

Electronic service matching: Failure of incentive compatibility in Vickrey auctions

Tobias Widmer^{a,*}, Joerg Leukel^a

^a*Faculty of Business, Economics and Social Sciences, University of Hohenheim*

*Corresponding author. Postal address: University of Hohenheim, Schwerzstr. 35, 70599 Stuttgart, Germany. Phone: +49 711 459 22274; Fax: +49 711 459 22961.

Email addresses: `tobias.widmer@uni-hohenheim.de` (Tobias Widmer),
`joerg.leukel@uni-hohenheim.de` (Joerg Leukel)

This is the authors' version of the following journal article:

Widmer, T. and Leukel, J. (2018). Electronic service matching: Failure of incentive compatibility in Vickrey auctions. *Operations Research Letters*, 46(3), pp. 318-323.
<https://doi.org/10.1016/j.orl.2018.03.004>

©2018. This manuscript version is made available under the CC-BY-NC-ND 4.0 license
<http://creativecommons.org/licenses/by-nc-nd/4.0/>

Abstract

We consider pricing schemes for matching customers and providers on double-sided markets for electronic services. While existing second-best solutions are incentive compatible, the associated payment functions are difficult to implement in real-world settings. Based on the VCG and the k -pricing mechanism, we propose two straightforward payment rules that offer a practical alternative to the second-best solution. Our experiments provide evidence that the VCG payments fail to implement incentive compatibility. This failure is due to the interdependency of the participants' utilities.

Keywords: Mechanism design, electronic services, pricing, double-sided markets, Vickrey auction

1. Introduction

Matching the right pairs of competitive customers and providers for electronic services on double-sided markets is an important optimization problem in Operations Research [2]. On these markets, multiple service providers offer electronic services of a specified quality of service (QoS), while multiple customers demand these services at a specific QoS. To this end, matching markets have recently emerged in different business areas including Cloud computing [1].

In common market settings, strategic participants may engage in bid manipulation in order to influence their transaction prices. While first-best solutions are not available in such environments [10], second-best mechanisms for matching customers and providers with interdependent utilities exist [15]. Although such mechanisms satisfy incentive compatibility and individual rationality, the associated payment structures are difficult to implement in real-world scenarios. Attempts to simplify the payment scheme, however, may open the way for strategic participants to increase their utilities by misrepresenting their bids. Thus, before modifying the payments, the mechanism designer must obtain an accurate estimation for the potential utility gain that participants can achieve due to strategic bid manipulation.

Prior research studies the average utility gain of participants with independent utilities on markets for electronic services. The mechanism proposed by Schnizler et al. uses k -pricing to provide a simple payment scheme that is well-suited for real-world electronic service exchange [12]. While their approach allows for estimating the utility gain of strategic participants, it does not consider interdependent utilities. Lee analyzes the manipulability of stable matching mechanisms to quantify the utility gain participants can expect through bid manipulation [8]. Yet concrete payment schemes for real-world markets are missing. In the context of generalized assignment problems, Fadaei and Bichler propose truthful approximation mechanisms in payment-free environments [3]. Fadaei and Bichler use the optimal welfare value as a benchmark to estimate the efficiency loss due to strategic bidding. Because they consider mechanism design without money, no payment rules are provided. Widmer and Leukel [15] provide a lower bound for the efficiency of a second-best mechanism that allocates electronic services with private quality information. Although they specify the incentive compatible payments of customers and providers in double-sided markets, these payment rules turn out to be inexpedient for implementing the associated mechanism in real-world environments.

The objective of our research is to study the efficiency loss of a mechanism with two straightforward payment schemes for electronic service matching in double-sided markets. We apply these two payment schemes to markets

where participants have interdependent utilities. The first payment scheme is based on the prominent Vickrey-Clarke-Groves (VCG) pricing rules [14], which satisfy incentive compatibility in single-unit and certain multi-unit procurement auctions [5]. The second payment scheme is based on k -pricing introduced by Satterthwaite and Williams [11], where the price is simply calculated as the arithmetic mean of customer valuation and provider cost. Mechanisms with k -pricing, however, are not incentive compatible [10]. In a set of simulation experiments, we study the potential utility gains that participants can expect through strategic bid manipulation.

We find that in our model with interdependent utilities, the prominent VCG mechanism is not incentive compatible. The reason for the failure of incentive compatibility is that strategic customers and providers can manipulate their perceived impact on social surplus through the utilities of their matched partners. Hence, participants are able to increase their utilities by manipulation even when VCG pricing is used. The evaluation results for both payment schemes provide an accurate estimation for these utility gains.

2. Formal framework

There are two disjoint sets of participants in the market, namely customers and providers. In order to obtain benchmark results for arbitrary market settings, we begin by assuming an equal number N of customer and providers on each market side in this article. This assumption is consistent with prior research investigating mechanisms for allocating electronic services [12]. We revisit this assumption in our discussion of future research directions (cf. section 4.3).

Each provider attempts to sell an electronic service to one customer, and each customer seeks to buy a service from one provider. Each customer i demands for a privately known QoS θ_i , and each provider offers their service at a privately known QoS σ_j .

If customer i receives the service of provider j , i produces the pairwise private valuation $v(\theta_i, \sigma_j)$. This valuation depends on i 's desired quality, as well as on the difference between its own desired quality and provider j 's actual quality. This situation depicts a market in which customer valuation functions are non-monotonic in the quality offered by the provider. For instance, a customer may prefer a service with medium over high computational capacity. A high-capacity service is well able to process many simultaneous requests from the customer's application that uses this service. If, however, this application does not have enough computational power or resources, the application will fail to answer these simultaneous requests in due time. This failure leads to higher buffering in the application and thus longer response times. Therefore, a customer's valuation must

take into account the application that uses the service and the tradeoff between being idle or buffering heavily [6]. Therefore, we assume that any mismatch in desired quality and actual quality creates adjustment problems for the customer. That is, $v(\theta_i, \sigma_j)$ is maximized when the supplied quality and the desired quality are equal (i.e., when $\theta_i = \sigma_j$). By assumption, the maximal value is increasing in θ_i .

On the supply side, if provider j sells a service to customer i , j accrues a service provision cost $c(\sigma_j, \theta_i)$, which depends on the actual quality and on the difference between the own actual quality and the customer's desired quality. If a provider produces a quality lower than the quality desired by a customer, this provider incurs higher cost from not fulfilling the requirements. If, in contrast, a provider maintains higher quality than desired, their cost increases due to idle resources [4]. Hence, we assume that $c(\sigma_j, \theta_i)$ is minimized when $\sigma_j = \theta_i$ and that the minimal value is increasing in σ_j . This assumption captures the fact that a mismatch in actual quality and desired quality creates higher provision cost resulting from after-sales customer service cost and missed opportunity cost. Both $v(\theta_i, \sigma_j)$ and $c(\sigma_j, \theta_i)$ are assumed to be thrice differentiable.

Customers and providers use quasi-linear utilities. Hence, customer i paying t_c for receiving an electronic service from j obtains a payoff of

$$u_c(\theta_i, \sigma_j) = v(\theta_i, \sigma_j) - t_c, \quad (1)$$

and provider j receiving t_p for delivering the service to customer i obtains a payoff of

$$u_p(\sigma_j, \theta_i) = t_p - c(\sigma_j, \theta_i). \quad (2)$$

This research takes the perspective of a social planner, who is interested in maximizing the sum of the participants' welfare. Therefore, the social planner aims at maximizing the social surplus among all participants. Let $x_{ij} \in \{0, 1\}$ denote the decision variable, which is 1 if customer i receives the electronic service from provider j in the final allocation, and 0 otherwise. Thus, the mechanism faces the following optimization problem:

$$\max_{x_{ij}} \sum_{i=1}^N \sum_{j=1}^N (v(\theta_i, \sigma_j) - c(\sigma_j, \theta_i)) x_{ij} \quad (3)$$

$$\text{s.t. } 0 \leq \sum_j x_{ij} \leq 1 \quad \forall i \quad (4)$$

$$0 \leq \sum_i x_{ij} \leq 1 \quad \forall j. \quad (5)$$

The expression in (3) adds up the pairwise match surplus across all customers and providers and determines the allocation that maximizes the social welfare. Notice that the payments t_c and t_p do not appear in (3) because they add up to zero due to the budget balance constraint. Constraints (4) and (5) ensure that each customer is matched to exactly one service provider in the final allocation.

3. Mechanism definition

3.1. Allocation rule

The allocation rule of the mechanism must ensure that the final allocation of customers and providers maximizes the social welfare defined in (3). In many auction settings, it is difficult to determine the winners of the auction due to computational complexity issues. In supermodular environments, however, it turns out that the allocation rule, which maximizes the social welfare, adopts a rather simple form. Let $\rho_\theta(\theta_i) = |\{\theta_k \in \theta : \theta_k \geq \theta_i\}|$ be the rank of desired quality θ_i within the vector of all customers' desired qualities $\theta = \{\theta_1, \dots, \theta_N\}$. Define $\rho_\sigma(\sigma_j)$ similarly for providers. Then, the allocation rule is given by

$$x_{ij} = \begin{cases} 1 & \text{if } \rho_\theta(\theta_i) = \rho_\sigma(\sigma_j) = k \wedge v(\theta_i, \sigma_j) - c(\sigma_j, \theta_i) \geq 0, \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

In other words, the mechanism accepts the bids θ_i of all customers and the bids σ_j of all providers, sorts each side in descending order, and allocates customers and providers that are on the same rank from top to bottom. That is, the allocation rule of the mechanism is positively assortative. It maximizes the optimization problem in (3) because the pairwise surplus $v(\theta_i, \sigma_j) - c(\sigma_j, \theta_i)$ is a supermodular function. If the pairwise surplus function is supermodular, the optimal match function is positively assortative [13].

3.2. Pricing

After having obtained the welfare maximizing allocation rule in (6), it is crucial to determine the payments made by the participants for electronic service allocation. For designing an efficient mechanism, these payments must guarantee that no participant has an incentive to deviate from their true bid. That is, the payments must ensure incentive compatibility. It is well-known that a pricing scheme based on the VCG mechanism [14] is incentive compatible for a single customer who buys one unit of a product from a set of providers [5]. Moreover, the VCG mechanism ensures incentive compatibility in settings with many customers and many providers that exchange specific electronic services. On such markets, however, budget balance cannot be achieved [12]. A practical alternative for VCG payments is a pricing scheme based on k -pricing. The k -pricing scheme determines the payments by equally splitting the difference between the bids of customers and providers [11]. Appropriate VCG and k -pricing payments for the proposed mechanism are introduced in the following sections.

3.2.1. VCG pricing

In a VCG mechanism, each participant's impact on the social welfare is internalized through their payments such that the other participants receive the same payoff, regardless of the participant's bid [9]. That is, the payments of each participant is set equal to their impact on social surplus relative to the reports. To be more specific, let $\hat{\theta}_i$ denote the desired QoS submitted to the mechanism by customer i and let $\hat{\sigma}_j$ denote the offered QoS submitted to the mechanism by provider j . Then i 's payment for service consumption is

$$t_c^i = - \sum_{\substack{k=1 \\ k \neq i}}^N \sum_{j=1}^N (v(\hat{\theta}_k, \hat{\sigma}_j) - c(\hat{\sigma}_j, \hat{\theta}_k)) x_{kj} + \max_{x'_{kj}} \left(\sum_{\substack{k=1 \\ k \neq i}}^N \sum_{j=1}^N (v(\hat{\theta}_k, \hat{\sigma}_j) - c(\hat{\sigma}_j, \hat{\theta}_k)) x'_{kj} \right). \quad (7)$$

The compensation payments for providers t_p^j are defined analogously:

$$\begin{aligned}
t_p^j &= \sum_{i=1}^N \sum_{\substack{k=1 \\ k \neq j}}^N (v(\hat{\theta}_i, \hat{\sigma}_k) - c(\hat{\sigma}_k, \hat{\theta}_i))x_{ik} \\
&\quad - \max_{x'_{ik}} \left(\sum_{i=1}^N \sum_{\substack{k=1 \\ k \neq j}}^N (v(\hat{\theta}_i, \hat{\sigma}_k) - c(\hat{\sigma}_k, \hat{\theta}_i))x'_{ik} \right). \tag{8}
\end{aligned}$$

As stated above, the VCG mechanism is incentive compatible in common auction theory settings, for instance, when customers and providers have unit demand and supply and private information about their QoS. Similarly, public good environments entail incentive compatibility of the VCG mechanism. Considering the interdependent structure of the participants' utilities in this work, however, requires a reassessment of VCG's incentive properties.

3.2.2. *k*-pricing

In the *k*-pricing scheme, the transfers are calculated based on the difference between the valuation and provision cost of the matched participants. In general, these transfers are given by $kv(\theta_i, \sigma_j) + (1 - k)c(\theta_i, \sigma_j)$ for any $k \in [0, 1]$. Since each final allocation is made up of exactly two participants, *k* is set to 0.5. Using this pricing scheme, both individual rationality and budget balance are satisfied by definition. Hence, the mechanism with *k*-pricing is given by the allocation rule defined in (6) with transfers

$$t_c^i = t_p^j = \frac{1}{2}(v(\theta_i, \sigma_j) + c(\sigma_j, \theta_i)). \tag{9}$$

By the Myerson-Satterthwaite-Theorem [10], incentive compatibility fails to hold when ex post optimality is required. Therefore, the mechanism with *k*-pricing cannot be incentive compatible.

Example 1. *Suppose there are two customers and two providers in the market who submit their QoS bids truthfully. As suggested in our prior work [15], customers use a valuation function equal to $v(\theta_i, \sigma_j) = 1 + \sqrt{\theta_i} - (\theta_i - \sigma_j)^2$ and providers use a cost function equal to $c(\sigma_j, \theta_i) = \sigma_j^2 + (\theta_i - \sigma_j)^2$. All QoS realizations θ_i and σ_j are uniformly distributed over the unit interval. Table 1 shows the average utilities of customers and providers as well as the average budget achieved by the VCG and the *k*-pricing mechanism. While the mechanism with *k*-pricing balances the budget, the VCG pricing scheme runs a budget deficit. Notice that the difference between welfare and budget obtained by the VCG mechanism equals the welfare*

Table 1: Average utility and budget with uniformly distributed QoS.

| Pricing scheme | Cust. util. | Prov. util. | Welfare | Budget |
|----------------|-------------|-------------|---------|---------|
| VCG | 1.3295 | 1.1130 | 4.8849 | -2.6335 |
| k -pricing | 0.5629 | 0.5629 | 2.2516 | 0 |

achieved by the k -pricing mechanism. Because k -pricing does not satisfy incentive compatibility, the VCG pricing scheme cannot be incentive compatible either. The experimental evaluation confirms this finding.

4. Experimental evaluation

This section reports an experimental evaluation of the two pricing schemes developed in this proposal. We describe the setup, report the results, and discuss the findings.

4.1. Experimental setup

The experimental setup is based on the setting used in Example 1. The participants' QoS are described by random variables drawn from the uniform distribution over the unit interval and the normal distribution truncated to the unit interval with mean $\mu = 0.5$ and standard deviation $\sigma = 0.1$.

Customers and providers are assumed to misrepresent their QoS bids. We only consider linear misrepresentations of participants; that is, participants use a manipulation factor to distort their true QoS. For instance, a manipulation factor of 0.8 means that all customers reduce their true QoS by 20% and all providers increase their true QoS by 20%. In each experiment, the number of manipulating participants is varied as well. For example, a ratio of 0.7 denotes that 70% of customers and 70% of providers engage in bid manipulation, while the remaining 30% report truthfully.

We study the average utility gain each participant can achieve by submitting manipulated QoS values to the respective mechanism. The average utility gain of a participant is calculated as the ratio between the average utility obtained by manipulation and the average utility obtained by truthful bidding. Consequently, an average utility gain of less than 1 implies that, on average, no participant can increase their utility by manipulation. For values greater than 1, however, participants have an incentive to manipulate their bid. An average utility gain of 1.12, for example, implies that each participant can increase their utility by 12% on average if they distort their bid.

4.2. Results

Figure 1 shows the average utility gain a single, high-quality customer can generate if they distort their QoS in the mechanism with VCG pricing. All other participants are assumed to report truthfully. If the market contains four participants ($N = 2$) with QoS drawn from the uniform distribution, the manipulating customer is able to increase their utility once their manipulation factor exceeds 0.45 (i.e., their true QoS is lowered by 55%) and obtains a relative maximum utility gain of 4.7% at a manipulation factor of 0.7 (i.e., lowered by 30%). If the participants use normally distributed QoS, the manipulating customer can increase their utility using a manipulation factor that exceeds 0.3. Like in the uniform case, their utility gain is maximized at a factor of 0.7, where their relative utility gain is 7.3%.

When the market contains 200 participants ($N = 100$) that use uniformly distributed QoS, the manipulating customer cannot improve their utility at all. If the QoS is drawn from the normal distribution, however, they are able to improve their utility at a manipulation factor that exceeds 0.75. Finally, if this customer uses a manipulation factor of 0.9, their relative maximum utility gain appears to be 1%.

In contrast to the previous experiment, the following experiment assumes a market on which multiple participants manipulate their bids. Figure 2 depicts the utility gain each customer can expect on average if 10%, 20%, ..., 100% of all participants manipulate their bids by manipulation factors between 0.7 (i.e., customers lower by 30%, providers raise by 30%) and 1 (i.e., no manipulation). If all participants engage in manipulation (manipulating participants = 100%), the average utility gain of a customer depicted in Figure 2 arrives at its maximum of 7.6% at a manipulation factor of 0.85. In other words, if all customers lower their QoS by 15% and all providers raise their QoS by 15%, each customer can increase their utility by 7.6% on average. If over 70% of all participants manipulate their QoS by more than 33%, no customer can expect any utility gain (values are below 1).

Figure 3 presents the average customer utility gain for participants whose QoS is normally distributed in the mechanism with VCG pricing. The maximum relative utility gain a customer can achieve on average is 7.9% (peak of surface) at a manipulation factor of 0.45 and with 60% manipulating participants. When both the number of manipulating participants and the manipulation percentage increase at the same time, no customer is able to improve their utility by dishonest bid reporting.

Figures 4 and 5 depict the average utility gain of providers in the mechanism with k -pricing assuming uniformly and normally distributed QoS. For QoS

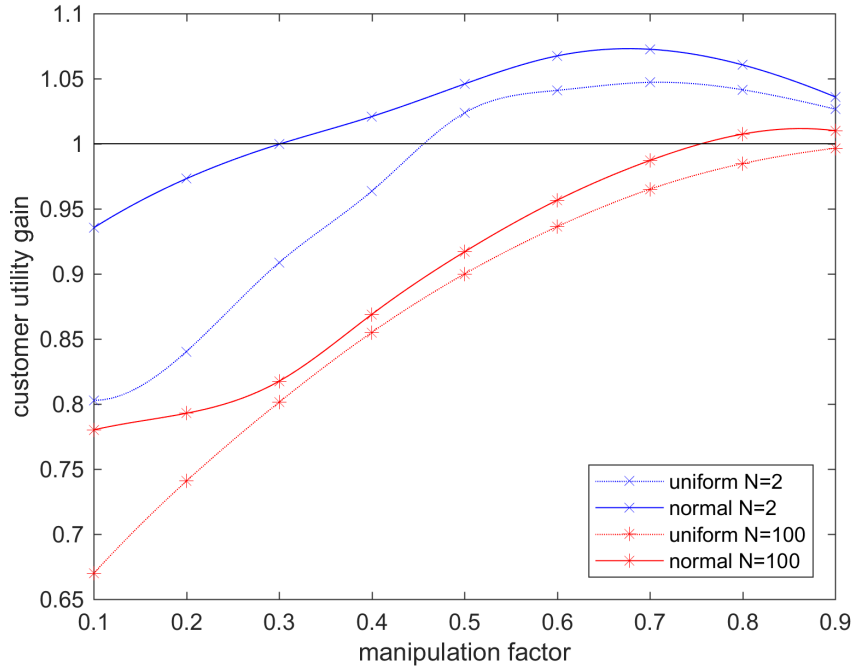


Figure 1: Average utility gain of a single manipulating customer (VCG).

realizations drawn from the uniform distribution (cf. Figure 4), each provider can expect a maximum utility gain of 3.3% on average when each participant manipulates their report by 15% (manipulation factor of 0.85) and all participants engage in manipulation (manipulating participants = 100%). Once customers and providers manipulate by more than 60%, however, no provider can anticipate any average gain in utility regardless of how many participants engage in misrepresentation.

Figure 5 shows the average provider utility gain for normally distributed QoS in the mechanism with k -pricing. If all participants manipulate their QoS at a manipulation factor of 0.75 (i.e., 25% manipulation), each provider can expect the maximum average utility gain of 1.5%. For high manipulation factors (greater than 35%), providers cannot improve their utility anymore (values are below 1).

4.3. Discussion

Our experiments demonstrate the impact of bid manipulation on the utilities of customer and providers in different market settings. Although the proposed pricing schemes fail to provide adequate incentives for participants to report their

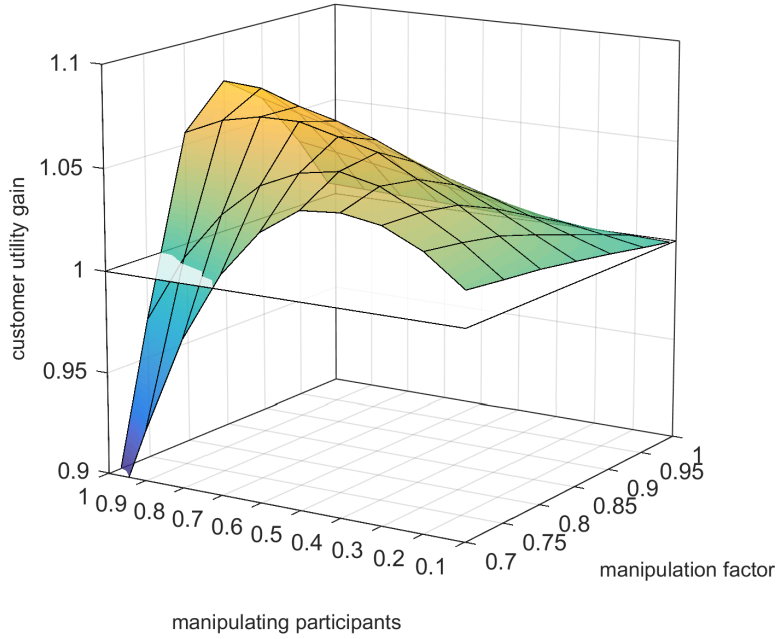


Figure 2: Customer utility gain with uniform QoS and $N = 10$ (VCG).

QoS truthfully, the results suggest that the relative utility gain of a participant does not exceed 8% on average in any market setting. This finding provides evidence for the efficacy of the proposed pricing mechanisms in double-sided markets. In the following paragraphs, we discuss the insights that can be obtained from our research.

First, we find that the mechanism with VCG pricing does not satisfy incentive compatibility in our model. This result is surprising because in common auction or public good environments, the VCG generally implements the efficient outcome [9]. As shown in the payment definitions (7) and (8), the VCG pricing scheme is based on the participant's impact on social surplus relative to the reports of other participants. Because utilities are interdependent, a customer can manipulate their perceived impact on social surplus through the cost function of their matched provider. Hence, customers who pay for the value they bring to their matched providers can always improve their utility by pretending to be of lower QoS (and vice versa). Consequently, the VCG pricing scheme is not incentive compatible in our model. Johnson obtains a similar result for matching markets in the context of position auctions [7].

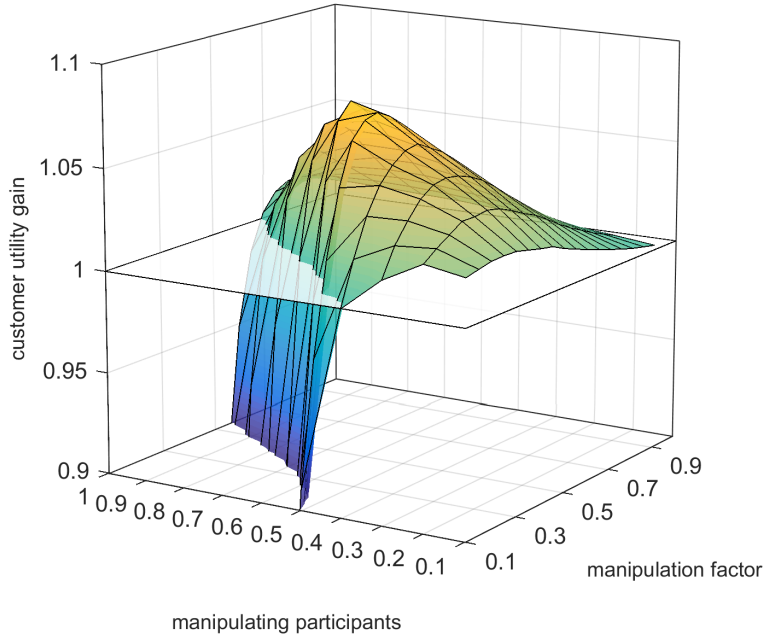


Figure 3: Customer utility gain with normal QoS and $N = 10$ (VCG).

Second, high-quality customers depicted in Figure 1 achieve higher relative utility gains in the mechanism with VCG pricing when QoS is normally distributed (7.3%) as compared to uniform QoS (4.7%). Because the underlying normal distribution assumes a mean of 0.5, there is more mass of QoS realizations in the middle of the unit interval. Therefore, high-quality customers have more leeway to manipulate their QoS as compared to the uniform distribution. However, as the market size increases to $N = 100$, the incentives for customers to manipulate their QoS decrease. In large markets, participants are faced with higher competition, which causes a decrease in their potential utility gain.

Third, we find that by using the VCG pricing scheme, customers have an incentive to manipulate their QoS, while providers do not. Under the mechanism with k -pricing, however, providers can improve their utility by manipulation, while customers cannot. In the mechanism with VCG pricing, customers have higher utilities than providers (cf. Table 1). These utilities denote a measure for the loss of incentive compatibility. Because the utilities of customers exceed those of providers, the mechanism's cost for providing adequate incentives for customers is higher, too. Hence, the mechanism with VCG payments is vulnerable in the face

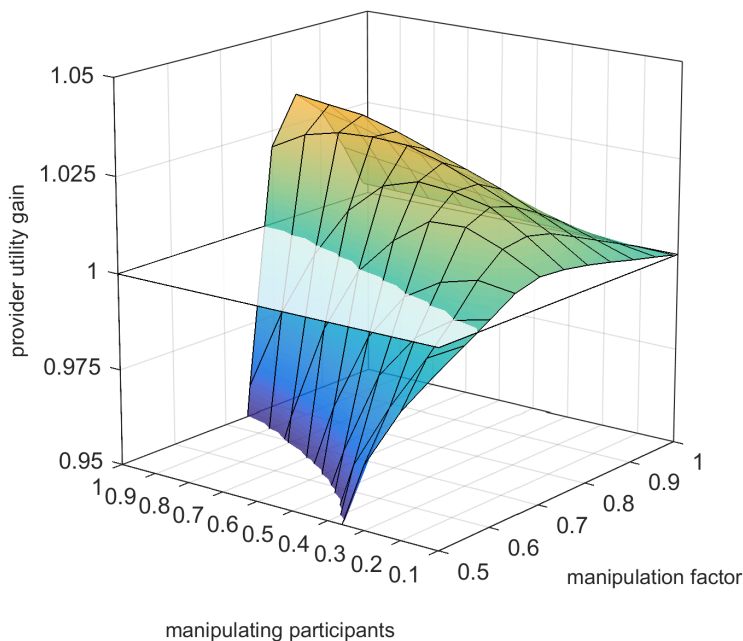


Figure 4: Provider utility gain with uniform QoS and $N = 10$ (k -pricing).

of bid manipulation of customers. On the other hand, providers can increase their utilities by strategic misrepresentation in the mechanism with k -pricing. Here, the payments for all participants is just the middle of each valuation-cost pair. Because the customer's valuation is greater than the provider's cost on average, there is more weight on the valuation when calculating the transaction price. Therefore, if all participants distort their bids using the same manipulation factor, the average transaction price rises. Higher transaction prices, however, entail higher utilities for providers, thus offering the possibility for strategic bid manipulation.

Fourth, the more participants manipulate their bids in the mechanism with k -pricing, the higher is the average utility gain for each provider (cf. Figure 4). Schnizler et al. report on similar findings for bids that follow the Decay distribution [12]. They find that each participant can increase their utility by 26% when all participants manipulate by 6%. In this work, however, the maximum utility gain of providers is only 1.5% at a manipulation factor of 0.75 (i.e., 25%). The reason for this difference is that our model is based on interdependent utilities of customers and providers. As such, the mechanism must consider the fact that a customer's utility is maximized when their desired QoS matches the actual QoS delivered by a

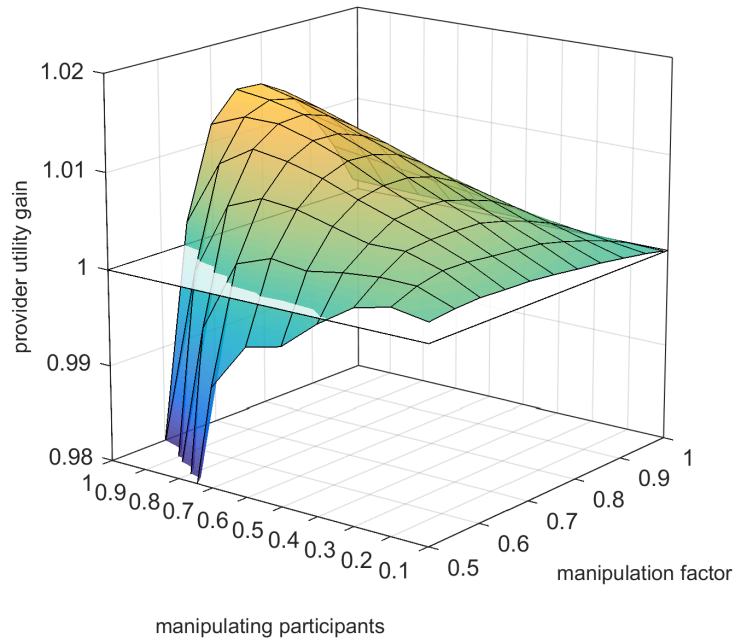


Figure 5: Provider utility gain with normal QoS and $N = 10$ (k -pricing).

provider (and vice versa). Therefore, participants must apply a 25% manipulation to arrive at the relatively low average utility gain of 1.5%.

Future research might be pursued in two directions. First, it would be interesting to study the effects of strategic bid manipulation on potential utility gains from the perspective of a profit-maximizing intermediary. This change would require modifying the allocation rule (6) and revisiting the payment schemes. To what extent the interdependent structure of the participants' utilities affects the strategic bidding behavior in a profit-maximizing implementation remains an open research question. Second, while we assume an equal number of customers and providers in the current model, future research can now analyze the effects of a varying number of customers and providers on the potential utility gain of strategic participants. Further experimentation is required to understand the efficacy of the mechanism for dynamic market settings.

Acknowledgements

This work has been partially supported by the German Federal Ministry of Education and Research under grant 16KIS0083.

References

- [1] AVT Technology Solutions. StreamOne Enterprise Solutions, 2018. URL <https://ses.techdata.com/catalog/>. [Online; accessed 26-January-2018].
- [2] G. I. Doukidis, K. Pramatari, and G. Lekakos. OR and the management of electronic services. *European Journal of Operational Research*, 187(3): 1296–1309, 2008.
- [3] S. Fadaei and M. Bichler. Generalized assignment problem: Truthful mechanism design without money. *Operations Research Letters*, 45:72–76, 2017.
- [4] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel. The cost of a cloud: research problems in data center networks. *ACM SIGCOMM Computer Communication Review*, 39(1):68–73, 2009.
- [5] S. Gupta. On a modification of the VCG mechanism and its optimality. *Operations Research Letters*, 44:415–418, 2016.
- [6] C.-W. Hang and M. P. Singh. From quality to utility: Adaptive service selection framework. In *Proceedings of the 8th International Conference on Service-Oriented Computing (ICSOC 2010)*, pages 456–470, San Francisco, California, USA, 2010.
- [7] T. Johnson. Matching through position auctions. Working Papers 001, Department of Economics, University of Notre Dame, 2011.
- [8] S. Lee. Incentive compatibility of large centralized matching markets. *The Review of Economic Studies*, 84(1):444–463, 2016.
- [9] P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.
- [10] R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:256–281, 1983.

- [11] M. Satterthwaite and S. Williams. *The Double: Auction Market - Institutions, Theories, and Evidence*, chapter The Bayesian Theory of the k-Double Auction, pages 99–123. Addison-Wesley, 1993.
- [12] B. Schnizler, D. Neumann, D. Veit, and C. Weinhardt. Trading grid services - a multi-attribute combinatorial approach. *European Journal of Operational Research*, 187:943–961, 2008.
- [13] R. Shimer and L. Smith. Assortative matching and search. *Econometrica*, 68(2):343–369, 2000.
- [14] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
- [15] T. Widmer and J. Leukel. Efficiency of electronic service allocation with privately known quality. *European Journal of Operational Research*, 255(3): 856–868, 2016.