Inf Syst E-Bus Manage (2013) 11:29–50
DOI 10.1007/s10257-011-0179-z

ORIGINAL ARTICLE

One step ahead: using predictive inference to select service providers in loosely coupled electronic markets

Daniel S. Soper

Abstract The rapid growth of service-oriented electronic markets implies a common belief among managers that participation in these dynamic, loosely coupled markets can yield many benefits for their firms. At present, however, very little is known about organizational behavior in these nascent markets. By utilizing a sophisticated simulation of a loosely coupled interorganizational service market, this paper demonstrates that customer organizations seeking to purchase and consume services in such markets can benefit from the application of predictive inference in the provider selection process. Specifically, it is shown that a customer organization employing a simple predictive method to select service providers can, in the aggregate, achieve notably superior outcomes in terms of price, quality of service received, and several other metrics when compared to competitors who act opportunistically in selecting their business partners. The implications of these findings for managers and researchers are presented and discussed in the context of the rising popularity of loosely coupled electronic markets.

Keywords Market prediction · Service provider selection · Electronic markets · Simulation

1 Introduction

It has now been a decade and a half since service-oriented electronic markets were first proposed (Schulte and Natis 1996), and in the intervening years many skeptical business writers have viewed these markets with disapprobation, occasionally even

D. S. Soper (✉)
Department of Information Systems and Decision Sciences,
Mihaylo College of Business and Economics, California State University, Fullerton,
800 N. State College Blvd, Fullerton, CA 92831, USA
e-mail: dsoper@fullerton.edu

dismissing the service-oriented business model as nothing more than a managerial fad (Smith 2008; Alexander and Korine 2008). With the rise of web services, however, service-oriented business has found itself in the company of a natural bedfellow, and together these two philosophies are undergirding a monumental shift in the way that organizations interconnect with one another to conduct business (Newcomer and Lomow 2004; Papazoglou and Georgakopoulos 2003; Lawler and Howell-Barber 2007). Indeed, the establishment of loosely coupled electronic markets as enabled by web services is now a *fait accompli*, and since the value of these markets was expected to reach $154 billion by the end of 2010 (Cantara et al. 2007), it seems increasingly unlikely that the means through which businesses interconnected in the past will ever again rise to prominence.

When combined with a service-oriented business model, web service technologies enable core electronic service capabilities to be deployed at the edge of an organization, where they can then be purchased by interested customers. This paradigm has many important advantages for both customers and service providers interacting in the market (Yoon and Carter 2007). Among customers, web service technologies have been credited with enabling customer information needs to be more readily met (Liang and Tanniru 2006–2007), lowering long-term IT expenses (Castro-Leon et al. 2007), and streamlining business processes (Bean 2009). Among service providers, these technologies have been argued to improve the leveraging of existing organizational IT infrastructures, thereby enhancing organizational productivity (Patrick 2005). Adoption of a service-oriented business model has also been shown to augment supply chain performance (Kumar and Dakshinamoorthy 2007), and to produce greater organizational agility and competitiveness (Baskerville et al. 2010). Because of these advantages, and because of the increasing ease with which organizations are able to deploy and consume services electronically, service-oriented electronic markets continue to grow and are likely to flourish for many years to come (Nitto et al. 2009; Rai and Sambamurthy 2006).

Service-oriented electronic markets can hence be expected to provide fertile ground for managerial and economic research into the foreseeable future. At this point, however, knowledge of these markets is relatively sparse (Rai and Sambamurthy 2006), and much remains to be learned, especially in the areas of customer and service provider behavior. With a view toward contributing to the literature in this burgeoning area of study and practice, the current paper considers how different service provider selection methods might affect a customer organization's competitive advantage when operating in a fully-automated, loosely coupled electronic market. Specifically, we consider the question of whether customer organizations participating in such markets can achieve superior outcomes by using predictive inference to select service providers when compared to competing customer organizations who behave opportunistically in the provider selection process.

The balance of this paper is organized as follows: In Sect. 2, characteristics of loosely coupled electronic markets are described, and the paper's primary research question is introduced. In Sect. 3, the simulation methodology by which the paper's principal research question was evaluated is laid out in detail, including descriptions of the simulation process and its associated assumptions, outcome metrics, and data.

Section 4 presents the results of the analyses and discusses those results from the perspective of organizational decision-making and competitive advantage. Finally, Sect. 5 brings the paper to a close by summarizing the results, describing the paper's limitations, and proposing several promising avenues for future research.

2 Loosely coupled electronic markets

Because markets such as that simulated herein rely on web services to enable customer/service provider connections, they have several distinguishing characteristics which are worthy of description. To begin, many of the services which are most amenable for provision and consumption in loosely coupled electronic markets are utilitarian in nature. Examples of such electronic services might include authentication, credit card processing, payroll, procurement, etc. Competition among providers of these services can be strong, as the commoditized nature of those services, in conjunction with the comparative ease of deploying them by way of standard web service technologies, present relatively few barriers to entry. Because a given market may contain many providers of ostensibly equivalent electronic services, customer organizations commonly enjoy the benefit of being able to select from among a bevy of options when choosing a provider for a needed service. Further, because so many of these services are utilitarian in nature, they can, like other utilities such as electricity or telephone service, be purchased on an as-needed or demand basis (Huhns and Singh 2005). What differentiates the purchase of services in electronic markets from the purchase of utilities, however, is the relative ease with which a customer in an electronic market can sever its relationship with one service provider and establish a new relationship with another.

One of the great benefits for customers in loosely coupled electronic markets is that connections with service providers can be established electronically and almost instantaneously on an automated basis (Cardoso and Sheth 2006). Thus, whenever a service provider other than that to which the customer is currently connected is found to be more desirable than the current provider, the customer can easily break its connection with the current provider and connect to the new provider, thereby helping the customer to extract maximum utility from its purchases. To expound upon this notion, consider Fig. 1 below which depicts a simple electronic market containing six organizations.

Each circle in the figure above represents an organization interacting in the market. Services in the market are identified by a unique number, and the services that a given organization needs or provides are indicated by the numbers following the "N" and "P", respectively. If an organization needs to purchase zero services, then it can be classified solely as a provider. Similarly, if an organization provides zero services, then it can be classified solely as a customer. Organizations which both provide and consume services in the market can be classified as being both providers and customers concurrently. To select a provider for a needed service, the customer thus needs only to evaluate the characteristics of the potential providers of that service against its own preferences, and select the provider which is most closely aligned with those preferences.

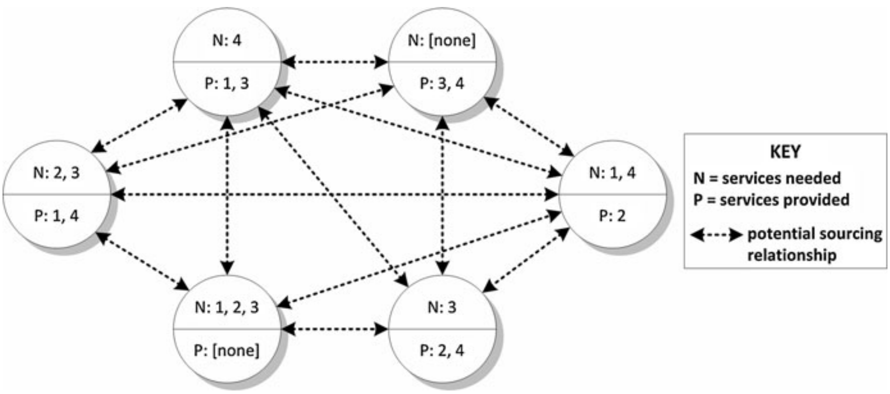

Fig. 1 A simple service-oriented electronic market

The market characteristics described above imply the possibility of many interesting organizational behaviors. Among service providers, for example, the nature of the market allows services to be re-priced on an ongoing basis in real time according to demand. In such a demand-based pricing model, the price of a service is adjusted dynamically in response to changes in demand with a view toward maximizing the service provider's revenue. Among customer organizations, the loosely coupled nature of the market implies the possibility of virtually all connections to service providers being transient. The ability to add and drop connections to service providers as needed thus enables customers to behave opportunistically by seeking out the best possible deal at any given point in time. Customer organizations clearly have several options with respect to the method used to select service providers, however knowledge regarding the long-term effectiveness of these methods is very limited, and has not yet been evaluated in scholarly research. With a view toward contributing to the knowledge in this regard, the balance of this paper will be dedicated to addressing the following research question:

> In a loosely coupled electronic market populated with opportunistic competitors, can a customer organization gain competitive advantage by using predictive inference to select its service providers?

3 Simulation methodology

To answer the research question posed above, a sophisticated custom software application was developed which allowed a moderately-large, service-oriented electronic market to be simulated over an extended period of time. Past research in this area has successfully applied the simulation methodology to study several aspects of loosely coupled electronic markets. Yu et al. 2007 for example, utilized a series of simulations to evaluate the performance of different heuristic algorithms in selecting services when operating under end-to-end quality of service constraints

(Yu et al. 2007). Maximilien and Singh relied on a simulation to demonstrate the tenability of using autonomous agents in conjunction with a quality of service ontology to select web services (Maximilien and Singh 2004). Simulation has also been leveraged to study the plausibility of using automated quality of service brokers in electronic markets (Serhani et al. 2005), to evaluate the performance of composite web services in such markets (Chandrasekaran et al. 2001), and to investigate the selection of web services by customers with dynamically-changing preferences (Lamparter et al. 2007). The current study builds on the existing body of simulation research by considering whether customer organizations with specific preferences can benefit from the application of predictive inference when selecting service providers in loosely coupled electronic markets.

The simulation environment used in this paper was constructed on the basis of several important assumptions which merit discussion. First, the market employed a demand-based pricing model wherein the price charged by a provider for a specific service varied dynamically in real-time according to demand (Pride et al. 2009). Second, all customers operating within the market were served by providers on a first-come-first-served basis, with no preferential pricing or other preferential treatment being afforded to any customer in exchange for their business. Moral and ethical issues aside, this simulated behavior is realistic since the preferential treatment of customers is quite rare in markets wherein organizational interactions take place principally on a short-term, transactional basis (Crane 2007). Finally, the number of organizations interacting in the market was fixed; i.e., no new organizations were allowed to join the market while the simulation was in progress, nor were any existing organizations allowed to leave. Having considered the foundational properties of the simulated market environment, we can now turn our attention to its more detailed characteristics.

3.1 Services and organizations

A total of 20 unique services were made available in the market which could be electronically purchased or sold by organizations interacting therein. Each of these services was assigned a unique identification number, thus allowing an organization's service needs and service provision capabilities to be described numerically. Recalling Fig. 1, a given organization might, for example need to purchase services 2 and 3, and might have services 1 and 4 available for sale to other organizations.

In sum, the simulated market environment contained 100 total organizations. These organizations were allowed to interact with one another in a fully-automated, loosely coupled way, enabling services to be bought and sold in the market by means of electronic connections to business partners that could be established and dissolved on an as-needed basis. Each organization was randomly assigned to need between 0 and 10 services, and to provide between 0 and 10 services. If an organization happened to find itself providing one or more services while also needing to purchase one or more services, care was taken to ensure that no overlap was present between such services; i.e., an organization was not allowed to both need and provide the same service. Further, if an organization happened to be assigned to both need and provide zero services, then the random assignment process was repeated until at least one service was either needed or provided. Within the confines of the electronic market, a given organization could hence find itself acting as a customer, a service provider, or as both a customer and a service provider concurrently. Excepting for the number of organizations and available services, the simulated market environment was thus identical in principle to that depicted in Fig. 1.

3.2 Service providers

Organizations which had services available to sell in the market were randomly assigned a maximum customer capacity between 10 and 100, inclusive. This constraint served to establish a ceiling on the number of customers to which a given provider could concurrently provide services, and was implemented with the intent of accurately representing technological or other resource limitations which might limit service capacity in the real world. Each service provider made available a limited description of itself which could be queried at any time by existing or potential customers. This description contained information about: (1) which services the provider had available for purchase, (2) the current price to process a single transaction, and (3) the quality of service (QoS) which was currently available. It is important to note that past research supports the use of price and QoS as criteria for evaluating potential service providers in loosely coupled electronic markets (Liu et al. 2004). For purposes of mathematical manipulation and interpretation, the values of the price and QoS metrics were constrained to a scale which ranged from 0 to 1, exclusive. Finally, the provider also made available information which indicated whether or not it was currently able to accept new customers. In this regard, service providers were allowed to accept new customers until their maximum capacity had been reached. Aside from the values described immediately above, no other information was made available by service providers to potential or existing customers.

At the outset of a simulation, each service provider was assigned initial values for both price and QoS for the services it provided. Because each service provider began the simulation with zero customers, the QoS for each provider was initially set to 0.999 in order to indicate that the maximum quality of service was currently available.[1] The degree to which the quality of service would degrade according to demand was determined as a linear function of the service provider's maximum customer capacity, such that:

$$\Delta\mathrm{QoS} = \frac{1}{\mathrm{Customers}_{\mathrm{max}}} \tag{1}$$

and:

$$\mathrm{QoS}_{\mathrm{c}} = 0.999 - (\Delta\mathrm{QoS} \times \mathrm{Customers}_{\mathrm{c}}) \tag{2}$$

[1] Three decimals of precision were used in measuring QoS, thus 0.999 was the highest possible value in the range 0–1, exclusive.

where ΔQoS is the decline in quality of service per additional customer, Customers_{max} is the service provider's maximum customer capacity, QoS_c is the current quality of service, and Customers_c is the number of customers currently connected to the service provider. The quality of service available was thus maximal when a service provider had zero customers, and was minimal when a service provider had reached its maximum customer capacity.[2] As additional customers were added, the QoS for providers with a comparatively high capacity would degrade at a slower rate than that realized by service providers with a comparatively low capacity, thus embedding the principles of economies of scale into the simulated market. Each provider was then randomly assigned a value which represented its price-to-QoS ratio. This ratio, whose values ranged from 0 to 1, exclusive, was used throughout the simulation to calculate the service provider's current price in light of demand and capacity. Specifically, the current price charged per transaction was determined as a function of the current QoS and the price-to-QoS ratio, such that:

$$\text{P}_c = (1 - \text{QoS}_c) \times \text{PQR}. \tag{3}$$

where P_c is the current price charged per transaction, QoS_c is the current quality of service available, and PQR is the price-to-QoS ratio. Because the current QoS changed dynamically in response to demand (per Eq. 2), so too did the price per transaction. Thus when demand was high, the price charged by an organization to process a transaction would also be comparatively high, with the reverse being true when demand was low. In this way, the simulated market incorporated the principles of demand-based pricing (Pride et al. 2009), as noted in the introduction to this section.

3.3 Customers

Customers in the electronic market were those organizations which needed to purchase one or more services from other organizations. Each customer was randomly assigned a set of partner selection preferences which represented its own unique views regarding the desirability of potential service providers, and which it used to guide its selection of those providers. Specifically, customer organizations were assigned partner selection preferences along three dimensions: (1) the price charged for a needed service, (2) the quality of service received, and (3) customer loyalty to past service providers. A random value between 0 and 1, exclusive, was assigned to each of these dimensions for each customer organization. Together, these randomly-assigned values represented the preferences of the customer with respect to its evaluation and selection of business partners, with higher values indicating a higher priority for a given dimension. For example, a customer organization whose price, QoS, and loyalty values were 0.85, 0.25, and 0.35 respectively, would heavily prioritize price when selecting business partners. With this system in place, a customer could evaluate each potential service provider along

[2] The maximum value of QoS_c was 0.999, which occurred when $\text{Customers}_c = 0$. The minimum value calculated for QoS_c was -0.001, which occurred when $\text{Customers}_c = \text{Customers}_{max}$. In this latter case a range constraint was applied, and QoS_c was reinitialized to a value of 0.0.

each of the three dimensions (price, QoS, and loyalty) according to its preferences and according to the current state of the market. Together, these considerations constituted a partner selection function whereby a score quantifying the desirability of any potential service provider could be calculated by a customer at any given point in time. Specifically:

$$\text{Score}_{\text{provider}} = \text{Score}_{\text{price}} + \text{Score}_{\text{QoS}} + \text{Score}_{\text{loyalty}} \tag{4}$$

where $\text{Score}_{\text{provider}}$ is the potential provider's overall desirability score, $\text{Score}_{\text{price}}$ is a multiplicative function of the customer's price preference and the potential provider's current price, $\text{Score}_{\text{QoS}}$ is a multiplicative function of the customer's QoS preference and the potential provider's current QoS, and $\text{Score}_{\text{loyalty}}$ is a multiplicative function of the customer's loyalty preference and the potential provider's current loyalty score. With respect to the latter, customer loyalty has been identified as a critical component in business-to-business sales and procurement relationships (Rauyruen and Miller 2007), and recent research has concluded that loyalty in e-service settings is driven by customer satisfaction, which itself is driven by trust (Gummerus et al. 2004). Customer satisfaction and trust, however, are tied to the emotions of human decision makers, which are not present in fully-automated electronic markets such as that simulated herein. This paper instead conceptualizes customer loyalty more in terms of familiarity with a service provider than in terms of an emotional connection between the customer and the provider. Each customer therefore maintained a record of the providers to which it had previously connected. Based upon these values, a customer could compute a loyalty score for each potential provider as a function of how many times the customer had purchased services from the provider in the past. For each potential provider, the loyalty score could range from 0 to 1, exclusive, and the simulation began with all potential providers being assigned a minimal loyalty score of 0.001. Whenever a customer chose to connect to a particular provider for the purpose of purchasing services, that particular provider's loyalty score would be incremented by 0.333. If after one of these operations the loyalty score exceeded the maximum allowed value, a range constraint was applied, and it was reinitialized to a value of 0.999. Under this model, maximum loyalty was therefore achieved after a customer had made three successive connections to the same provider for the purchase of the same service. Whenever a customer selected a provider for the purchase of a specific service, the loyalty scores for all other past providers of that same service would decay. Specifically, the loyalty scores of all past providers of a particular service were reduced by 33.3% whenever one of their competitors was chosen to provide that service instead of themselves.

Possible values of the overall provider score for a potential service provider (*q.v.*, Eq. 4) thus ranged from 0 to 3 (exclusive), with higher values indicating that a provider was more desirable to the customer. In this regard, a potential service provider could be said to be more desirable as its characteristics become more closely aligned with the preferences of the customer. These provider scores were used by customers in the selection of providers insofar as the default behavior of customer organizations during a simulation was to first compute a provider score for each available provider of the needed service, and then opportunistically choose the provider with the highest score. This approach would ostensibly lead a customer to always select the provider which was most closely aligned with its preferences at that moment in time. This was not always possible, however, because providers were not allowed to accept any additional customers if they were already operating at maximum capacity. In such situations, the next best available provider would be selected (i.e., the provider with the next highest provider score).

3.4 The organization of interest (OoI)

As a means of evaluating this paper's principal research question—that is, whether or not a customer organization can gain competitive advantage by using predictive inference to select business partners in a loosely coupled electronic market—one of the customer organizations was selected at random to behave differently than its peers. This organization, which we shall refer to as the Organization of Interest (OoI), was allowed to use simple linear regression to predict what might happen over time if it were to connect to a potential service provider. That is, rather than selecting a service provider based on its provider score at the time of connection, the OoI would select service providers on the basis of the expected value of a chosen optimization metric over the life of the customer/service provider purchasing relationship. The outcomes of the predictive approach used by the OoI, as quantified by several metrics which shall be described shortly, could then be compared against the outcomes achieved by the other customer organizations whose behavior was more opportunistic. In this way, it was possible to assess the efficacy and usefulness of the predictive inference approach to service provider selection as compared to that of opportunistic service provider selection in electronic markets.

3.5 The simulation process

In addition to the initial configuration settings described above, it is also important to consider the overall process which guided each simulation. To begin, the customer/service provider purchasing relationship in the simulation was centered on the completion of what we shall call *jobs*, each of which consisted of a number of transactions which needed to be processed as a single work unit. The duration of each job could thus be described in terms of the number of transactions that it contained, and in the simulation each job was randomly assigned to consist of between 1 and 100 transactions, inclusive. Because of the nature of the demand-based pricing model, the price charged by a provider for the processing of a single transaction could change dynamically from transaction to transaction according to real-time demand, as could the quality of the service provided. After a job had been completed, the customer was assigned a new job for the given service, with the number of transactions in the job again being a randomly assigned integer between 1 and 100. At this point, the customer would reevaluate all of the potential providers for the current service with a view toward considering any changes that had taken place in the market while the previous job was being processed. Depending upon the newly-calculated provider scores, either the current provider was retained for the next job, or a better provider (in terms of alignment with customer preferences) was selected.

Each simulation was carried out in three distinct phases: (1) the stabilization phase, (2) the training phase, and (3) the testing phase. In the case of the first of these phases, a period of 1,000 transaction processing cycles was carried out for purposes of bringing the market to a point of stabilization. This was necessary because each simulation began without any existing connections between organizations, and because all customers entered the market simultaneously at the outset of the simulation. Naturally, one would expect abnormal market behavior under such unusual conditions, and so the stabilization phase was included as a means of ensuring that the market had normalized by the time that data gathering began. The second phase in the simulation was the training phase, which lasted for 5,000 transaction processing cycles. During this phase, all customers in the market—including the OoI—behaved according to the default behavior described previously (i.e., they behaved opportunistically in the selection of their service providers). The OoI, however, was allowed to gather data about potential providers during this time. Specifically, the OoI gathered pricing and QoS data from each potential provider for each transaction processing cycle, thus allowing for the construction of a large dataset which quantified the changes in those metrics for each provider over time. When the training phase was complete, the simulation entered its third and final stage, the testing phase. During this phase, which lasted for an additional 5,000 transaction processing cycles, all customer organizations except the OoI behaved according to the opportunistic provider selection method described previously. The OoI, however, was allowed to use simple linear regression in the selection of its service providers. Relying upon the dataset which had been accumulated during the training phase, the OoI would examine the properties of a potential provider (i.e., the provider's current QoS level and price charged per transaction), and would then identify any cases in the dataset which were similar to the currently-observed values. Specifically, cases identified as similar were those whose metric values were within 5% of the currently-observed values. Beginning with each of these identified cases, the $n - 1$ observations immediately following the identified case were extracted and added to a temporary dataset, where n was the number of transactions in the customer's current job. Thus, if a customer had 30 transactions to process as part of its current job, a longitudinal series of 30 observations would be extracted from the overall dataset for each situation in which the properties of the provider as recorded in the dataset closely approximated those which were being currently observed for that provider. This subset of data was then used by the OoI to generate a simple linear regression equation which would predict what might happen over the course of the job if the customer were to connect to this provider. Using this equation, the OoI could then compute an expected average value of the metric of interest over the life of the current job. Unlike the other customers in the market, the OoI would then select the service provider which was expected to be optimal over the life of the customer/supplier purchasing relationship, rather than selecting the service provider which was deemed optimal at that moment in time. This process is illustrated in Fig. 2 below.

The figure above depicts two potential provider options for a service which might be considered by a customer in the market. The diamonds in the figure represent the price charged by the provider over the life of a job containing 20 transactions, and

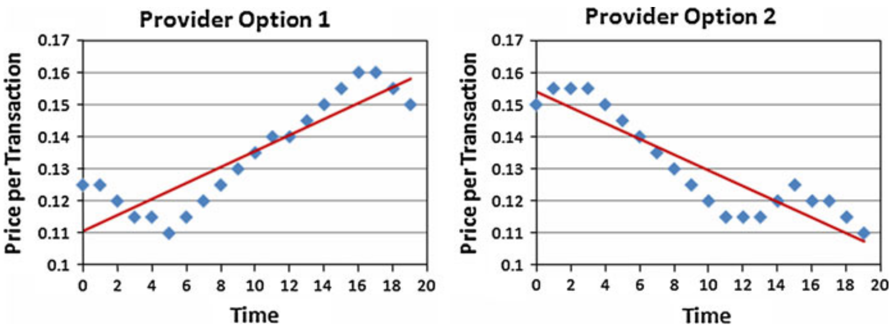

Fig. 2 Comparison of service provider selection methods

the line represents the least-squares regression equation for those transactions. If an opportunistic customer was optimizing for price, it would select Provider Option 1 whose current price (Time 0) was 0.125 units versus a current price of 0.15 units for Provider Option 2. By considering the expected average price over the life of the job, however, the OoI would select Provider Option 2, because it would be expected to produce superior outcomes in the long run.

A final consideration which merits mentioning here is the dissatisfaction rate which was built into the simulated electronic market. It has been established that loyalty can decrease when a customer is dissatisfied with their purchase (Cooil et al. 2007), and to model this behavior in the simulation a random 1% dissatisfaction rate was introduced into the marketplace. When a customer was dissatisfied with the service it had received—a situation which occurred for 1% of the jobs processed during the simulation—the customer would penalize the service provider by refusing to purchase any services from that provider during the subsequent job processing cycle, even if the provider was determined to be optimal at that moment in time. In this way, customer loyalty would degrade with respect to the offending service provider in response to the customer having been dissatisfied with the service received.

3.6 Outcome metrics

A total of four outcome metrics were used in the evaluation of this paper's principal research question: (1) price, (2) QoS, (3) price-to-QoS ratio, and (4) provider score. During the simulation, the OoI was assigned the task of attempting to optimize for one of these four metrics. First, the price metric was operationalized as the average price paid per transaction by a customer in the service-oriented electronic market over the life of the testing phase of the simulation. Customers that optimize for price in such a market are principally interested in purchasing needed services for the lowest average cost possible, regardless of the quality of service received. By way of example, we might draw a parallel between such customers and travelers who select an airline flight based solely on the cost of the ticket. Similarly, the QoS metric was operationalized as the average quality of service received by a customer during the testing phase of the simulation. Customers that optimize for quality of service prefer to obtain services at the highest available QoS level, regardless of the costs incurred in the purchase of those services. Continuing our previous example, we might draw a parallel between such customers and travelers who prefer to always fly first-class, regardless of the cost of doing so. The third outcome metric was the price-to-QoS ratio, which was operationalized as the average price paid per unit of QoS by a customer throughout the testing phase of the simulation. This ratio would be of interest to customers which do not have specific QoS or pricing objectives, but instead are seeking to obtain the best quality of service possible at the lowest cost. In our continuing example, we might draw a parallel between these customers and travelers who would consider purchasing a business-class ticket if the perceived additional benefits of doing so would outweigh the additional costs.

For each of these first three metrics (i.e., price, QoS, and the price-to-QoS ratio), it was necessary to consider two separate scenarios. In the first of these scenarios, the performance of the OoI, which was optimizing for one of the outcome metrics, would be compared against the average performance realized for that metric by all other customers operating in the market. In this way it would be possible to quantify any benefits of the predictive inference method in the context of the whole market, which, as described previously, was populated with competing customers, each of which was behaving according to its own randomly-assigned preferences. In the second of these scenarios, the performance of the OoI would be evaluated against that of another randomly-selected customer organization which, although behaving opportunistically, had been assigned the same optimization objectives as the OoI. In this way, a true comparison could be made between the overall performance of the predictive inference method versus that of the opportunistic method of selecting providers.

The final outcome metric which contributed to the evaluation of this paper's primary research question was that of the provider score. In this scenario, the OoI was randomly assigned preferences for price, QoS, and loyalty in exactly the same fashion as any other customer in the market. When all customers, including the OoI, were using a randomly-assigned customer preference function to compute a provider score for each potential provider, the test of the efficacy of the predictive inference approach can be said to be the most robust, because aside from the provider selection method (i.e., predictive inference vs. opportunistic selection), the customers all behaved identically. Put another way, the best test of the predictive inference approach to service provider selection was to allow all customers in the market to optimize for a randomly-assigned set of preferences, and then observe any differences in performance which were directly attributable to the method of service provider selection used.

4 Results and discussion

A total of 60 complete simulation cycles were carried out for each of the first three outcome metrics described in the previous section (i.e., price, QoS, and price-to-QoS ratio), with an additional 30 simulation cycles carried out for the evaluation of the randomly-assigned customer preference function (provider score). Together, these 210 simulation cycles produced a vast quantity of output, yielding a combined dataset which contained more than 80 million unique transactions. Of these, approximately 74.1 million transactions were associated with customer organizations which behaved opportunistically by selecting service providers based upon consideration of whether those providers were optimal at the time of selection. The remaining approximately 6.1 million transactions were associated with customer organizations which utilized predictive inference with a view toward optimizing for expected levels of a target outcome metric over the life of the customer/service provider purchasing relationship. The effectiveness of this latter method relative to that of the former was ascertained by means of a series of one-way analyses of variance (ANOVAs) wherein the output derived from the opportunistic service provider selection method was compared to that of the service provider selection method which was based on predictive inference (Keller 2008). Details for the statistical analyses comparing the two service provider selection methods for each of the four outcome metrics are provided in the subsections which follow immediately hereafter.

4.1 Price

Customer organizations which optimize for price in a service-oriented electronic market are principally interested in purchasing needed services for the lowest average cost possible, regardless of the quality of service received. The results obtained from the simulations in which the OoI had been configured to utilize predictive inference with a view toward optimizing for price are provided in Tables 1 and 2 below. These values are contrasted with the average price paid per transaction for customer organizations which behaved opportunistically in the selection of their service providers.

As shown in Table 1, when using predictive inference to optimize for the lowest price possible, the OoI was able to achieve an average price per transaction which was substantially less than the average price per transaction paid by the other customer organizations interacting in the market (0.005 vs. 0.411). Further, this difference in average price between the two provider selection methodologies was found to be highly significant ($p < 0.001$), as indicated by the ANOVA results reported in Table 2. For loosely coupled electronic markets in which the majority of customer organizations behave opportunistically according to their own set of preferences, these results imply that an organization whose preference is to obtain

Table 1 Comparative performance by average transaction price

Selection method	Average price per transaction	
	Predictive inference	Opportunistic
N	750,000	21,900,000
Mean	0.005	0.411
Standard deviation	0.016	0.397

Table 2 ANOVA for difference by average transaction price

	Average price per transaction				
	Sums of squares	*df*	Variance	*F*	*p*
Between	119,533.391	1	119,533.391	784,346.822	<0.001
Within	3,451,828.942	22,649,998	0.152		
Total	3,571,362.333	22,649,999			

needed services at the lowest average cost possible can achieve success using the predictive inference approach to selecting service providers.

The sizeable degree of difference between the average price paid by the OoI and that paid by the other customers can be partially attributed to the fact that all of the other customers in the market were selecting business partners according to a randomly-assigned, multidimensional preference function. The majority of these customer organizations were thus not optimizing exclusively for price, but were also taking other factors such as QoS and loyalty into account when selecting their business partners—factors which may have led them to pay more for the same service. The extent to which the utilization of predictive inference contributed to the observed difference could not thus be determined from the results reported above alone. For this reason, an additional 30 simulations were conducted in which one of the opportunistic customer organizations was assigned preferences that instructed it to optimize exclusively for price. By comparing the OoI to an opportunistic customer which was optimizing for price, a more realistic assessment of the comparative advantage of the predictive inference method could be proffered. These results are presented in Tables 3 and 4 below.

As shown in Table 3, when directly comparing the performance of two customer organizations—both of which had optimized their preferences with a view toward purchasing services at the lowest price possible—the predictive inference method was observed to acquire services at a price which was substantially lower than the average price paid per transaction by the customers which acted opportunistically. As this difference was observed to be statistically-significant ($p < 0.001$, per Table 4), it can be concluded that when optimizing for price, customer organizations which utilize predictive inference when selecting their service providers will enjoy a distinct advantage over customers which act opportunistically.

Table 3 Comparative methodological performance among customers optimizing for price

Selection method	Optimizing for price	
	Predictive inference	Opportunistic
N	750,000	850,000
Mean	0.005	0.075
Standard deviation	0.028	0.073

Table 4 ANOVA for difference in methodological performance among customers optimizing for price

	Optimizing for price				
	Sums of squares	*df*	Variance	*F*	*p*
Between	1,952.344	1	1,952.344	610,387.546	<0.001
Within	5,117.644	1,599,998	0.003		
Total	7,069.988	1,599,999			

4.2 Quality of service (QoS)

By optimizing for quality of service, a customer organization is signaling its desire to obtain services at the highest available QoS level, regardless of the costs incurred in the purchase of those services. As with price, we begin by first presenting the results of the simulations in which the OoI used predictive inference to select service providers with a view toward obtaining the highest average QoS possible, while the balance of the customer organizations behaved opportunistically according to their own randomly-assigned preference functions.

As shown in Table 5, the OoI was able to substantially outperform the overall average of the other customers interacting in the electronic market with respect to its ability to secure the highest possible quality of service for needed services. Specifically, the OoI was able to procure needed services at an average QoS level of 0.932, a value some 44% greater than the overall average QoS level of 0.645 which was obtained by customers acting opportunistically. The values reported in Table 6 indicate that this difference in performance was highly statistically-significant ($p < 0.001$), thereby lending support to the conclusion that an organization which uses predictive inference to select service providers with a view toward optimizing for the quality of services received can procure those services at a QoS level which is markedly higher than the average level obtained by competing customer organizations which behave opportunistically according to their own preferences.

While these results are interesting, they do not speak to whether a customer organization using predictive inference in the selection of its business partners could outperform a competing opportunistic customer when both organizations are attempting to maximize the average QoS received. For this reason another series of 30 simulations was conducted in which one of the opportunistic customer organizations was intentionally assigned the same task as the OoI; namely, to attempt to maximize the average QoS received. The results comparing these two

Table 5 Comparative performance by average quality of service received

Selection method	Average quality of service received	
	Predictive inference	Opportunistic
N	1,000,000	26,200,000
Mean	0.932	0.645
Standard deviation	0.040	0.376

Table 6 ANOVA for difference by average quality of service received

	Average quality of service received				
	Sums of squares	*df*	Variance	*F*	*p*
Between	79,340.728	1	79,340.728	582,372.060	<0.001
Within	3,705,651.057	27,199,998	0.136		
Total	3,784,991.785	27,199,999			

methodological approaches to service provider selection in the context of quality of service are presented in Tables 7 and 8 below.

As compared to the opportunistic method, the results shown in Table 7 highlight the degree of superiority of the predictive inference method of service provider selection among competing customer organizations seeking to maximize the quality of service received in loosely coupled electronic markets. Per Table 8, the observed difference of approximately 8% was found to be highly statistically-significant. This observation provides support to the notion that the predictive inference method can effect substantially more desirable outcomes than can be realized by opportunistic service provider selection behavior among competing customers in a dynamic, loosely coupled market environment.

4.3 Price-to-QoS ratio

The price to quality of service ratio reflects the per-unit QoS cost which would be incurred by a customer organization for the purchase of a given service from a provider. This outcome metric would be of interest to customer organizations which do not have specific QoS or pricing objectives, but instead are seeking to obtain the

Table 7 Comparative methodological performance among customers optimizing for QoS

Selection method	Optimizing for quality of service	
	Predictive inference	Opportunistic
N	1,000,000	2,000,000
Mean	0.944	0.874
Standard deviation	0.066	0.042

Table 8 ANOVA for difference in methodological performance among customers optimizing for QoS

	Optimizing for quality of service				
	Sums of squares	*df*	Variance	*F*	*p*
Between	3,266.667	1	3,266.667	1,243,023.982	<0.001
Within	7,883.994	2,999,998	0.003		
Total	11,150.661	2,999,999			

best quality of service possible at the lowest cost. With this objective in mind, we shall first contrast the performance of the predictive inference method used by the OoI with that of the other customer organizations, each of which was acting opportunistically in the market according to its own randomly-assigned preference function. These results are provided in Tables 9 and 10 below.

Compared to the average price-to-QoS ratio achieved by customer organizations acting opportunistically in the electronic market, the results reported in Table 9 indicate that the OoI, which was using predictive inference in the selection of its business partners, was able to achieve a much lower price-to-QoS ratio. It can therefore be concluded that customer organizations that are interested in optimizing for the price-to-QoS ratio—and which utilize predictive inference when selecting service providers—can expect to pay markedly less for services on a per-unit of QoS basis than the average paid by other customers in the market which are acting opportunistically. To investigate the efficacy of the predictive inference method in more detail with respect to its usefulness in obtaining a comparatively low price-to-QoS ratio, the performance of the OoI was evaluated against that of one of the opportunistic customer organizations in a series of 30 additional simulations. For these simulations, the opportunistic organization was intentionally assigned service provider selection preferences which would optimize for the price-to-QoS ratio, thus allowing the methodological performance of the predictive inference method to be directly compared against that of the opportunistic method in a valid way. The results obtained from these simulations are reported in Tables 11 and 12 below.

As with the results obtained for the evaluation of the price and QoS outcome metrics, the predictive inference method was once again observed to produce superior outcomes in terms of the average price-to-QoS ratio achieved when both the OoI and the opportunistic customer organization were optimizing for that ratio. Specifically, the values reported in Table 11 indicate that the average price-to-QoS

Table 9 Comparative performance by average price-to-QoS ratio

Selection method	Average price-to-QoS ratio	
	Predictive inference	Opportunistic
N	600,000	15,050,000
Mean	0.007	101.324
Standard deviation	0.126	285.288

Table 10 ANOVA for difference by average price-to-QoS ratio

	Average price-to-QoS ratio				
	Sums of squares	*df*	Variance	*F*	*p*
Between	5,922,949,804.196	1	5,922,949,804.196	75,674.377	<0.001
Within	1,224,908,034,443.540	15,649,998	78,268.894		
Total	1,230,830,984,247.730	15,649,999			

Table 11 Comparative methodological performance among customers optimizing for price-to-QoS ratio

Selection method	Optimizing for price-to-QoS ratio	
	Predictive inference	Opportunistic
N	800,000	600,000
Mean	0.009	0.020
Standard deviation	0.071	0.014

Table 12 ANOVA for difference in methodological performance among customers optimizing for price-to-QoS ratio

	Optimizing for price-to-QoS ratio				
	Sums of squares	*df*	Variance	*F*	*p*
Between	41.486	1	41.486	13,993.830	<0.001
Within	4,150.395	1,399,998	0.003		
Total	4,191.881	1,399,999			

ratio achieved by the OoI using predictive inference was more than 50% lower (0.009 vs. 0.020) than that achieved by a competing customer employing an opportunistic service provider selection method. Given that this difference was observed to be statistically-significant at the 0.001 level (per Table 12), it may be concluded that the predictive inference method of selecting service providers produces statistically-superior outcomes in terms of the average price-to-QoS ratio achieved when compared to the performance of the opportunistic method.

4.4 Provider score

When compared to the three outcome metrics previously described in this section, a randomly-assigned customer preference function is perhaps the most robust test of the efficacy of the predictive inference approach to selecting service providers in a dynamic, loosely coupled electronic market. The reason for this lies in the approach used in the simulations to assign purchasing preferences to customer organizations acting opportunistically in the market. For these organizations, the extent to which a given provider was considered desirable was based upon its provider score, which itself incorporated randomly-assigned customer preferences for price, quality of service, and partner loyalty (*q.v.*, Eq. 4). Among all of the customer organizations participating in the market, we would therefore expect that on average only a small proportion had been randomly assigned preferences which would wholly prioritize price, QoS, or the price-to-QoS ratio over all of the other constituents of the preference function. Thus, as is the case among real-world organizations, most of the customers participating in the market would not select partners based solely upon price, QoS, or partner loyalty, but rather on some combination of those considerations. For this reason, perhaps the best test of the predictive inference approach was to compare the outcomes of customers using that approach with those

Table 13 Comparative performance by provider score

Selection method	Provider score	
	Predictive inference	Opportunistic
N	1,150,000	7,450,000
Mean	1.225	1.028
Standard deviation	0.311	0.360

Table 14 ANOVA for difference by provider score

	Provider score				
	Sums of squares	*df*	Variance	*F*	*p*
Between	38,662.338	1	38,662.338	308,796.250	<0.001
Within	1,076,748.924	8,599,998	0.125		
Total	1,115,411.262	8,599,999			

which selected service providers opportunistically when all of the customers, regardless of their service provider selection method, were using randomly-assigned preferences. In this scenario, a customer would use Eq. 4 to calculate a provider score for each available provider of a given service. The larger the value of this provider score, the more closely the provider could be said to match the customer's overall preferences. The results of the 30 simulations in which all customer organizations, including the OoI, used a randomly-assigned preference function, are provided in Tables 13 and 14 below.

As shown in Table 13, customer organizations which used predictive inference to select service providers were able to achieve an average provider score of 1.225 over the life of the purchasing relationship, as opposed to an average provider score of 1.028 realized by customer organizations which behaved opportunistically. When optimizing for a randomly-assigned preference function, the predictive inference method was thus able to produce outcomes which were approximately 19% better than those produced by the opportunistic method—a difference which was highly significant, as indicated by the statistical results reported in Table 14. We may therefore conclude that, even under the most robust and equitable testing conditions, customer organizations which use predictive inference to select service providers in a loosely coupled electronic market can be expected to achieve outcomes which are markedly and statistically superior to those achieved by competing customer organizations which select their service providers opportunistically.

5 Summary and concluding remarks

This paper considered the question of whether customer organizations participating in a fully-automated, loosely coupled electronic market could achieve superior outcomes by using predictive inference to select service providers when compared to competing customer organizations who behave opportunistically in the selection of such providers. Rather than selecting a service provider based upon consideration of whether it is optimal at the time of selection, it was shown that superior outcomes could be achieved by instead selecting service providers based upon the expected levels of a desired outcome metric over the life of the transient customer/service provider purchasing relationship. Specifically, it was shown that a customer organization could achieve superior outcomes by utilizing simple linear regression-based prediction in the selection of service providers, regardless of whether the customer was interested in optimizing for price, quality of service, the price to quality of service ratio, or even a randomly-assigned, multidimensional customer preference function.

These results carry several important implications for both practitioners and researchers. With respect to the former, the results imply not only that predictive inference is a viable method of selecting business partners in loosely coupled electronic markets, but also that the use of such techniques can effect a sizeable and statistically-significant advantage for organizations who are inclined to adopt a prediction-based approach. Managers wishing to enter the rapidly-growing market for electronic services should therefore consider the potential benefits that a predictive approach might bring to their efforts. Managers already operating in these markets should similarly consider incorporating predictive capabilities into their existing systems, since it is early adopters who will likely extract the greatest benefits from this approach. With respect to researchers, the paper's findings also point to the viability of using large-scale simulations to study electronic markets. Indeed, there may be many emergent phenomena in loosely coupled electronic markets for which a simulation methodology would be particularly well-suited.

As with all research efforts, there are some limitations to this work which must be acknowledged. First, the results reported here were derived from a simulated market environment. While every effort was taken to ensure that the simulation was as sophisticated and representative of reality as possible, it must ultimately be left to the reader to decide whether those efforts sufficiently captured the subtleties of such markets and realistically reproduced the behaviors of organizations interacting therein. Second, this paper's results were obtained from a marketplace in which only a single customer organization utilized predictive inference to select providers of services that it needed to purchase. Although an empirical question, it seems reasonable to expect that the scope of the benefits realized by using prediction to select service providers would be inversely related to the proportion of customer organizations interacting in the market which have chosen to adopt a predictive approach. Finally, this research relied only upon simple linear regression as a predictive methodology for inferring partnership outcomes. While the magnitude of the benefits obtained by using this simple method are indeed encouraging, it may well be possible to achieve even better outcomes through the use of more advanced predictive methodologies such as polynomial regression, multilevel change modeling, or pattern analysis. Together, all of these limitations represent tenable opportunities for future research in the area of service provider selection in dynamic, loosely coupled electronic markets.

The merger of service-oriented business architectures with web service technologies has engendered myriad new and potentially revolutionary possibilities for modern business. Given the results reported in this paper, one might readily imagine a new business model, predicated on the monitoring and collection of metrics of interest within a service-oriented electronic market, in which a broker could identify and evaluate service providers and sell the resulting data to interested customers for a fee. Further, one might imagine the powerful possibility of operating a business which is nothing more than a collection of services purchased from other service providers and assembled in such a way as to add value. The era is now upon us when, like assembling the pieces of a jigsaw puzzle, a visionary entrepreneur could leverage methods such as those described in this paper to identify and interconnect ostensibly independent electronic services with a view toward operating a business characterized by both substantial profits and very low overhead. There are, of course, many other as yet unforeseen business models which will undoubtedly emerge in response to this rapidly developing paradigm, and these portend a panoply of future avenues for fruitful research in the area of loosely coupled electronic markets. May both scientists and managers alike move prudently and deliberately as we step together across the threshold into this brave new era of business.

References

Alexander M, Korine H (2008) When you shouldn't go global. Harv Bus Rev 71–77

Baskerville RL, Cavallari M, Hjort-Madsen K, Pries-Heje J, Sorrentino M, Virili F (2010) The strategic value of SOA: a comparative case study in the banking sector. Int J Infor Technol Manag 9(1):30–53

Bean J (2009) SOA and web services interface design: principles, techniques, and standards. Morgan Kaufmann, Burlington

Cantara M, Adams C, Clark LM (2007) Worldwide IT professional services market size and forecast for SOA, web services and web 2.0, 2005–2010. Garner, Inc., Stamford

Cardoso J, Sheth AP (2006) Semantic web services. Processes and applications. Springer, New York

Castro-Leon E, He J, Chang M (2007) Scaling down SOA to small businesses. In: Paper presented at the IEEE international conference on service-oriented computing and applications. Newport Beach, CA

Chandrasekaran S, Miller JA, Silver GS, Arpinar B, Sheth AP (2001) Performance analysis and simulation of composite web services. Electron Mark 13(2):120–132

Cooil B, Keiningham TL, Aksoy L, Hsu M (2007) A longitudinal analysis of customer satisfaction and share of wallet: investigating the moderating effect of customer characteristics. J Mark 71(1):67–83

Crane A (2007) Business ethics. Oxford University Press, Oxford

Gummerus J, Liljander V, Pura M, Av Riel (2004) Customer loyalty to content-based web sites: the case of an online health-care service. J Serv Mark 18(3):175–186

Huhns MN, Singh MP (2005) Service-oriented computing: key concepts and principles. IEEE Int Comput 9(1):75–81

Keller G (2008) Statistics for management and economics, 8th edn. South-Western Cengage, Mason

Kumar S, Dakshinamoorthy V (2007) Does SOA improve the supply chain? An empirical analysis of the impact of SOA adoption on electronic supply chain performance. In: Paper presented at the Hawaii international conference on system sciences (HICSS). Waikoloa, HI

Lamparter S, Ankolekar A, Studer R, Grimm S (2007) Preference-based selection of highly configurable web services. In: 16th international World Wide Web conference. Banff, Alberta, Canada

Lawler JP, Howell-Barber H (2007) Service-oriented architecture: SOA strategy, methodology, and technology. Auerbach Publications, New York

Liang T-P, Tanniru M (2006–2007) Customer-centric information systems. J Manag Infor Syst 23(3):9–15

Liu Y, Ngu AH, Zeng LZ (2004) QoS computation and policing in dynamic web service selection. In: Proceedings of the 13th international World Wide Web conference, Manhattan, NY

Maximilien EM, Singh MP (2004) A framework and ontology for dynamic web services selection. IEEE Internet Comput 8(5):84–93

Newcomer E, Lomow G (2004) Understanding SOA with web services. Addison-Wesley, Reading

Nitto ED, Sassen A-M, Traverso P, Zwegers A (2009) At your service: service-oriented computing from an EU perspective. MIT Press, Cambridge

Papazoglou MP, Georgakopoulos D (2003) Service-oriented computing. Commun ACM 44(10):25–28

Patrick P (2005) Impact of SOA on enterprise information architectures. In: Paper presented at the ACM SIGMOD international conference on management of data. Baltimore, MD

Pride WM, Hughes RJ, Kapoor JR (2009) Business. South-Western Cengage, Mason

Rai A, Sambamurthy V (2006) The growth of interest in services management: opportunities for information systems scholars. Infor Syst Res 17(4):327–331

Rauyruen P, Miller KE (2007) Relationship quality as a predictor of B2B customer loyalty. J Bus Res 60(1):21–31

Schulte RW, Natis YV (1996) Service oriented architectures, part 1. Gartner

Serhani MA, Dssouli R, Hafid A, Sahraoui H (2005) A QoS broker based architecture for efficient web services selection. In: IEEE international conference on web services. Orlando, Florida

Smith DB (2008) Is SOA real? In: Paper presented at the seventh IEEE international conference on composition-based software systems (ICCBSS). Madrid, Spain

Yoon T, Carter P (2007) Investigating the antecedents and benefits of soa implementation: a multi-case study approach. In: Paper presented at the Americas conference on information systems. Keystone

Yu T, Zhang Y, Lin K-J (2007) Efficient algorithms for web services selection with end-to-end QoS constraints. ACM Trans Web 1(1)