commerce Research and Applications Journal (ECRA) on Emerging Economic, Strategic and Technical Issues in Online Auction and Electronic Markets* (Elsevier Science Publishers B. V.), 9 (5), 460-472.

Blau, Benjamin, Jan Krämer, Tobias Conte, and Clemens van Dinther (2009), "Service Value Networks," Proceedings of the 11th IEEE Conference on Commerce and Enterprise Computing, Vienna, Austria, 194-201.

Blau, Benjamin, Clemens van Dinther, Tobias Conte, Xu Yongchun, and Christof Weinhardt (2009), "How to Coordinate Value Generation in Service Networks–A Mechanism Design Approach," *Business and Information Systems Engineering (BISE)*, 1 (5), 343-356.

Blau, Benjamin (2009), "Coordination in Service Value Networks - A Mechanism Design Approach," Dissertation, Faculty of Economics and Business Engineering, Karlsruhe Institute of Technology (KIT).

Brandenburger, Adam M. and Barry J. Nalebuff (1997), "Co-Opetition: A Revolution Mindset That Combines Competition and Cooperation: The Game Theory Strategy That's Changing the Game of Business." Doubleday, New York.

Brenner, Thomas (2006), "Agent Learning Representation: Advice on Modelling Economic Learning." In *Handbook of Computational Economics*, Vol. 2, *Agent-Based Computational Ecoomics*, Tesfatsion Leigh and Judd Kenneth L., eds. Elsevier/North-Holland. 895-947.

Brenner, Thomas and Ulrich Witt (2003), "Melioration Learning in Games with Constant and Frequency-Dependent Pay-Offs," *Journal of Economic Behavior & Organization*, 50 (4), 429-448.

Brenner, Thomas (1999), *Modelling Learning in Economics*. Edward Elgar, Cheltenham.

Busquets, Javier, Juan Rodon, and Jonathan Wareham (2009), "Adaptability in Smart Business Networks: An Exploratory Case in the Insurance Industry," *Decision Support Systems*, 47 (4), 287-296.

Camerer, Colin F. (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton: Princeton University Press, Russell Sage Foundation.

Clarke, Edward H. (1971), "Multipart Pricing of Public Goods," *Public Choice* (Springer Netherlands) 11 (1), 17-33.

Collier, J.E. and C.C. Bienstock (2006), "Measuring Service Quality in E-retailing," *Journal of Service Research*, 8 (3), 260-275.

Conte, Tobias, Benjamin Blau, Gerhard Satzger, Clemens van Dinther, and Christof Weinhardt (2011), "Rewarding Participation in Service Value Networks," *e-Service Journal*, 7 (2), 2-27.

Conte, Tobias, Benjamin Blau, Gerhard Satzger, and Clemens van Dinther (2009), "Enabling Service Networks Through Contribution-Based Value Distribution," Proceedings of the 15th Americas Conference on Information Systems, San Francisco, CA.

Cook, Karen S. and Richard M. Emerson (1978), "Power, Equity, Commitment in Exchange Networks," *American Sociological Review* 43 (5), 721-738.

Curbera, Francisco, Rania Khalaf, Nirmal Mukhi, Stefan Tai, and Sanjiva Weerawarana (2003), "The Next Step in Web Services," *Communications of the ACM*, 46 (10), 29-34.

Dai, Yuan-Shun and Xiao-Long Wang (2006), "Optimal Resource Allocation on Grid Systems for Maximizing Service Reliability Using a Genetic Algorithm," *Reliability Engineering & System Safety*, 91 (9), 1071-1082.

Danaher, P. J. (1998), "Customer Heterogeneity in Service Management," *Journal of Service Research*, 1 (2), 129-139.

De Jong., Kenneth A and William M. Spears (1991), "An Analysis of the Interacting Roles of Population Size and Crossover in Genetic Algorithms," PPSN I: Proceedings of the 1st Workshop on Parallel Problem Solving from Nature, London, UK: Springer-Verlag, 38-47.

D'Souza, Bryan and Timothy W. Simpson (2003), "A Genetic Algorithm Based Method for Product Family Design Optimization," *Engineering Optimization*, 35 (1), 1-18.

Duffy, John (2006), "Agent-Based Models and Human Subject Experiments," In *Handbook of Computational Economics*, Vol. 2, *Agent-Based Computational Economics*, TesfatsionL. and JuddK. L., eds. Elsevier/North-Holland, 949-1011.

Emerson, R. M. (1976), "Social Exchange Theory," *Annual Review of Sociology*, 2, 335-362.

Erev, Ido and Alvin E. Roth (1998), "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria," *American Economic Review*, 88, (4), 848-881.

Fassnacht, M. and I. Koese (2006), "Quality of Electronic Services," *Journal of Service Research*, 9 (1), 19-37.

Fogel, David B. (2002), *Blondie24: Playing at the Edge of AI*. Kaufmann, San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

Franklin, Daniel (2005), "Business 2010 - Embracing the Challenge of Change," The Economist Intelligence Unit. http://www.eiu.com/site_info.asp?info_name=eiu_SAP_business2010&rf=0.

Gartner-Inc (2010), "Press Release: Gartner Reveals Five Business Process Management Predictions for 2010 and Beyond," http://www.gartner.com/it/page.jsp?id=1278415. Contact: Christy Pettey and Laurence Goasdu.

Goldberg, David E. (1989), *Genetic Algorithms in Search, Optimization and Machine Learning*. Boston, MA, USA: Addison-Wesley.

Groves, Theodore (1973), "Incentives in Teams," *Econometrica* (The Econometric Society), 41 (4), 617-631.

Gustafsson, A. and M.D. Johnson (2004), "Determining Attribute Importance in a Service Satisfaction Model," *Journal of Service Research*, 7 (2), 124.

Håkansson, Håkan and Jan Johanson (1993), "Industrial Functions of Business Relationships," in *Advances in International Marketing*, Vol. 5, S. Cavusgil and D. Sharma, eds. Greenwich: JAI Press, 15-31.

Hamel, G., Y.L. Doz, and C.K. Prahalad (1989), "Collaborate with your Competitors and Win," *Harvard Business Review*, 67 (1), 133-139.

Heck, E. van and P. Vervest (2007), "Smart Business Networks: How the Network Wins," *Communications of the ACM*, 50 (6), 29-37.

Heim, G.R. and K.K. Sinha (2001), "A Product-Process Matrix for Electronic B2C Operations." *Journal of Service Research*, 3 (4), 286-299.

Heuser, Lutz, Stefan Lacher, and Simone Perlmann (2007), "Flexible Prozessgestaltung als Basis Innovativer Geschäftsmodelle - Von der Service-Orientierten Architektur zur Vision des Business Webs," Proceedings of the Conference Wirtschaftsinformatik 2007.

Hoetker, G. (2006), "Do Modular Products Lead to Modular organizations?," *Strategic Management Journal*, 27 (6), 501.

Holland, John H. (1992), *Adaptation in Natural and Artificial Systems*. Cambridge, MA, USA: MIT Press.

Holloway, B.B. and S.E. Beatty (2008), "Satisfiers and Dissatisfiers in the Online Environment," *Journal of Service Research*, 10 (4), 347-364.

Hoogeweegen, Martijn R. and Peter H.M. Vervest (2005), "How Much Business Modularity." in *Smart Business Networks*, Chap. 23, Vervest Peter, Heck Eric van, Preiss Kenneth and Pau Loius-Francois, eds. Springer Berlin Heidelberg: Springer, 339-348.

Hosting (2009), "2009 Cloud Computing Trends Report." 2009 Cloud Computing Trends Report.

Hu, Junling and Michael P. Wellman (1998), *Multiagent Reinforcement Learning: Theoretical Framework and an Algorithm*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 242-250.

Kimbrough, Steven Orla (2012), *Agents, Games, and Evolution: Strategies at Work and Play*. Chapman and Hall/CRC, London: CRC Press.

Kimbrough, Steven O. and Frederic H. Murphy (2009), "Learning to Collude Tacitly on Production Levels by Oligopolistic Agents," *Computational Economics* 33 (1), 47-78.

Krämer, Jan, Tobias Conte, Benjamin Blau, Clemens van Dinther, and Christof Weinhardt (2010), "Service Value Networks: Unleashing the Combinatorial Power of Service Mashups," Working Paper Series on Social Science Research Network, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1582902.

Law, Averill M. (2007), *Simulation Modeling and Analysis*, 4th ed. Boston, MA: McGraw-Hill.

Lusch, Robert, Stephen Vargo, and Mohan Tanniru (2010), "Service, Value Networks and Learning," *Journal of the Academy of Marketing Science* 38, (Springer Netherlands), 19-31.

MacKenziee, Matthew C., Ken Laskey, Francis McCabe, Peter F. Brown, and Rebekah Metz (2006), "Reference Model for Service Oriented Architecture 1.0.," Reference Model for Service Oriented Architecture 1.0.

McAfee, Andrew (2005), "Will Web Services Really Transform Collaboration?" *MITSloan Management Review*, 46 (2), 78-84.

Menascé, Daniel A. (2002), "QoS Issues in Web Services," *IEEE Internet Computing*, 6 (6), 72-75.

Miles, Raymond E. and Charles C. Snow (1992), "Causes of Failure in Network Organizations," *California Management Review*, 34, 53-72.

Miles, Raymond E. and Charles C. Snow (1986), "Organizations: New Concepts for New Forms," *California Management Review*, 28 (3), 62-74.

Ostrom, A. L., Bitner M. J., Brown S. W., et al. (2010), "Moving Forward and Making a Difference: Research Priorities for the Science of Service," *Journal of Service Research*, 13 (1), 4-36.

Papazoglou, Michael P. (2008), *Web Services: Principles and Technology*. Prentice Hall.

Papazoglou, Michael P. and J. Dubray (2004), "A Survey of Web Service Technologies," Technical Report, University of Tronto, Department of Information and Communication Technology.

Papazoglou, Michael P. and D. Georgakopoulos (2003), "Service Oriented Computing," *Communications of the ACM*, 46 (10), 25-28.

Roth, Alvin E. and Ido Erev (1995), "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior* 8, 164-212.

Roth, S., H. Woratschek, and S. Pastowski (2006), "Negotiating Prices for Customized Services," *Journal of Service Research*, 8 (4), 316-329.

Rust, Roland T. and P. K. Kannan (2002), "The Era of e-Service," in *E-Service: New Directions in Theory and Practice*, Rust Roland T. and Kannan P. K., eds. New York: M. E. Sharpe, 3-21.

Rust, Roland T. and Katherine N. Lemon (2001), "E-Service and the Consumer," *International Journal of Electronic Commerce* (M. E. Sharpe, Inc.), 5 (3), 85-101.

Rust, Ronald T. (2001), "The Rise of e-Service." *Journal of Service Research*, 3 (4), 283-284.

Schroth, Christoph and Till Janner (2007), "Web 2.0 and SOA: Converging Concepts Enabling the Internet of Services," *IT Professional*, 9 (3), 36-41.

Skyrms, Brian (2010), *Signals: Evolution, Learning & Information*. Oxford: Oxford University Press.

Sommerville, Ian (2007), *Software Engineering*, 8th ed. Boston, MA, USA: Addison Wesley.

Spohrer, J., L. C. Anderson, N. J. Pass, T. Ager, and D. Gruhl (2008), "Service Science," *Journal of Grid Computing*, 6 (3), 313-324.

Tesfatsion, Leigh (2002), "Agent-Based Computational Economics: Growing Economies From the Bottom Up," *Artificial Life* (MIT Press), 8 (1), 55-82.

Thorndike, E. L. (1927), "The Law of Effect," *American Journal of Psychology*, 39 (1/4), 212-222.

van Dinther, Clemens (2010), "Designing Service Value Networks—Coordination, Competition, Cooperation," Habilitation, Karlsruhe Institute of Technology, Germany.

Verma, R., G.M. Thompson, and J.J. Louviere (1999), "Configuring Service Operations in Accordance with Customer Needs and Preferences," *Journal of Service Research*, 1 (3), 124-141.

Vickrey, William (1961), "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16 (1), 8-37.

Welch, P. D. (1983), "The Statistical Analysis of Simulation Results," in *The Computer Performance Modeling Handbook*, Lavenberg, S. S., ed. New York: Academic Press, 268-328.

Wilcoxon, Frank (1945), "Individual Comparisons by Ranking Methods," *Biometrics Bulletin* (International Biometric Society), 1 (6), 80-83.

Wooldridge, Michael and Nicholas R. Jennings (1995), "Intelligent Agents: Theory and Practice," *Knowledge Engineering Review*, 10 (2), 115-152.

Zeng, Liangzhao, Boualem Benatallah, Marlon Dumas, Jayant Kalagnanam, and Quan Z. Sheng (2003), "Quality Driven Web Services Composition," WWW '03: Proceedings of the 12th international conference on World Wide Web (Budapest, Hungary), ACM, 411-421.

Bios

Christian Haas is a PhD student at the Karlsruhe Institute of Technology, Germany. He received a Master of Science in Computer Science from the Georgia Institute of Technology in 2008, and a Diploma in Business Engineering from the Karlsruhe Institute of Technology in 2010. His research interests include the application of learning methodologies in strategic contexts, incentive engineering in social networks, and service science in general. In his PhD, he is currently working on two-sided market mechanisms and dynamic behavior in social collaboration systems

Steven O. Kimbrough is Professor of Operations and Information Management at the University of Pennsylvania. He also holds a secondary appointment at, and is active in, the Philosophy Department at Penn. He works as a decision scientist, with an overarching interest in rationality and computation. His principal research areas are in procedural (post-classical) game theory, evolutionary computation, and applied logic. He also works in the area of text mining and analytics. He is an inventor on three related patents in this area.

Clemens van Dinther received the venia legendi for Business Administration from the Karlsruhe Institute of Technology (KIT). He teaches at KIT and as a guest lecturer at RWTH Aachen. Currently, Clemens is working at ista Deutschland GmbH as Director for Energy Management and Business Process Outsourcing. His research interests are in the application of electronic services, computational economics, strategic decisions, application of learning methods, Smart Grid, and Regulation in the Energy Sector.