

Strategic Bidding in Decentralized Collaborative Vehicle Routing

Los, Johan; Schulte, Frederik; Spaan, Matthijs T.J.; Negenborn, Rudy R.

DOI

[10.1007/978-3-031-05359-7_21](https://doi.org/10.1007/978-3-031-05359-7_21)

Publication date

2022

Document Version

Final published version

Published in

Dynamics in Logistics

Citation (APA)

Los, J., Schulte, F., Spaan, M. T. J., & Negenborn, R. R. (2022). Strategic Bidding in Decentralized Collaborative Vehicle Routing. In M. Freitag, A. Kinra, H. Kotzab, & N. Megow (Eds.), *Dynamics in Logistics : Proceedings of the 8th International Conference LDIC 2022, Bremen, Germany* (pp. 261-274). (Lecture Notes in Logistics). Springer. https://doi.org/10.1007/978-3-031-05359-7_21

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Strategic Bidding in Decentralized Collaborative Vehicle Routing

Johan Los¹, Frederik Schulte¹(✉), Matthijs T. J. Spaan²,
and Rudy R. Negenborn¹

¹ Department of Maritime and Transport Technology, Delft University of Technology, Mekelweg 2, 2628 CD Delft, The Netherlands
{F.Schulte,R.R.Negenborn}@tudelft.nl

² Department of Software Technology, Delft University of Technology, Van Mourik Broekmanweg 6, 2628 XE Delft, The Netherlands
M.T.J.Spaan@tudelft.nl

Abstract. Collaboration in transportation is important to reduce costs and emissions, but carriers may have incentives to bid strategically in decentralized auction systems. We investigate what the effect of the auction strategy is on the possible cheating benefits in a dynamic context, such that we can recommend a method with lower chances for carriers to cheat. We consider both a first-price auction system and a second-price auction scheme. Contrary to what was expected, a second-price auction scheme gives more room for successful strategic behaviour, while it also results in more rejected orders. A first-price auction scheme might be useful in practice if the profit shares that are allocated to the winner of an auction are selected carefully.

Keywords: Collaborative vehicle routing · Decentralized collaborations · Strategic behaviour · Auctions · Multi-agent system

1 Introduction

Inner cities seriously suffer from pollution problems, with traffic being one of the major causes. In particular, the pollution by freight transportation could be significantly decreased, since operations are often inefficient: carriers generally do not cooperate, although the total vehicle mileage could be reduced if different carriers combine similar tasks into one route. It is hence desirable that carrier cooperation becomes common practice in freight transportation. Preliminary studies have shown that reductions of 20–30% in vehicle mileage can be obtained if only 2 or 3 carriers cooperate (Gansterer and Hartl 2018b).

In the current article, we consider large-scale collaboration in dynamic pickup and delivery problems with time windows and capacity constraints. For this problem type, centralized collaboration approaches are not applicable due to scale problems and privacy and autonomy concerns. Thus, decentralized auction approaches have been developed (Máhr et al. 2010; Mes et al. 2013; Los et al. 2020b,

2020c). In such cases with large amounts of collaborating carriers, the positive collaboration effects could increase even more: Los et al. (2020a, 2022) found cost reductions of up to almost 80% with cooperation between 1000 carriers.

Unfortunately, carriers may have incentives to cheat a decentralized cooperation system. It is often assumed that (estimates of the) real valuations for orders are reported within the auctions. In practice, however, carriers and shippers might bid strategically and try to increase their individual profits at the cost of the others. In this context, Los et al. (2022) have shown that carriers can successfully outplay other carriers, resulting in a decrease of the total cooperation benefits.

Nonetheless, strategic behaviour is not straightforward: Los et al. (2022) show that it is highly dependent on the auction conditions whether false bidding pays off. Also Gansterer and Hartl (2018a) present a small computational example with a central combinatorial auction in which the cheating carrier always incurs a loss compared to truthful bidding. They emphasize that no general conclusions can be drawn from the example, but suggest that it might be rather difficult in practice to find a profitable cheating strategy.

Although not much is known about strategic behavior in auction-based transportation collaborations (Gansterer and Hartl 2020), it is important to investigate whether decentralized auction systems are incentive compatible, that is, can withstand strategic behaviour, before they can be applied in practice. Hence, the goal of this paper is to investigate what the effect is of the auction strategy on the benefits of cheating, such that we can find a method that has a less likely chance for parties to cheat.

To compare different auction strategies, we adopt and adapt the auction system of Los et al. (2022). We shortly describe the method here, to be able to highlight the two points of focus of this paper. In the decentralized auction approach, each order is proposed in a reverse first-price auction immediately after its release, and all carriers can bid on the order. It is assumed that they bid exactly their marginal costs for the order, that is, the extra travel costs that would result from including the order at the most efficient position in one of their routes. If the lowest bid is below the shipper's reservation price, the lowest bidding carrier is compensated by the amount of its bid, and commits to transport the order. Since the world is dynamic, better allocations might appear later on. Hence, frequent reauctions are used: an order can be transferred to another carrier if the new bid is lower than the actual costs for the current carrier. Los et al. (2022) show that carriers can sometimes successfully outplay other carriers by asking for a lower value than their true marginal costs. Although they will incur a small loss by doing so, they will often be compensated: they can either directly get a share of the auction profits, or they might be compensated later on if they outsource the task again.

In this article, we consider two ways to influence the direct compensation. First, Los et al. (2022) conjecture that benefits from strategic bidding might largely depend on the share of the profit generated by a successful auction that a winning carrier obtains. We investigate this hypothesis by comparing different profit allocation schemes in which we vary the direct profits that are assigned to the winning carrier.

Next, we investigate whether a solution to the problem of direct compensation can lie in applying a second-price auction scheme. Within a second-price auction, the lowest bidding carrier still wins the order, but is compensated with the (higher) amount of the second best bid. Under certain conditions, participants in reverse second-price auctions do not have any incentives to deviate from their true value (Vickrey 1961): the winning carrier either would have won the auction anyhow (if its true value is below the second price), or makes a loss (if its true value is above the second price). Although this incentive compatibility property holds for auctions with a single indivisible good, it is not guaranteed for our scenario where we have multiple dependent auctions. Carriers can still be compensated indirectly, either by reselling the orders, or by obtaining other orders that have positive interaction effects with the orders already in their routes. Still, to find a strategic policy seems more difficult with second-price auctions than with first-price auctions because the direct compensation is only dependent on the price of the second best bid. Hence, we extend the auction approach with a second-price auction scheme and examine how it performs under strategic behaviour.

2 Related Work

A mechanism for exchanging orders between carriers should not only be robust with respect to strategic behaviour, but requires some other qualities as well. Ideally, it has the following four properties from standard auction theory:

- **Efficiency:** The mechanism leads to a routing solution that cannot be further improved.
- **Individual rationality:** For each carrier, participating in the collaboration does not result in worse results than not participating.
- **Incentive compatibility:** Carriers do not have incentives to report other values than their true valuations.
- **Budget balance:** No extra money from outside the system is needed.

It is, however, not possible to obtain all four properties simultaneously in standard environments (Myerson and Satterthwaite 1983). A couple of studies investigate trade-offs of these properties in static carrier collaboration situations: Xu et al. (2017) propose a bundle double auction for a problem where each carrier can exchange only one full truckload, and show that their method realizes budget balance, incentive compatibility and individual rationality, but only asymptotical efficiency. They extend the model to the exchange of multiple truckloads and propose two extended mechanisms that either are not incentive compatible for our-sourcing carriers or not asymptotically efficient anymore. Gansterer et al. (2019) analyze combinatorial auctions where carriers can act as buyers and sellers at the same time. The marginal costs for insourcing an order then do not only depend on their current orders, but also on which orders they will outsource, making the problem more complex. The authors compare a Vickrey-Clarke-Groves mechanism and

a team bidder approach: both are incentive compatible and efficient, but the properties of individual rationality and budget balance are violated. In an experimental study, they show the trade-offs of both approaches.

In addition to the possible interactions between insourced and outsourced orders, our case is even more complicated because we consider a dynamic environment where future orders might influence the value of current orders. As far as we know, only Figliozzi (2006) studies incentive compatible mechanisms for dynamic carrier collaboration. He uses a second-price auction scheme for each newly arrived order and claims that the approach is incentive compatible, individually rational, and budget balanced, but not fully efficient. Efficiency is hindered by possible future orders (as is common for dynamic systems) and also by the fact that no reassignment is made if the current costs for the owner of the order are lower than the value of the second bid, but higher than the value of the first bid. (In that case, the owner would make a loss by paying the second price, but a better allocation could be made.)

The claim for incentive compatibility, however, can be opposed. It is argued that a carrier will not place a bid lower than its true marginal costs for transporting the order, since it will make a loss if also the second price is below its true marginal costs. This indeed holds under the assumptions that the marginal costs “include all relevant costs (including opportunity costs) associated with servicing (or not servicing) an additional shipment or shipments” and that “all participating carriers compute these costs accurately” (Figliozzi 2006, p. 35). However, these assumptions are too strong: it is impossible to give a certainly accurate prediction of opportunity costs in dynamic systems, simply because it is not known what orders might appear later on, and even more because it is not known whether these can be lucratively obtained or outsourced via the auction system. (Furthermore, in large-scale systems, an exact computation of the insertion costs may take too much time to be practical.) Hence, carriers may strategically bid lower values to obtain orders at a loss if they expect that advantageous interaction effects can occur later on.

In this article, we experimentally investigate to what extent the first-price auction system developed by Los et al. (2022) and a second-price auction system (which is comparable to that of Figliozzi (2006)) are incentive compatible in the context of dynamic pickup and delivery problems.

3 Auction Approaches

We build upon the auction approach developed by Los et al. (2022), in which a platform auctioneer offers bundles of orders. First, we summarize their first-price auction approach; for details, both regarding the problem definition and regarding the method, we refer to Los et al. (2022). Next, we describe how the approach can be transformed into a second-price auction system.

3.1 First-Price Auctions

The auction for a bundle of orders B at time t is as follows:

1. **Requesting transportation.** The auctioneer requests all active carriers to bid for the transportation of bundle B .
2. **Computing marginal costs.** Each carrier c computes its individual marginal costs $MC_c^t(B)$ for bundle B at time t .
3. **Bidding.** All carriers place a bid for the bundle, based on $MC_c^t(B)$.
4. **Comparing.** The auctioneer compares the lowest bid b_0 with the current costs $CC^t(B)$ for the bundle. The current costs consist of the sum of the marginal costs for already assigned orders and the reservation prices for yet unassigned orders, which are requested from the involved shippers and carriers.
5. **Updating contracts.** The bundle is exchanged if and only if b_0 is lower than the current costs. The auctioneer gets in total $CC^t(B)$ from the outsourcing shippers and carriers and pays the winning carrier an amount of b_0 . The gain of $CC^t(B) - b_0$ is shared among the participants as cooperation incentive. Los et al. (2022) use the following parameters:
 - **Winner gain share (WGS):** This parameter defines what fraction of the gain $CC^t(B) - b_0$ is allocated to the winning carrier.
 - **Contracted gain share (CGS):** This parameter defines the total fraction of the gain $CC^t(B) - b_0$ that is allocated to the currently contracted carrier(s) and/or shipper(s) for the orders within B . Each of them gets an equal share.

The platform keeps the remaining gains if $WGS + CGS < 1$.

To model strategic bidding in the first-price auction system, we vary the value of a bid in step 3 of the procedure. Formally, the bid of a strategic carrier c bidding for bundle B at time t will be $\sigma_c MC_c^t(B)$, where σ_c is a parameter representing the degree of strategic bidding for carrier c .

3.2 Second-Price Auctions

We extend the approach to a second-price auction. Instead of getting the value b_0 (and possibly an extra gain, dependent on the profit distribution function), the winning carrier gets the amount of the second-lowest bid b_1 . The use of second price auctions raises a new problem. The amount of the second price needs to be paid by someone. In a budget balanced setting, still the shippers or already contracted carriers must pay this price. The second price, however, is more likely to be higher than their current costs than the first price is. This might result in less (re)allocations, and hence a worse final solution than with first-price auctions. A solution could lie in the bundling approach proposed by Los et al. (2020a, 2022). If bundles of orders from different owners are offered, the interaction advantages of the orders might result in lower bids by the carriers, while the separate current costs are not influenced by interaction effects. Hence, paying the second price could be less problematic if the platform generates bundles of orders from different owners.

Still, the risk that an auction does not succeed due to false current costs is higher than with first-price auctions. To prevent the current owners of the orders from reporting too low current costs (which is the value they need to pay), we let the auctioneer ask them a certain amount such that the second price can be paid to the winning carrier. The current owners only have to accept or refuse the proposed price from the auctioneer. Thus, we propose the following second-price auction procedure for a bundle of orders B at time t :

1. **Requesting transportation:** The auctioneer requests all active carriers to bid for the transport of bundle B .
2. **Computing marginal costs:** Each carrier c computes its individual marginal costs $MC_c^t(B)$ for bundle B at time t .
3. **Bidding.** All carriers place a bid with value $\sigma_c MC_c^t(B)$ for the bundle, where σ_c again represents the degree of strategic bidding.
4. **Comparing:** The auctioneer compares the received bids; let b_0 be the lowest bid provided by carrier c_0 and b_1 the second lowest bid.
5. **Proposing prices:** The auctioneer needs to pay b_1 to c_0 for a (re)allocation, and hence must make sure to get at least b_1 from the current owner(s) of the orders in B . If a lower amount is gathered, the auctioneer will make a loss, and has no incentive to make a reallocation. All amounts above b_1 can be kept as profit for the auctioneer. Thus, the auctioneer proposes a price a_c for all carriers $c \in C_B^t$, and a price a_s for all shippers $s \in S_B^t$ such that $\sum_{c \in C_B^t} a_c + \sum_{s \in S_B^t} a_s \geq b_1$, where C_B^t represents the set of all carriers contracted at time t for at least one order in B and S_B^t represents the set of shippers having an order in B that is yet unassigned at time t . Prices could be determined in different ways. In this article, we use a straightforward approach that divides b_1 proportionally to the distance between pickup and delivery locations of the orders in B , and adds a small profit factor to it, defined as follows:
 - **Platform gain share (PGS):** This parameter defines what fraction of the second bid b_1 is additionally requested from the current owners of the orders as a gain for the platform.

Let O_c^t denote the set of orders that carrier c has in its route plans at time t , let O_s^t represent the set of yet unassigned orders of shipper s at time t , and let $t_{p_o d_o}$ stand for the travel time between the pickup and delivery location of order o . The requested prices are then given by

$$a_c = \frac{\sum_{o \in B \cap O_c^t} t_{p_o d_o}}{\sum_{o \in B} t_{p_o d_o}} (1 + \text{PGS}) b_1 \quad \forall c \in C_B^t \quad (1)$$

and

$$a_s = \frac{\sum_{o \in B \cap O_s^t} t_{p_o d_o}}{\sum_{o \in B} t_{p_o d_o}} (1 + \text{PGS}) b_1 \quad \forall s \in S_B^t \quad (2)$$

such that $\sum_{c \in C_B^t} a_c + \sum_{s \in S_B^t} a_s = (1 + \text{PGS}) b_1$.

When the auctioneer has proposed the prices, the current owners of the orders can check whether the requested prices are less than or equal to their current costs or reservation prices. If so, they will accept the proposed prices.

6. **Updating contracts:** If all current owners accept the proposed prices, the bid is accepted. The platform informs all involved shippers and carriers, who update their contracts and routing plans. The auctioneer receives the payments from the outsourcing shippers and carriers as proposed (i.e., it receives $(1 + \text{PGS})b_1$ in total), and pays b_1 to the winning carrier c_0 . The remaining gain of $\text{PGS}b_1$ is kept by the auctioneer. If one of the current owners does not accept to outsource its orders at the proposed price, no (re)allocations and no payments take place.

The approach guarantees that the second price is paid to the winning carrier if a (re)allocation takes place, and that the auctioneer does not incur a loss (if $\text{PGS} \geq 0$). The drawback, however, is that current owners need to accept the prices that are proposed by the auctioneer to have a successful (re)allocation. This becomes less likely with higher PGS values.

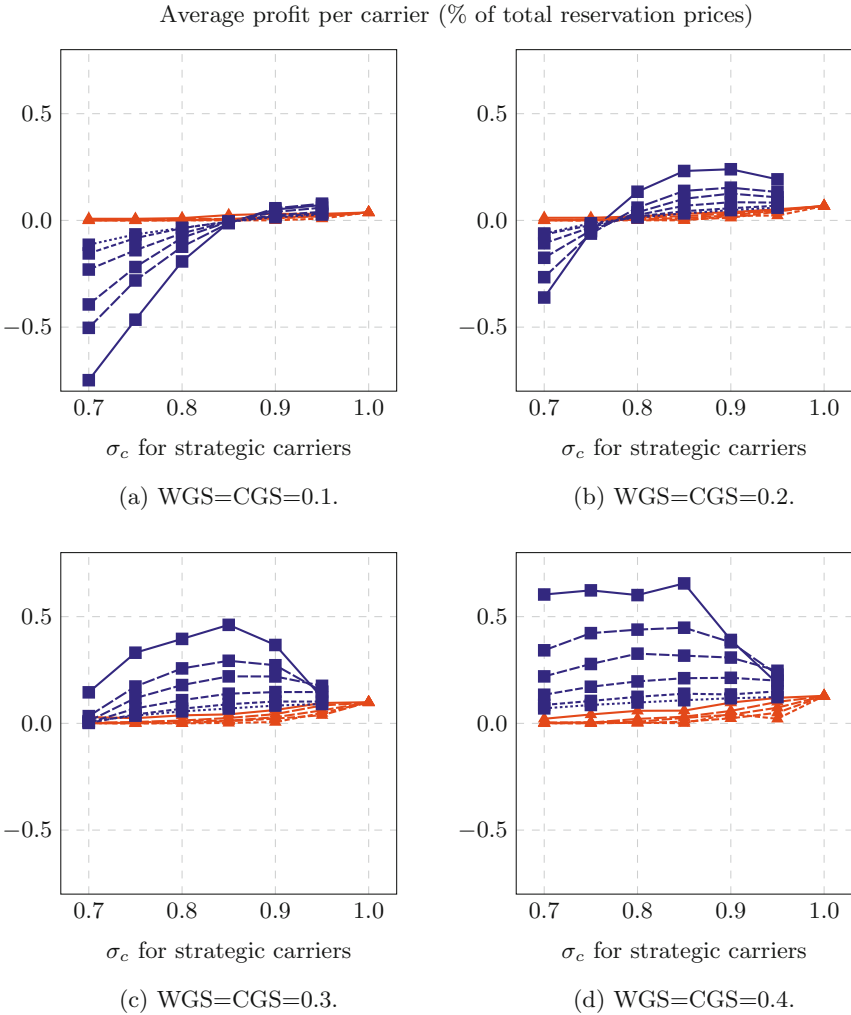
4 Computational Study

In this section, we empirically test what the influence of strategic bidding is within the proposed decentralized auction approaches. Throughout the computational study, we use a real-world data set from a Dutch transportation platform company, and generate instances of 2000 orders each. To prevent any bias from unprofitable initial contracts, we use problem instances without initial assignment. Per instance, there are 250 carriers with 1–3 vehicles each. Restricted availability time windows are applied in one third of the cases. Further instance characteristics are the same as described by Los et al. (2022).

4.1 Strategic Behaviour in a First-Price Auction System

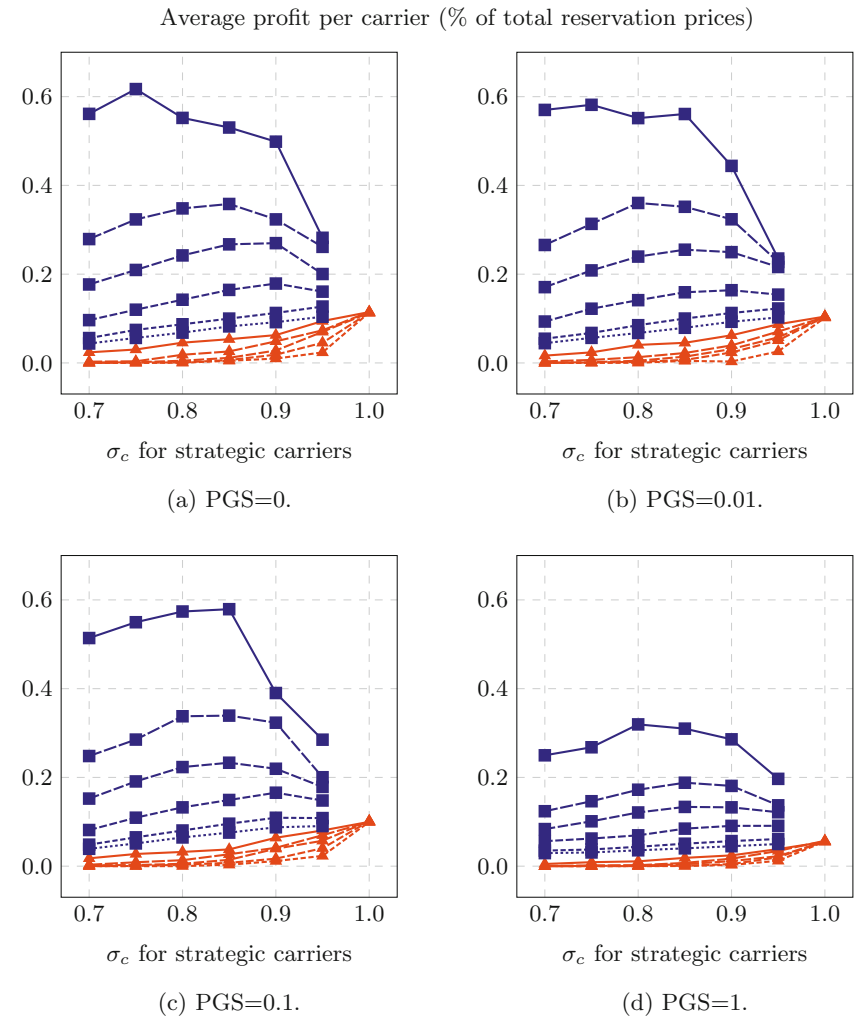
As hypothesized by Los et al. (2022), lower values of σ_c are expected to be beneficial for individual carriers when higher values for WGS are used: if the system assigns large shares of the gains to the winning carriers, cheating might appear too easy. We tested this hypothesis with different percentages of carriers (10%, 20%, 30%, 50%, 80%, or 100%) that place strategic bids ($\sigma_c \in \{0.7, 0.75, 0.8, 0.85, 0.9, 0.95\}$), and consider four different values for winner gain share ($\text{WGS} \in \{0.1, 0.2, 0.3, 0.4\}$). The results (average profits as a percentage of the sum of the reservation prices for the transported orders) are given in Fig. 1, where the average profit with only truthful carriers is given as a reference at $\sigma_c = 1$.

While the turning point below which strategic bidding does not pay off is $\sigma_c = 0.9$ for $\text{WGS} = 0.1$, this decreases to $\sigma_c = 0.8$ for $\text{WGS} = 0.2$. For $\text{WGS} = 0.3$ and $\text{WGS} = 0.4$, even the lowest tested value of $\sigma_c = 0.7$ is still beneficial for strategic carriers. Thus, indeed, the higher the value of WGS, the easier it is for carriers to find a beneficial strategic bidding policy.



	Truthful carriers	Strategic carriers
10% strategically bidding carriers		
20% strategically bidding carriers		
30% strategically bidding carriers		
50% strategically bidding carriers		
80% strategically bidding carriers		
100% strategically bidding carriers		

Fig. 1. Average carrier profits if part of the carriers bid a fraction of their real (estimated) insertion costs in the first-price auction system, for increasing values of WGS and CGS.



	Truthful carriers	Strategic carriers
10% strategically bidding carriers		
20% strategically bidding carriers		
30% strategically bidding carriers		
50% strategically bidding carriers		
80% strategically bidding carriers		
100% strategically bidding carriers		

Fig. 2. Average carrier profits if part of the carriers bid a fraction of their real (estimated) insertion costs in the second-price auction system, for different values of PGS.

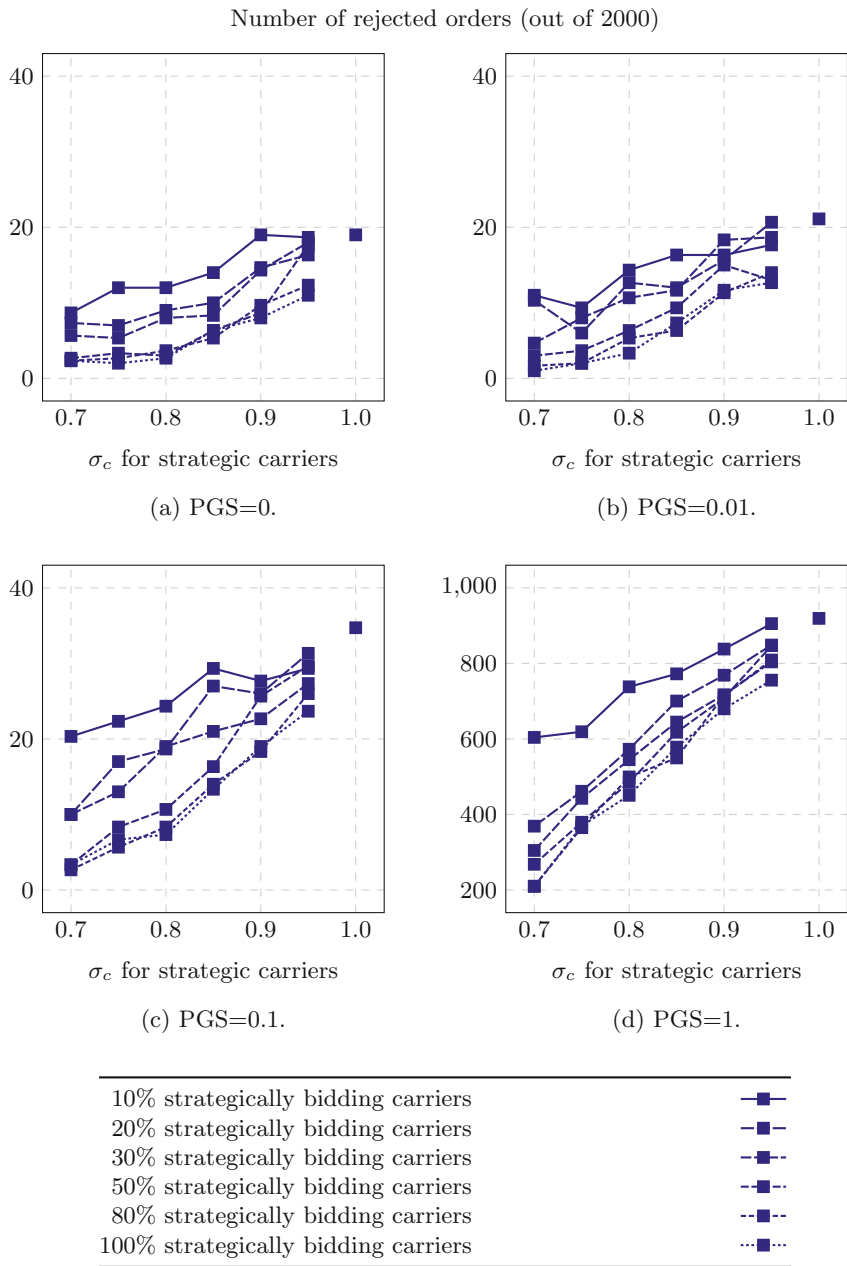


Fig. 3. Average number of rejected orders if part of the carriers bid a fraction of their real (estimated) insertion costs in the second-price auction system, for different values of PGS. (Note that the scale of the y-axis is different for Fig. 3d.)

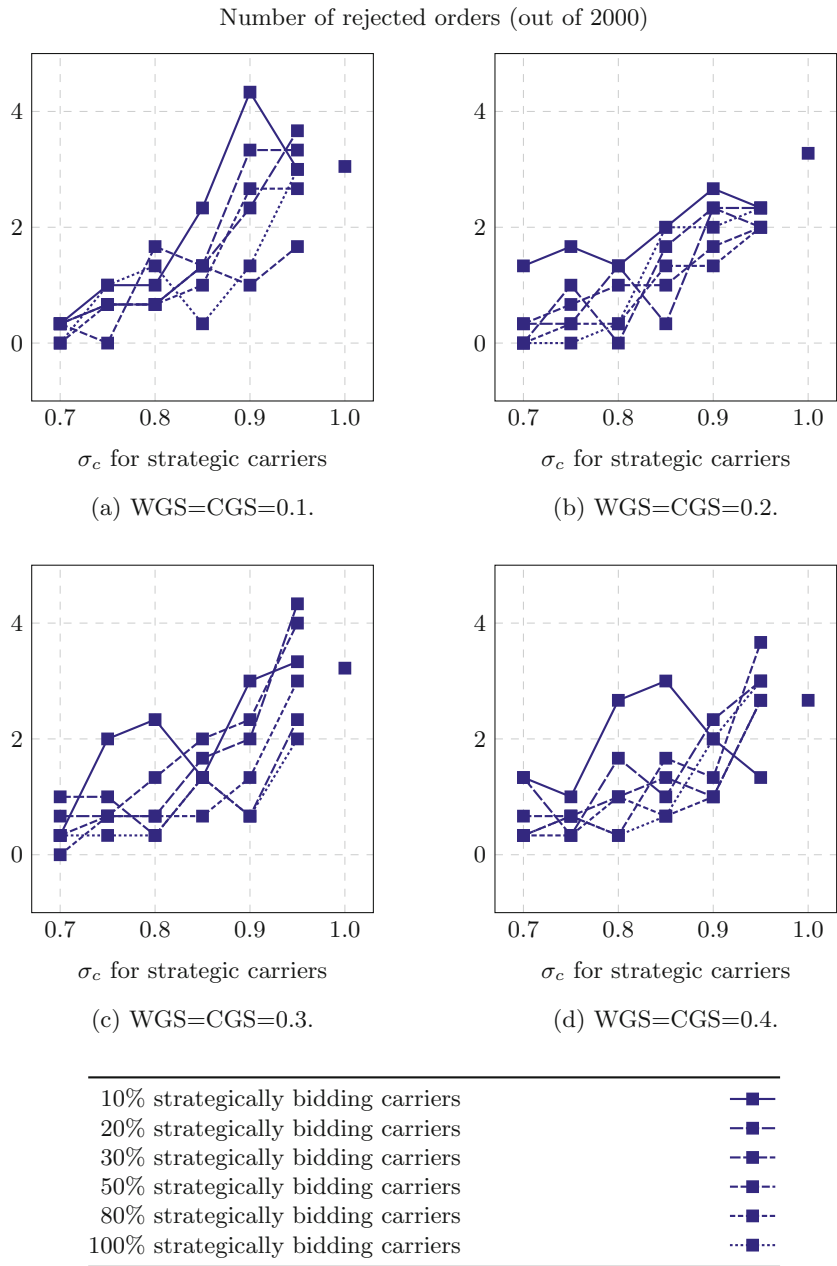


Fig. 4. Average number of rejected orders if part of the carriers bid a fraction of their real (estimated) insertion costs in the first-price auction system, for different values of WGS and CGS.

4.2 Strategic Behaviour in a Second-Price Auction System

We now investigate whether a second-price auction system in a dynamic world can reduce the motivation to bid strategically. To remove any interference between the bid value that a carrier c submits for a bundle B and the expected price a_c that the platform proposes to this carrier if c already owns any order $o \in B$, we restrict our experiments in such a way that bids are only made for bundles that do not contain any currently owned orders. We run the algorithm with different percentages of carriers (10%, 20%, 30%, 50%, 80%, or 100%) that place strategic bids ($\sigma_c \in \{0.7, 0.75, 0.8, 0.85, 0.9, 0.95\}$), and consider four different values for platform gain share ($\text{PGS} \in \{0, 0.01, 0.1, 1\}$). For $\text{PGS} = 0$, the auctioneer asks in each auction round exactly b_1 in total from the current owners, and hence, makes no profit itself. In the extreme case of $\text{PGS} = 1$, on the other hand, the auctioneer asks $2b_1$ from the current owners, and tries to make a profit of b_1 itself each auction.

The average profits for the carriers that bid strategically and for the carriers that bid truthfully are given in Fig. 2. (Again, as a reference, the average profit with only truthful carriers is given at $\sigma_c = 1$.)

Strikingly, strategic bidding always results in higher profits than true bidding, regardless of the number of strategic carriers, the value of σ_c , or the value of PGS. The profits of strategic carriers, however, highly depend on the total number of strategic carriers within the system. If 80% or 100% of the carriers act strategically, their profits are easily becoming lower than the profits in a scenario with only truthful carriers, leading to a kind of prisoner's dilemma: irrespective of what the others do, strategic bidding results in higher individual profits than truthful bidding, but carriers are better off when they all bid truthfully than when they all bid strategically.

In Fig. 3, we show the corresponding average numbers of rejected orders. With more strategically bidding carriers and lower σ_c values, the number of rejected orders decreases, as expected, since the value of b_1 is likely to get lower. For increasing values of PGS, the number of rejected orders increases, with almost half of the orders rejected under some conditions for $\text{PGS} = 1$. This can be explained by the (too) high prices that the auctioneer asks from the current owners of the orders. As a reference, we show the numbers of rejected orders within the first-price auction system in Fig. 4. The number of rejected orders is lower in the first-price system than in the second-price system and also not dependent on the WGS and CGS parameters, while it is heavily dependent on the PGS parameter in the second-price system.

5 Conclusions

We analyzed the potential incentives for carriers to bid strategically in an auction system that they can use to jointly solve their dynamic pickup and delivery problems. We considered the influence of different auction strategies on the individual benefits of cheating, to be able to recommend a method that has a less likely chance for parties to cheat.

Whether strategic bidding can pay off in a first-price auction setting turns out to depend highly on the share of the gains per auction that is attributed to the winning carrier:

- If this share is relatively low, carriers can benefit if they slightly lower their bids. If they make their bids too low, however, they easily will make a loss. The exact value of the turning point will not be clear beforehand, making it difficult for carriers to cheat. The drawback of a system with a low gain share for carriers is that they have little incentive to participate in the system.
- If the gain share for a winning carrier is relatively high, carriers might be interested in participating in the cooperation system. The problem is that it then will be easy for them to cheat the system: they can bid lower prices to get more orders and will be compensated for their too low bids. At the same time, the total routing solution will become worse, since the orders will often not be assigned to the carriers that can perform them at least costs.

It might thus be possible for a platform provider to use a first-price auction system, but the procedure details must be selected carefully to prevent strategic bidding. With too high gain shares for carriers, strategic bidding easily pays off for them, while they might have no incentive to participate with too low gain shares.

Motivated by the strategy-proofness of second-price auctions for single indivisible items, we experimentally tested whether second-price auctions could be applied successfully in our multi-item dynamic context. The hypotheses were that a second-price auction could reduce the individual profit of strategic bidding and that auctioning bundles of orders might solve the budget balance problem as well in this case. It turned out, however, that strategic bidding always pays off for carriers within the proposed system: average profits for strategic carriers are higher than the related profits for truthful carriers. Apparently, the long-term advantages of having a larger set of orders outweigh the lower compensations when acquiring them. Carriers will only have a disadvantage if too many other carriers also cheat: when 80–100% of the carriers bid strategically, the average profits are lower than in a completely truthful setting.

Future research should investigate different payment methods to compensate carriers for exchanging orders in a decentralized auction system. Within our second-price auction approach, the auctioneer always proposed a fixed proportional price for outsourcing carriers and shippers. Different ways to get the required amount from the current owners might improve the results. Even the auctioneer could make a loss in some auctions where it is difficult to gather the amount of the second price, if this can be compensated in other auction rounds. Furthermore, different ways of assigning the profits from a single auction round might contribute to a system in which it is hard for carriers to benefit from strategic behaviour.

Acknowledgements. This research is supported by the project “Dynamic Fleet Management (P14-18 - project 3)” (project 14894) of the Netherlands Organization for Scientific Research (NWO), domain Applied and Engineering Sciences (TTW).

References

- Figliozi, M.A.: Analysis and evaluation of incentive-compatible dynamic mechanisms for carrier collaboration. *Transp. Res. Rec.* **1966**(1), 34–40 (2006)
- Gansterer, M., Hartl, R.F.: Centralized bundle generation in auction-based collaborative transportation. *OR Spectr.* **40**, 613–635 (2018a)
- Gansterer, M., Hartl, R.F.: Collaborative vehicle routing: a survey. *Eur. J. Oper. Res.* **268**(1), 1–12 (2018b)
- Gansterer, M., Hartl, R.F.: Shared resources in collaborative vehicle routing. *TOP* **28**, 1–20 (2020)
- Gansterer, M., Hartl, R.F., Vetschera, R.: The cost of incentive compatibility in auction-based mechanisms for carrier collaboration. *Networks* **73**(4), 490–514 (2019)
- Los, J., Schulte, F., Gansterer, M., Hartl, R.F., Spaan, M.T.J., Negenborn, R.R.: Decentralized combinatorial auctions for dynamic and large-scale collaborative vehicle routing. In: *Proceedings of the 11th International Conference on Computational Logistics, ICCL 2020, Enschede, The Netherlands*, pp. 215–230 (2020a)
- Los, J., Schulte, F., Gansterer, M., Hartl, R.F., Spaan, M.T.J., Negenborn, R.R.: Large-scale collaborative vehicle routing. *Ann. Oper. Res.* (2022). <https://doi.org/10.1007/s10479-021-04504-3>
- Los, J., Schulte, F., Spaan, M.T.J., Negenborn, R.R.: Collaborative vehicle routing when agents have mixed information sharing attitudes. *Transp. Res. Procedia* **44**, 94–101 (2020b)
- Los, J., Schulte, F., Spaan, M.T.J., Negenborn, R.R.: The value of information sharing for platform-based collaborative vehicle routing. *Transp. Res. Part E Logist. Transp. Rev.* **141**, 102011 (2020c)
- Máhr, T., Srour, J., de Weerd, M., Zuidwijk, R.: Can agents measure up? A comparative study of an agent-based and on-line optimization approach for a drayage problem with uncertainty. *Transp. Res. Part C Emerg. Technol.* **18**(1), 99–119 (2010)
- Mes, M., van der Heijden, M., Schuur, P.: Interaction between intelligent agent strategies for real-time transportation planning. *CEJOR* **21**, 337–358 (2013)
- Myerson, R.B., Satterthwaite, M.A.: Efficient mechanisms for bilateral trading. *J. Econ. Theor.* **29**(2), 265–281 (1983)
- Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *J. Financ.* **16**(1), 8–37 (1961)
- Xu, S.X., Huang, G.Q., Cheng, M.: Truthful, budget-balanced bundle double auctions for carrier collaboration. *Transp. Sci.* **51**(4), 1365–1386 (2017)