Association for Information Systems
AIS Electronic Library (AISeL)

Wirtschaftsinformatik Proceedings 2013 | Wirtschaftsinformatik

2013

Individualized Quality-Differentiated Services: A Market Model and Comparison of Negotiation Mechanisms

Steffen Haak
Research Center for Information Technology (FZI), Karlsruhe, Germany, haak@fzi.de

Henner Gimpel
Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany, henner.gimpel@kit.edu

Follow this and additional works at: http://aisel.aisnet.org/wi2013

Recommended Citation
Haak, Steffen and Gimpel, Henner, "Individualized Quality-Differentiated Services: A Market Model and Comparison of Negotiation Mechanisms" (2013). *Wirtschaftsinformatik Proceedings 2013*. 9.
http://aisel.aisnet.org/wi2013/9

This material is brought to you by the Wirtschaftsinformatik at AIS Electronic Library (AISeL). It has been accepted for inclusion in Wirtschaftsinformatik Proceedings 2013 by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

Individualized Quality-Differentiated Services: A Market Model and Comparison of Negotiation Mechanisms

Steffen Haak[1] and Henner Gimpel[2]

[1] Research Center for Information Technology (FZI), Karlsruhe, Germany
haak@fzi.de
[2] Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany
henner.gimpel@kit.edu

Abstract. Individualized services are common in distributed computing systems: Consumers demand custom solutions and service providers tailor their offering. In negotiating a service contract, providers and consumers frequently fail in identifying the optimal combination of non-functional properties and price. A key reason is that negotiators simultaneously try to create and claim value leading to strategic bidding and inefficient outcomes.
We present a theoretical comparison of three negotiation mechanisms in the scenario of self-interested, rational agents bilaterally negotiating over service quality and price. Two mechanisms are stylized representations of mechanisms commonly used in theory and practice, one is newly introduced. The mechanisms are characterized in terms of truthfulness of agents, efficiency of outcomes, and distribution of welfare.
The analysis is an extension to the field of bilateral multi-attribute negotiations, relevant for researchers and practitioners designing markets for individualized IT services and for human and computer agents acting in such markets.

Keywords: Consumerization, Service Quality, Service Management, Negotiation Strategy, Automated Negotiation

1 Introduction

Today, consumerization of products and services is well-set in almost any domain. For services, individual and close interaction between consumer and provider happens almost by definition. With technological progress and rising customer expectations, off-the-shelf offerings to fit the greatest common denominator of consumer demands go out of date in favor of individualized solutions. In this trend, quality marks the center of attention besides price and functional requirements. This is especially the case for IT-based services, with Cloud services being in the focus of current attention.

11th International Conference on Wirtschaftsinformatik,
27th February – 01st March 2013, Leipzig, Germany

1.1 Individualized Quality-Differentiated Services

There are different types of Cloud services available, forming a continuum from very small atomic services that serve a very distinct purpose to very large services, built manually and on demand for important consumers. Atomic services are highly standardized and typically address the volume business segment, trying to sell the same service to as many customers as possible while keeping transaction costs at a minimum. An example for such a very small service is *WhatsApp*, a messaging service for smartphone users. It is offered in a single quality for a fixed price. Very large services on the other hand are situated in the value business segment, trying to meet the very individual needs of substantial consumers. Examples include IT outsourcing by *IBM*, *SAP* and others.

In this work, we focus on medium sized service offers in the middle of the aforementioned continuum. With *Dropbox* and *Amazon Web Services* as examples, these services are characterized by a set of different service levels. They try to compromise value and volume business by advertising different versions and qualities of the same basic functionality. With increasing number of versions, services are provided in a mass customization manner, which becomes very profitable if transaction costs can be effectively reduced by automation. *1&1 Dynamic Cloud Server*, for example, offers almost a continuum of different service configurations for consumers to choose from through a simple Web interface.[1] In this example, the configurations are rather close to technical specifications, the extension to Service Level Agreements (SLAs) is, however, straightforward.

In the trend of an increasing demand for short-lived and ad hoc outsourcing, it becomes even more important to reduce transaction costs through automating the agreement process for customized services. The agreement on quality attributes of a customized service is neither black nor white: The consumer may have an optimal service offer in mind which would perfectly fulfill her needs along multiple quality dimensions, paired with a maximum willingness to pay. Offerings that yield slightly lower quality may still be good enough, yet come along with a decreased willingness to pay. In the following we assume that the consumer's preferences over quality attributes can be expressed in a scoring rule or function [1-3].

1.2 Negotiation of Service Agreements

It has long been postulated that the wide-spread, fine-grained use of IT services in cross-organizational environments requires low transaction costs in negotiating service agreements [4]. This implies multiple challenges: Negotiating parties need to have a common understanding of the negotiation object, e.g., via a domain ontology. Parties need a negotiation mechanism to follow in their negotiation and, finally, each party needs a decision making model to automate its internal decision making. See [4-6] for early approaches to these challenges. Over the last years, economic mecha-

[1] http://hosting.1und1.de/cloud-server-config (accessed on Nov. 26, 2012)

nisms for service negotiation and management became more prominent; see [7] for a recent survey.

In this paper, we address the challenge of selecting a negotiation mechanism, by proposing and comparing three mechanisms. Multiple mechanisms need to be designed and deployed as consumers and providers have different goals, objectives, strategies, and requirements [7]. A multi-dimensional negotiation between consumer and provider typically includes an integrative and a distributive element. The integrative element shall identify the optimal service level that maximizes the difference from consumer utility and provider costs. In other words and when expressing consumer utility in monetary terms, it specifies the economic surplus created in the negotiation. The distributive element determines how the economic surplus is actually distributed between the negotiating parties. Thus, in this part, the market participants claim their stake in the value created by the agreement. Negotiators acting strategically typically address integrative and distributive elements in parallel. They thereby limit their own ability to mutually maximize the economic surplus.

A negotiation mechanism or, more general, a market mechanism is typically judged with respect to four desirable economic characteristics: individual rationality, incentive compatibility, (ex-post) allocation efficiency, and budget balance [8]. The economic outcome of such a setting is restricted by a multitude of strong theoretic results: Given quasi-linear preferences, it is impossible to design a mechanism that achieves individual rationality, efficiency, and budget balance at once, regardless if incentive compatibility is fulfilled or not [8].[2] However, incentive compatibility is the prerequisite for an efficient outcome. Hence, the named characteristics need to be balanced. Efficient mechanisms can lead to a considerable need to plough in money, as the Vickrey-Clarke-Groves (VCG) mechanisms give proof of [e.g., 9]. Likewise, individually rational and budget balanced mechanisms can result in highly inefficient outcomes [10]. Additionally, budget balance and individual rationality are required to enable sustainability and implementability over time [11].

In this work, we evaluate three suitable mechanisms for bilateral negotiations on multiple attributes with respect to their economic properties. Two are stylized representations of existing mechanisms, one – DISCOUNTBIDDING – is newly introduced in this article. We maintain the requirements of individual rationality and budget balance, abandon incentive compatibility, and study the effect of strategic bidding on efficiency and distribution of economic surplus in different settings. Following, we define the market scenario in more detail and review related work. We then formalize the negotiation mechanisms, analyze strategic behavior and outcomes, and compare the mechanisms in settings under complete and incomplete information with different levels of risk and risk aversion. In the last section, we summarize the results, sketch the limitations, and outline future work.

[2] The result is derived in Theorem 1 of [8]. Their Theorem 2, provides the grounds for testing whether a given incentive-compatible, individually rational mechanism maximizes the expected economic surplus. This does, however, not apply to the present scenario, as its robustness depends on knowledge of both distributions of player types, it is restricted to single-attribute negotiations and a trusted third party is required.

2 Scenario & Related Work

In order to evaluate the performance of different negotiation mechanisms, we consult a market scenario that is typical for custom services such as Cloud services: Multiple consumers are interested to procure a service of a particular functionality. Price and quality attributes play a major role. The consumers are endowed with heterogeneous preferences; their individual willingness to pay subject to a certain quality level is not fully known to the public.

Multiple providers are present in the market. They offer one or more custom services that fit the functional requirements of the consumer, yet differ in their quality of service and price. A key assumption is that providers have different technologies and internal processes and, thus, different costs for service provision. Or they serve as so called service supply hubs, forming a single interface to the customer and using multiple service suppliers in the backend [12]. These factors and associate costs to the service provider are not fully known to the consumers.

Consumers and providers interact bilaterally. There is no third party running a central exchange. All parties act rationally and strategically in their own interest. Since consumers' and providers' types are not publicly known, the market exhibits uncertainty of the optimal matching. To facilitate online scenarios with low transactions costs, interaction is restricted to two-stage mechanisms, in preference to more complicated negotiation mechanisms. This simplification follows the standard assumption in bargaining models that delay of reaching an agreement is costly and, thus, rational agents should reach an agreement immediately with their first offers (cf. the seminal paper by [13] and the literature building on it). It remains future work to extend the analysis to more complex, iterative mechanisms.

Academia suggests several market mechanisms to be used in such a scenario. The simplest and most widely used mechanism is a *fixed price* for a specifically defined service without any further information exchange. Such a fixed price can either be set by the provider or the consumer. In both cases, the bidding language is very simple, yet very little information is transferred hindering the discovery of the optimal service specification given the provider's and consumer's types. In game theoretic terms, such a fixed price offer is an ultimatum game, the simplest form of a negotiation mechanism. A subtle variation commonly observed in practice is to not offer a single quality-price combination, but a small set of such combinations to choose from. This increases the likelihood of discovering the optimal service specification at the cost of complexity of decision making for both sides.

More complex mechanisms to tackle the multi-dimensionality are multi-attribute negotiations and auctions. They allow for the negotiation on non-price attributes [e.g., 14]. Many examples for electronic markets that can handle the complex sale or procurement of multi-attribute services and products through automated negotiation exist [15-17]. Such mechanisms' bidding languages are rich – a lot of information is exchanged. Multi-attribute auctions based on the family of VCG mechanism are incentive compatible and efficient, yet they suffer from the impossibility to balance the budget. Approaches to achieve budget balance by foregoing efficiency mostly result in highly complex mechanisms, either with respect to their accomplishment or, due to

complex transfer functions, with respect to the strategies. For instance, [2] introduces an iterative protocol while [11] proposes budget balanced approximations of VCG mechanisms, at the cost of sophisticated transfer functions which lead to complex strategy considerations.

In summary, common mechanisms in the scenario sketched above either reveal very little information, or their bidding language and/or transfer function is complex. The information exchanged is oftentimes biased by strategic behavior of the participants. This circumstance suggests that there is a trade-off between the simplicity of a mechanism and the information content exchanged between the negotiating parties. Moreover, strategic misrepresentation of preferences can be observed due to the mixture of integrative and distributive parts of the negotiations.

We seek a mechanism that fulfills individual rationality and budget balance and, at the same time, keeps efficiency as well as truthful information content high, while maintaining a simple bidding language. In order to evaluate how different mechanisms perform in the present setting, we scrutinize the following three mechanisms:

1. **TUPLEBIDDING:** One party poses a fixed price offer, i.e. a single price-quality tuple. The other party either accepts or not.
2. **SCORINGBIDDING:** One party proposes a complete scoring function over the set of possible price-quality combinations. The other party either selects one tuple described by this function or rejects to agree at all.
3. **DISCOUNTBIDDING:** Like SCORINGBIDDING, but in addition to the scoring function, the proposing party requests a price discount it requires on any tuple described by the scoring function.

TUPLEBIDDING and SCORINGBIDDING represent extreme cases on the continuum of simple and rich information exchange. They are stylized representations of commonly used fixed price and multi-attribute mechanisms. Intermediate versions like proposing multiple price-quality tuples and extensions like repeated offer exchanges are possible but offer only limited insight in qualitative differences of the mechanisms' mechanics.

DISCOUNTBIDDING is newly introduced as an extension of SCORINGBIDDING. The basic idea is to separate the integrative and distributive element of the negotiation with the scoring function allowing to identify the optimal price-quality combination and the discount factor allowing to claim value. All three mechanisms are formally characterized and evaluated in the following.

3 Formalization and Evaluation of Negotiation Mechanisms

In this section, we first formalize the model of preferences and the negotiation mechanisms. We then evaluate the mechanisms under complete information, i.e. the provider and the consumer knowing each other's types. While this setting is rather hypothetical, it serves as benchmark for the more realistic case of incomplete information. An evaluation of the mechanisms under incomplete information with different degrees of risk and risk aversion will provide a thorough understanding. A comparison of all mechanisms concludes this section and discusses which mechanism will prevail depending on who decides on the mechanism.

3.1 Model of Preferences and Key Assumptions

Consider two parties, service provider P and consumer C, both being individually rational and maximizing their utility. P and C negotiate over quality q and price p of a service. Quality q may be an aggregated abstraction of various quality attributes of interest to C [cf. 12]. Following standard micro-economic theory [e.g., 18] we assume that the consumer C has decreasing marginal returns from increasing quality, i.e. a concave scoring function. Analogously, we assume that the provider P has a production technology which yields decreasing marginal returns. An example: This assumption asserts that a provider who guarantees availability of a Cloud service requires more resources to increase availability from 98% to 99% than he needs to increase it from 94% to 95%. Adding the assumption of approximately linear costs of resources, this yields a convex cost function, which is again typical in micro-economic modeling [e.g., 18].

Operationalization of the assumptions: Both quality and price are individually normalized to the unit interval: $q \in (0,1)$, $p \in (0,1)$. This simplification allows for more comprehensible analytical considerations without significantly limiting interpretability of results. The provider's cost function and the consumer's scoring function are both modeled as monomials for three reasons: (1) It is a common approach in the related literature to capture each quality-attribute with one monomial [e.g., 1, 19]. (2) It allows for both simple and comprehensible analytical considerations along with computational tractability and (3) it ensures proximity to realistic scenarios under consideration of the preference elicitation challenges. In effect, we assume the provider P has a cost function $\mathrm{C}(\mathrm{q}) = \mathrm{q}^{\mathrm{b}}$ with $\mathrm{b} \in [1,\infty)$ representing the cost of providing a service of given quality. $b = 1$ yields a linear, $\mathrm{b} > 1$ a convex cost function with positive and increasing marginal costs of quality. The provider P has the quasi-linear utility function $U_{\mathrm{P}}(q) = p - C(q)$. For consumer C we assume the scoring function $S(q) = q^a$ with $a \in (0,b)$. $a > 0$ ensures strong monotonicity in quality, i.e. C prefers a higher quality over a lower quality, and $a < b$ ensures the existence of a mutually beneficial agreement. In practice and for a given provider and consumer it might be the case that no mutually beneficial agreement on a *q*-*p*-tuple exists. However, in this case no mechanism could yield an individually rational agreement. We omit this case and assume the existence of a mutually beneficial agreement for comparing the ability of different negotiation mechanisms in identifying the optimal agreement.

Under incomplete information, we allow for the consumer being risk neutral or risk averse. To capture C's risk preferences, we employ – analogous to the provider – a quasi-linear utility function additionally wrapped by a term that introduces constant relative risk aversion [e.g., 20]. C's utility function becomes $U_{\mathrm{C}}(q) = (S(q) - p)^{\frac{1}{1+r}}$; $r = 0$ implies C being risk-neutral, while an increasing $r > 0$ implies C being increasingly risk averse.

The economic surplus – also termed welfare – generated by an agreement is defined as sum of utilities $W(q) = U_{\mathrm{C}}(q) + U_{\mathrm{P}}(q)$. For a risk neutral consumer or for complete information, this simplifies to $W(q) = q^a - q^b$, i.e. welfare is determined by the

quality that the two negotiating parties agree on. The price is merely the mechanism for distributing welfare.[3]

The optimal quality q^* maximizing welfare is

$$q^* = \underset{q}{\operatorname{argmax}}\, W(q) = \left(\frac{b}{a}\right)^{\frac{1}{a-b}} \in (0,1) \quad \text{with} \quad w^* = \left(\frac{b}{a}\right)^{\frac{a}{a-b}} - \left(\frac{b}{a}\right)^{\frac{b}{a-b}} \tag{1}$$

Fig. 1 sketches an example of the above model with $a = 0.5$, $b = 2$ and resulting $q^* = 0.397$. The price is arbitrarily chosen as $p = 0.35 \in [C(q^*), S(q^*)]$ and both parties' utilities are positive.

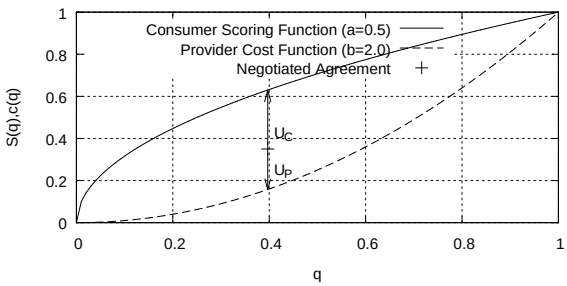

Fig. 1. Example for scoring and cost functions, optimal quality, price and utilities

Graphically, the assumptions on $C(q)$ and $S(q)$ assure the existence of an area of individually rational potential agreements in between the two curves. The objective of a negotiation is to determine a point $\{q, p\}$ within this area. It is in both parties' interest and Pareto optimal to choose $q = q^*$. This maximizes welfare. The price p anywhere on the intersection of q^* with the area distributes this welfare. For an omniscient arbitrator with given fairness perception, this is relatively easy. For P and C under incomplete information and strategic behavior, this is, however, very complex.

3.2 Negotiation Mechanisms

Three negotiation mechanisms can help P and C in their complex task. Without loss of generality, we assume the consumer C going first and the provider P responding to the offer. As the scenario and model are symmetric with respect to the roles, the evaluation with an inverted order of action follows analogously.

TUPLEBIDDING: C submits a binding bid $\{q, p\}$. P can accept this tuple as agreement or reject it leading to no agreement. Assuming myopic utility maximization within the negotiation, P accepts if $p \geq C(q)$ and rejects otherwise. For brevity of the analysis, we assume acceptance in case of indifference.

[3] In order to compare welfare across levels of risk aversion and in order not to overvalue the utility of C, we adapt the welfare function in case of $r > 0$ by back-transforming $U_C(q)$.

SCORINGBIDDING: C submits a binding bid of a scoring function $\hat{S}(q)$ or (in a less generic setting) the parameters of a scoring function, e.g., $\hat{a}$ for $\hat{S}(q) = q^{\hat{a}}$. C can (and will) choose $\hat{S}(q) \neq S(q)$, i.e. she can strategically misrepresent her type. P sets a quality $\hat{q}$ and price $\hat{p} \leq \hat{S}(\hat{q})$. $\{\hat{q}, \hat{p}\}$ is the agreement. If there exists a $\hat{q}$ with $\hat{S}(\hat{q}) \geq \mathrm{C}(\hat{q})$, P maximizes his utility with $\hat{p} = \hat{S}(\hat{q})$ and $\hat{q} = \underset{q}{\operatorname{argmax}}\, \hat{S}(q) - C(q)$. Otherwise, P rejects any agreement.

DISCOUNTBIDDING: C submits a binding bid of a scoring function $\hat{S}(q)$ or a single parameter, e.g., $\hat{a}$. In addition, C submits a discount value d. P sets a quality $\hat{q}$ and price $\hat{p} \leq \hat{S}(\hat{q}) - d$. $\{\hat{q}, \hat{p}\}$ is the agreement. If there exists a $\hat{q}$ with $\hat{S}(\hat{q}) - \mathrm{d} \geq \mathrm{C}(\hat{q})$, P maximizes his utility with $\hat{p} = \hat{S}(\hat{q}) - d$ and $\hat{q} = \underset{q}{\operatorname{argmax}}\, \hat{S}(q) - d - C(q)$. Otherwise, P rejects any agreement.

3.3 Evaluation

For the case of complete information, we derive optimal bidding behavior, utility, and welfare for both C and P analytically. With complete information, C knows P's cost function $C(q)$ and profit maximizing decision function. Analytically, one can derive the following equilibria: In either of the three mechanisms, provider P will be left with a utility of zero. Consumer C can claim the entire maximum possible surplus in both TUPLEBIDDING and DISCOUNTBIDDING, yielding a utility $\mathrm{U_C} = \left(\frac{\mathrm{b}}{\mathrm{a}}\right)^{\frac{\mathrm{a}}{\mathrm{a-b}}} - \left(\frac{\mathrm{b}}{\mathrm{a}}\right)^{\frac{\mathrm{b}}{\mathrm{a-b}}} = \mathrm{w}^*$. SCORINGBIDDING is sub-optimal from both welfare and consumer's perspective, with $\mathrm{U_C} = \mathrm{e}^{-\frac{\mathrm{a}}{\mathrm{b}}} - \frac{1}{\mathrm{e}} < \mathrm{w}^*$.[4]

Let us consider the more realistic case of incomplete information and risk aversion. Incomplete information implies risk in the decision making of C. C continues being perfectly informed with one exception: the provider's cost function, here operationalized as parameter b. C only knows that b is uniformly distributed: $b \sim \mathrm{U}(m-s, m+s)$ with m and s being public knowledge. s is a proxy for C's risk; the higher s, the higher the risk. P still knows b, i.e. his cost function. The case of complete information could be solved analytically; for incomplete information, we choose a numerical analysis with a wide range of parameters a, b, s, r as input to the simulation in order to test for sensitivities and assure robustness of the results. For brevity and simplicity, the following presentation focuses on a limited set of parameters ($a = 0.5$, $b = 2$, $r = 1$, $s \in [0,1]$).[5]

Three reasons make us believe that the numerical solution is accurate: (1) For the border case of complete information, the numerical and the analytical solution coin-

[4] Proofs can be obtained from the authors upon request.
[5] A wider range of results is available from the authors upon request.

cide; (2) all strategic effects and utility comparisons vary either smoothly within the parameter range or, in case of step functions, are straightforward to explain; (3) the qualitative effects hold true for all parameter configurations tested.

TupleBidding. Obviously, C's bid consisting of q and p has two effects: It determines the likelihood of reaching an agreement and the utility of an agreement, should it be achieved. In bidding, C trades-off these two directly opposed effects for maximizing her expected utility. In the border case of complete information ($s = 0$), bidding is straightforward and the numerical simulation coincides with the analytical solution presented above. Any $\{q, p\}$-tuple belongs to either of two sets: Either it leads to agreement with certainty or to disagreement with certainty. Among the tuples that lead to an agreement, C maximizes her utility by choosing the efficient allocation.

When introducing risk ($s > 0$), C's optimization problem becomes more tricky. For increasing s, C initially prefers bidding more conservatively, i.e. she foregoes utility from the agreement but assures reaching agreement with 100% certainty. She does so by simultaneously lowering q and increasing p (both to her disadvantage in case of agreement), i.e. by moving the bid to the upper left in Fig. 2. At one point – called "tipping point" t from here on –, s becomes excessive from C's viewpoint. C stops retracting and starts bidding more aggressive. From $s > t$ onwards, C simultaneously increases q and decreases p; she demands higher utility in case of agreement and takes the risk of not reaching an agreement.

Under complete information ($s = 0$), the efficient agreement is reached, C claims all the value, and P has a utility of zero. Welfare, i.e. the sum of utilities, is at its maximum. With increasing risk $s > 0$, C's utility decreases monotonically and increasingly rapidly – up to the tipping point (Fig. a). At the tipping point, C changes her strategy (see above): The monotonic decline of expected utility persists, but is slowed down. As C's risk increases, P's utility increases – again, up to the tipping point, when P's expected utility reaches its maximum and then sharply declines with further increasing risk (Fig. b). The effect on welfare is straightforward: Welfare is optimal for $s = 0$. With risk, the parties on average no longer agree on the efficient allocation; welfare declines gradually. With risk s beyond the tipping point t, when P's utility declines, the decline of welfare increases its speed (Fig. c).

ScoringBidding. C faces the same trade-off as with price-quality tuples: utility in case of agreement vs. the likelihood of reaching agreement. C's strategy in bidding a scoring function follows the same pattern as in bidding a price-quality tuple: With increasing risk (s), the consumer bids more conservatively. Technically, he decreases $\hat{a}$ and bids more truthfully. He does so sufficiently to assure certainty of reaching an agreement. Again, this holds up to the tipping point t, at which C starts gambling, i.e. she gradually increases $\hat{a}$ for claiming more value at the risk of not reaching agreement. With increasing risk aversion, the cost of gambling rises and, thus, the tipping point rises. Interestingly, the tipping point coincides with the tipping point for bidding price-quality tuples for any r. C's and P's expected utility as well as welfare all three qualitatively resemble the patterns from TUPLEBIDDING for varying s and r.

C's expected utility is maximal for $s = 0$ and decreases monotonically for increasing s (Fig. a). The decrease speeds up for increasing s up to the tipping point and slows down from $t > s$ onwards. P's expected utility is zero for $s = 0$, increases monotonically for increasing s up to the tipping point. It sharply kinks at $s = t$ and decreases thereafter (Fig. b).

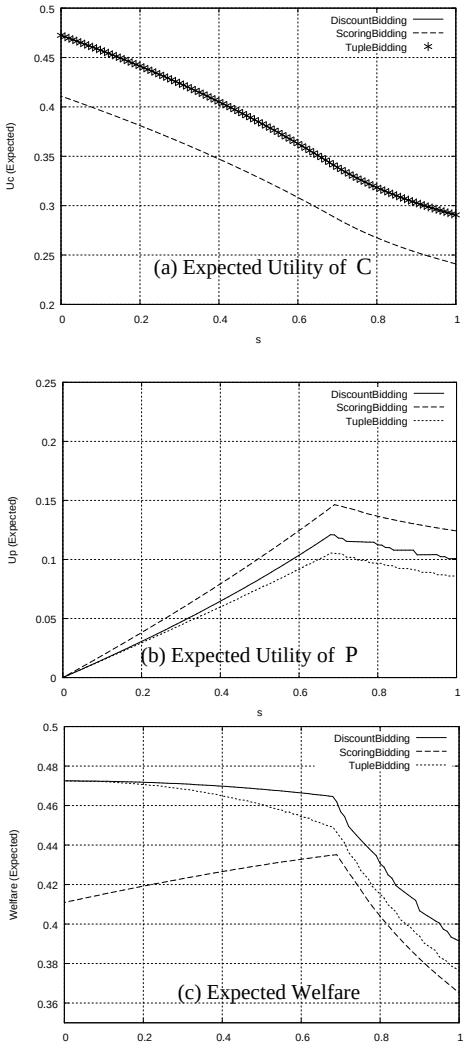

Fig. 2. Comparison of negotiation mechanisms with respect to utility and welfare

Interestingly, welfare is optimal for $s = t$. Strategic bidding by C hinders the integrative part of the negotiation. Risk initially lowers C's strategic misrepresentation and thereby allows P to come closer to the efficient quality. At $s = 0$, C submits a bid of $\hat{a}$ as close as possible to b, optimizing her own utility, yet sacrificing efficiency by overbidding by almost four times the real value a (for $a = 0.5$ and $b = 2.0$). As before, with increasing risk ($s > 0$), C will lower her bid, thus increasing efficiency simultaneously. After passing the utmost efficient point for $s = t$, C starts gambling by raising $\hat{a}$ and welfare starts to rapidly decline again (Fig. c).

DiscountBidding. C's trade-off is the same as for the other mechanisms, only the vehicle of bidding more or less conservatively differs. It can be implemented by either lowering $\hat{a}$, or by lowering d, or both. The interesting question is which of these vehicles C chooses to maximize her expected utility.

Analogously to the analytical results in case of complete information, bidding the scoring function truthfully ($\hat{a} = a$) and claiming value via d is the strategy maximizing C's expected utility. For the functional forms of $C(q)$ and $S(q)$ being strongly convex/concave, bidding the scoring function truthfully is the only optimum. Considering, for instance, the border case of $s = 0$: C knows the efficient allocation that maximizes her utility and needs to derive a bid $\{\hat{a}, d\}$ to meet that very allocation. With $\hat{a} \neq a$, the provider's optimization will yield a sub-optimal quality; value is lost and cannot be regained by any d. Numerical results show that the same holds true under risk. For more general scoring and bidding functions, however, bidding the scoring function truthfully might not be the only optimum – it is, however, always

one optimum, as long as P chooses the efficient quality q for a truthful bid. Again, the qualitative patterns in strategies and utilities resemble what is known from the other mechanisms: With increasing s, C initially bids more and more conservative, lowering d. The tipping point t exists at which this pattern inverts and C switches to gambling, i.e. to increasing d. Up to t, C's utility decreases in s and P's utility increases. For increasing $s > t$, the decline of C's utility slows down (Fig. a) and P's utility decreases as well (Fig. b). Welfare decreases in s with a kink at t (Fig. c). Increasing risk aversion r increases t but does not qualitatively affect behavior or outcomes. All variation in C's bidding depending on s and r takes place in the discount d; the revelation of the scoring function remains truthful ($\hat{a} = a$).

In summary, discount bidding effectively disentangles the integrative and distributive elements of the negotiation. Bidding the scoring function truthfully allows maximizing the consumer's expected utility and expected welfare in any given state of the world, i.e. for any combination of provider cost function (parameter b), consumer scoring function (parameter a), and consumer risk aversion (parameter r). A comparison of DISCOUNTBIDDING and SCORINGBIDDING shows that this truthfulness is only possible, as the discount factor d provides the consumer a vehicle to claim value, i.e. d is the distributive element of the negotiation. Truthfulness in the scoring function is an interesting and potentially beneficial property of the discount bidding mechanism. The even more interesting question is, however, how the three mechanisms compare in terms of utility for one or the other party, as this will drive adoption in the marketplace. The following section presents this comparison.

3.4 Implication for Adoption of Negotiation Mechanisms

This section provides a comparison of the three mechanisms regarding utility and welfare and discusses which mechanism will prevail depending on who decides on the mechanism. With complete information TUPLEBIDDING allows the consumer to choose the efficient allocation q^* while claiming the entire surplus by setting the price to the value of P's cost function. The same is possible when DISCOUNTBIDDING is implemented, as C can use the discount d to obtain the entire surplus, while ensuring efficiency by truthfully bidding $\hat{a}$. Thus, under complete information, both mechanisms result in the same allocation, utility and welfare.

Intuitively, SCORINGBIDDING cannot achieve the same efficiency, as C is always tempted to misrepresent her type by overbidding $\hat{a}$, thus leading to a quality other than q^*. Despite the fact that it yields the same utility for P, it turns out that SCORINGBIDDING leads to a lower utility for C, which also implies a lower welfare than when using TUPLEBIDDING or DISCOUNTBIDDING as negotiation mechanism.

With incomplete information, Fig. a shows that for any given risk s, the consumer's expected utility is equal for DISCOUNTBIDDING and TUPLEBIDDING (the lines are exactly on top of each other) and strictly lower for SCORINGBIDDING. This result is independent of the consumer's risk aversion. The implication is twofold: (1) When the consumer can choose the mechanism, she will prefer either DISCOUNTBIDDING or TUPLEBIDDING over SCORINGBIDDING. The consumer may have the ability to choose

the mechanism when either there are multiple providers offering different mechanisms, or when she has sufficient purchasing power to dictate the mechanism. (2) When the consumer has the chance to reduce her risk at a reasonable cost in terms of time and money, she will do so. She may have the chance to acquire information either by standard market research, or by learning from repeated negotiations with a single or with multiple providers.

From the provider's view, there is a clear ranking of the three mechanisms (Fig. b): SCORINGBIDDING is preferred over DISCOUNTBIDDING which is preferred over TUPLEBIDDING. This holds for all $s > 0$; for $s = 0$, the provider is indifferent between the three mechanisms. The implication is, again, twofold: (1) When the provider chooses the mechanism, he will choose SCORINGBIDDING. (2) When the provider can influence the consumer's risk, he will do so. He will try to provoke conservative bidding by the consumer but will not exaggerate the risk to a level where the consumer starts gambling and the provider risks not reaching an agreement at all. The provider can, e.g., influence risk by withholding information on the exact technology employed and on his costs. Both are common in real-world settings. The provider has, however, an interest on providing some information on technology and, thus, on implied costs. Again, this is a common behavior in real-world settings. The ranking of mechanisms from the provider's viewpoint is independent of the consumer's risk aversion. The provider's strategic manipulation of the consumer's risk, e.g., by providing or withholding information on the technology and associated costs, depends on the consumer's risk aversion. This, in turn creates risk on the provider side with regard to the exact degree of the consumer's risk aversion. The analysis of equilibria in this extended game is beyond the scope of this paper.

The welfare perspective combines the above perspectives and, again, provides a ranking of mechanisms. Interestingly, for the case of the consumer being risk averse, this ranking depends on the level of risk (Fig. c). For low risk, welfare from TUPLEBIDDING is higher than from SCORINGBIDDING. For high risk and risk aversion, this ranking inverts. In any case, DISCOUNTBIDDING results in at least the same welfare for the special case of $s = 0$ and strictly higher expected welfare for the general case of $s > 0$. This result leads to the following implication: When a third party can choose the mechanism and intends to maximize welfare, it will select DISCOUNTBIDDING, independent of risk and risk aversion. Note, however, that in the scenario studied in this paper, it is rather unlikely that a third party like a regulator imposes a mechanism to the bilateral negotiation.

Looking beyond the scope of a single negotiation, two observations stand out: Firstly, the provider will strategically manipulate the consumer's risk prior to the negotiation, if he has the chance of doing so (which will hold true in most real-world settings). His challenge is to find the optimal degree of risk, as neither a low nor a high risk by the consumer is optimal for the provider. Secondly, DISCOUNTBIDDING is an attractive mechanism, but far from certain to be adopted in practice. In the extended game of selecting a mechanism, TUPLEBIDDING is not Pareto-optimal, it is dominated by DISCOUNTBIDDING. DISCOUNTBIDDING will prevail when either the consumer or a third party have the discretion or power to impose a mechanism. Bidding scoring functions will be adopted when the provider decides on the mechanism.

A positive side effect exists that might increase the provider's utility from DISCOUNTBIDDING in the long term: DISCOUNTBIDDING promotes truthful revelation of the consumer's scoring function and, thereby, allows the provider to optimize his technology portfolio and cost structure in the long-term. This effect – that is not reflected in the provider's utility function in this paper – might lead to all parties unanimously preferring discount bidding over the other mechanisms.

4 Conclusion & Future Work

This paper is set in the scenario of IT-based individualized services that require matching of functional requirements and in addition the negotiation of non-functional aspects. Specifically, we studied bilateral negotiations on quality and price of a service between the service provider and the consumer. In this setting, a negotiation has an integrative facet – many possible qualities of service are sub-optimal, but by means of communication within the negotiation mechanism, the parties can identify a Pareto-optimal quality of service. On the other hand, the negotiation has a distributive facet – either party has an interest in claiming as large a share in the value from an agreement as possible. Strategic bidding typically leads to negotiators mixing the integrative and distributive facets which results in inefficient outcomes.

In such a scenario, the selection of an "optimal" or at least "satisfying" negotiation mechanism is a challenge. In the light of mechanism design theory [8], we postulated individual rationality and budget balance for analyzing economic properties of different negotiation mechanisms, namely in how far agents' negotiation strategies deviate from truthful revelation of their types and in how far efficiency of negotiated agreements deviates from the efficient agreement an omniscient arbitrator would define.

On this theoretical background, we compared three negotiation mechanisms: TUPLEBIDDING, SCORINGBIDDING and DISCOUNTBIDDING. TUPLEBIDDING serves as a proxy for commonly used fixed price mechanisms; SCORINGBIDDING resembles the widely used approach of agents bidding a scoring function, e.g., in multi-attribute auctions. DISCOUNTBIDDING was newly introduced in this paper – it allows bidding a scoring function and additionally a discount that the consumer demands from the provider. The intuition is that this approach disentangles the integrative and distributive facets of the negotiation and increases efficiency. Our results confirmed this intuition: For complete information, we derived these results analytically from the game theoretic equilibrium; for the extended case of incomplete information, risk and risk aversion, we used a numerical simulation to characterize the mechanisms.

TUPLEBIDDING is not Pareto-optimal but dominated by DISCOUNTBIDDING. Thus, it is unlikely to prevail in a marketplace for custom services as soon as this market matures. Nowadays, comparable mechanisms are used by some Cloud service providers, for example. It has the advantage of a simple bidding language and very little communication effort. However, as providers and consumers get more sophisticated and as automated negotiations become more prevalent, the disadvantageous economic properties will weigh heavier and TUPLEBIDDING might become less relevant.

SCORINGBIDDING emphasizes the integrative facet and yields higher expected utility for the provider than either of the other mechanisms. When the provider can dictate the choice of the mechanism, he will presumably favor SCORINGBIDDING. However, the consumer's strategic misrepresentation of her scoring function leads to sub-optimal agreements. DISCOUNTBIDDING captures even more of the integrative facet: it promotes truthful revelation of the consumer's scoring function and thereby allows reaching an efficient agreement more often than with SCORINGBIDDING. The expected welfare from DISCOUNTBIDDING is higher than from SCORINGBIDDING for any level of risk and risk aversion. Compared to SCORINGBIDDING, the discount factor in DISCOUNTBIDDING shifts utility from the provider to the consumer. Whenever the consumer or an independent third party can dictate the negotiation mechanism, she will tend to favor DISCOUNTBIDDING. A positive long-term effect of DISCOUNTBIDDING is that truthful revelation of the consumer's scoring function allows the provider to adapt his technology and service offering. In the long run, this may even overturn the provider's favoritism for SCORINGBIDDING.

All mechanisms show a tipping point in the consumer's behavior depending on the risk: with risk below this tipping point, the consumer bids conservatively and assures reaching an agreement; beyond this risk, she bids aggressively and risks not reaching an agreement. For a given level of risk aversion, this tipping point is – somewhat surprisingly – identical for all three mechanisms.

In each mechanism, the provider has an incentive to strategically manipulate the consumer's risk via the information provided on his technology and costs. Neither full transparency nor opacity are optimal for the provider; he will have to carefully chose the level of information depending on the consumer's risk aversion. The higher the consumer's risk aversion, the less information will be given by the provider. All results hold inversely when inverting the roles of consumer and provider.

The presented work has four main limitations: (1) We assume individually rational utility maximizing agents. (2) The results depend on the model of preferences, especially on the functional forms of cost, scoring, and utility functions. While we believe that similar results can be obtained for other preferences, this has not been proven yet. (3) We study "only" three distinct mechanisms without deriving an "optimal mechanism" in the mechanism design sense. However, given the complexity and impossibility theorems, we believe this comparison of existing mechanisms and introduction of DISCOUNTBIDDING as additional mechanism is a valuable contribution to the field. (4) All three mechanisms only allow for a single offer and its acceptance or rejection by the counterparty. More complex negotiation mechanisms with an alternating offer exchange are possible, so is the introduction of a central marketplace.

Future work will address these limitations and, for instance, study other formalizations of preferences, behavior by human agents, more complex negotiation mechanisms, and the comparison of bilateral mechanisms with a centralized exchange.

References

1. Bichler, M., Kalagnanam, J.: Configurable offers and winner determination in multi-attribute auctions. European Journal of Operational Research 160, 2, 380–394 (2005)

2. Parkes, D., Kalagnanam, J.: Models for iterative multiattribute procurement auctions. Management Science 51, 3, 435–451 (2005)
3. Asker, J., Cantillon, E.: Properties of scoring auctions. The RAND Journal of Economics 39, 1, 69–85 (2008)
4. Gimpel, H., Ludwig, H., Dan, A., Kearney, B.: PANDA: Specifying Policies for Automated Negotiations of Service Contracts. In: Orlowska, M.E., Weerawarana, S., Papazoglou, M.P., Yang, J. (eds.): Service-Oriented Computing – ICSOC 2003. LNCS, Vol. 2910, pp. 287–302. Springer, Heidelberg (2003)
5. Ströbel, M.: Engineering electronic negotiations. Kluwer, New York (2002)
6. Jennings, N.R., Faratin, P., Lomuscio, A.R., Parsons, S., Sierra, C., Wooldridge, M.: Automated. Negotiation: Prospects, Methods and Challenges. International Journal of Group Decision and Negotiation 10, 2, 199–215 (2001)
7. Garg, S.K., Buyya, R.: Market-Oriented Resource Management and Scheduling: A Taxonomy and Survey. In: Obaidat, M.S., Misra, S. (eds.): Cooperative Networking, pp. 277–306. John Wiley & Sons, Chichester (2011)
8. Myerson, R., Satterthwaite, M.: Efficient mechanisms for bilateral exchange. Journal of Economic Theory 28, 265–281 (1983)
9. Krishna, V.: Auction Theory. Academic Press, San Diego (2002)
10. McAfee, R.P.: A dominant strategy double auction. Journal of Economic Theory 56, 2, 434–450 (1992)
11. Parkes, D.: Iterative combinatorial auctions: Achieving economic and computational efficiency. Ph.D. thesis, University of Pennsylvania (2001)
12. Leimeister, J.M.: Dienstleistungsengineering und -management. Springer, Heidelberg (2012)
13. Rubinstein, A.: Perfect Equilibrium in a Bargaining Model. Econometrica 50, 1, 97–110 (1982)
14. Son, S., Kim, K.M.: A Price- and-Time-Slot-Negotiation Mechanism for Cloud Service Reservations. IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics 42, 3, 713–728 (2012)
15. Bichler, M., Kalagnanam, J.: Software frameworks for advanced procurement. Communications of the ACM 49, 12, 104–108 (2006)
16. Gimpel, H., Mäkiö, J.: Towards Multi-Attribute Double Auctions for Financial Markets. Electronic Markets – The International Journal on Networked Business 16, 2, 130–139 (2007)
17. Bichler, M., Pikovsky, A., Setzer, T.: An analysis of design problems in combinatorial procurement auctions. Business and Information Systems Engineering 1, 111–117 (2009)
18. Varian, H.R.: Intermediate Microeconomics: A Modern Approach, vol 7. Norton, W. W. & Company, New York (2005)
19. Blau, B., Conte, T., Weinhardt, C.: Incentives in Service Value Networks – On Truthfulness, Sustainability, and Interoperability. ICIS 2010 Proceedings, Paper 8. Association of Information Systems, Saint Louis (2010)
20. Eisenführ, F., Weber, M., Langer, R.: Rational Decision Making. Springer, Heidelberg (2010)