Submitted to *Management Science*
manuscript 0001

Coordination is Hard: Electronic Market Mechanisms For Increased Efficiency in Transportation Logistics

Paul Karaenke, Martin Bichler
Department of Informatics, 85748 Garching, Technical University of Munich, Germany, {karaenke, bichler}@in.tum.de

Stefan Minner
School of Management, 80333 Munich, Technical University of Munich, Germany, stefan.minner@tum.de

The lack of coordination among carriers leads to substantial inefficiencies in logistics. Such coordination problems constitute fundamental problems in supply chain management for their computational and strategic complexity. We consider the problem of slot booking by independent carriers at an operator of several warehouses, and investigate recent developments in the design of electronic market mechanisms promising to address both types of complexity. Relax-and-round mechanisms describe a class of approximation mechanisms that is truthful in expectation and runs in polynomial time. While the solution quality of these mechanisms is low, we introduce a variant able to solve real-world problem sizes with high solution quality while still being incentive-compatible. We compare these mechanisms to core-selecting auctions which are not incentive-compatible, but provide stable outcomes with respect to the bids. In addition to a theoretical analysis we report results from extensive numerical experiments based on field data. The experimental results yield a clear ranking of the mechanisms in terms of waiting time reductions and computation times.

Key words: Electronic markets and auctions, supply chain and logistics, IT-enabled supply chains

---

Version May 26, 2018

1. Introduction

Many information systems nowadays are designed to coordinate activities or allocate scarce resources. The design of respective information systems has a number of challenges because incentives of the participants need to be considered, but also computational problems play a role. This has led to a fruitful line of research on the design of electronic markets (Banker and Kauffman 2004, Bichler et al. 2010, Adomavicius et al. 2012). We focus on a specific problem in retail logistics, which is representative for coordination problems as they can be widely found in business practice. The problem is well motivated and has substantial impact on economic efficiency: long waiting times at loading docks of retail warehouses and unequal workload over the day. According to a survey among more than 500 transportation companies in Germany, 18% of them have an average waiting time of more than two hours and 51% have an average waiting time between one to two hours at each warehouse (Bundesverband Güterkraftverkehr Logistik und Entsorgung e.V. 2013). Such waiting times are a significant problem for carriers and warehouse operators. The German Federal Office for Transportation reports that the main reasons for waiting times include shortages of staff and infrastructure, and in particular the uncoordinated arrivals of trucks (Bundesamt für Güterverkehr 2011). This lack of coordination causes substantial inefficiencies in retail transportation logistics.

Different solutions have been proposed to address the problem including time-slot management systems, which are based on first-come first-served (FCFS) assignments of carriers to time slots at a warehouse (Elbert et al. 2016). Time-slot management allows warehouse managers to allocate sufficient capacity such that carriers with reservations do not experience significant waiting times.

Fundamentally, the allocation of carriers to time slots at different warehouses is a computationally hard coordination problem among carriers and warehouses. The carriers decentrally solve vehicle routing problems, which determine when a truck arrives at a certain warehouse. They want to have reservations for these time slots. The warehouses face capacity planning problems for their loading docks because they don't know when carriers arrive. If all information about supplier preferences for acceptable routes and warehouse capacities was available, then a central coordinator or clearing house (potentially organized via a booking platform) could select routes and allocate time slots to carriers such that waiting times are minimized. However, this information is not available in FCFS time-slot management systems, where the assignment of time slots to carriers is based on the sequence of reservations, not on a global optimization. In addition, carriers need to pay a reservation fee no matter if there is excess demand for a time slot or not.

1.1. Coordination via market design

We explore new types of market mechanisms able to address the computational and strategic problems that need to be solved in order to coordinate carriers optimally. We show that such mechanisms can decrease waiting times significantly. The goal of these auctions is to maximize allocative efficiency, i.e., to find the optimal allocation of time slots to carriers given their preferences, not to maximize revenue of the coordinator. Payments are only used to set incentives for truthful bidding. Ideally, there are no incentives for strategic manipulation such that bidders just have to reveal their valuations or opportunity costs for alternative routes and corresponding time slots truthfully.

If bidders had a dominant strategy to reveal their valuation for different routes (and corresponding time slots) truthfully, then this would not only decrease their bid preparation costs, it would also increase efficiency of the mechanism. If truthfulness is a concern, there is a well-known recipe in the mechanism design literature: the VCG mechanism is the unique social-welfare maximizing auction mechanisms, where truthful bidding is a dominant strategy equilibrium. Unfortunately, the VCG mechanism only exhibits dominant strategies if the optimal allocations can be computed exactly. If this is not the case, there might be problems computing individually rational VCG payments and possibilities for manipulation. For example, consider a single-object VCG auction where the second highest bid is winning (due to the approximation). In this case, truthful bidding is not optimal for the bidder with the highest valuation. If bidders were truthful, the payment of the winner would be higher than his valuation.

For our time slot assignment problem, the coordinator has to solve an NP-hard allocation problem as we will discuss in Section 2 and cannot expect to solve the underlying allocation problems to optimality with realistic input sizes. Therefore, we draw on two recent developments in the literature: *Core pricing with near-optimal allocations* and *truthful approximation mechanisms.*

While the VCG mechanism is well-known, *core pricing* has become popular only in the recent years, and it is nowadays used world-wide in spectrum auctions to address problems with a VCG payment rule (Ausubel and Milgrom 2002, Day and Raghavan 2007, Day and Milgrom 2008). In particular, with a VCG payment rule, the winners might have to pay less than what a losing coalition of bidders was willing to pay, which can be considered unfair by participants. Such an allocation is not in the core, i.e. it is not stable because sub-coalitions of bidders might find a better allocation with the auctioneer. VCG and core payment rules rely on an exact solution to the allocation problem. As is often the case for NP-hard problems, solutions with a small gap of the mixed-integer program can typically be found in minutes, while the exact solution of the same problem can take days or is even intractable. Goetzendorff et al. (2015) recently introduced two algorithms to approximate the VCG and core payments with such near-optimal solutions, and also make sure that the payments do not exceed the valuations.

On the one hand, VCG payment rules applied to near-optimal allocations are not truthful, as the example of the single-object VCG auction showed. On the other hand, profitable manipulation would require a lot of information in combinatorial auctions with many packages of items that a bidder can choose from. Even if the auctioneer does not select the optimal allocation, the "second-price" logic of a VCG and core payment rules might well set sufficient incentives for bidders to bid truthful in larger VCG or core-selecting auctions with near-optimal solutions. Strategic manipulation is very hard in such environments to say the least.

Algorithmic mechanism design is another stream of literature analyzing *approximation mechanisms*, which give up on social welfare maximization, but which are truthful. *Truthfulness in expectation* is a solution concept for randomized mechanisms, where a bidder maximizes his expected payoff with truthful bidding. This means such mechanisms are incentive-compatible. In a seminal paper, Lavi and Swamy (2011) introduce a general framework which is able to transform any approximation algorithm for a packing problem into a mechanism that is truthful in expectation and runs in polynomial time. This is remarkable, as earlier contributions provided approximation mechanisms for specific problems, but no black-box mechanism. More recently, other black-box mechanisms for restricted valuations (e.g., bidders only interested in one package or bidders with submodular valuations) have been proposed, but the approach by Lavi and Swamy (2011) stands out as it is applicable to any packing problem with an α-approximation algorithm. We refer to their approach as the *RaR* (relax-and-round) *framework*: Instead of solving the underlying packing problem directly, its relaxation is expressed as a linear program. The fractional solution is scaled down by the worst-case guarantee of an approximation algorithm that "proves" a certain integrality gap. This algorithm is used to construct a randomized mechanism, which provides lotteries over feasible integral solutions obtained by rounding (Lavi and Swamy 2011).

Approximation mechanisms have largely been discussed in theory. We are not aware of any applications in the field, of experiments, or an average-case analysis of such mechanisms. Understanding the average-case performance of approximation mechanisms is important for applications and the lack of empirical literature might explain why such mechanisms cannot be found in the field. In this paper, we study if and when the RaR framework leads to good average-case solution quality and acceptable computation times. We also propose an extension, which addresses the low average-case solution quality of the standard RaR framework, and compare it to a core-selecting auction with near-optimal solutions. These approaches should contribute to an important, but challenging problem in the management sciences, the design of mechanisms to solve computationally hard allocation problems, which are robust against strategic manipulation. A solution to this problem would have substantial impact far beyond the retail transportation logistics problem in this paper.

1.2. Contributions

We focus on a one-to-many time slot management problem. There is a single company (e.g. a retailer) operating multiple warehouses. On behalf of the company's suppliers, carriers visit several warehouses and bid for unloading capacity. The company provides a platform for coordinating the carriers, but instead of a FCFS time slot management system we aim for an economic mechanism in which carriers have simple strategies to reveal their opportunity costs for tours and respective time slots truthfully.

We analyze and compare different approaches to achieve coordination among carriers in retail transportation logistics in situations, where we cannot expect to compute the optimal allocation exactly. While coordination in retail logistics is an important application, the problem of coordinating distributed agents clearly is of general interest and similar problems arise in other domains such as scheduling and timetabling.

The *RaR framework* by Lavi and Swamy (2011) is truthful in expectation. Unfortunately, as we show in this paper, the average solution quality of the RaR framework is as low as the worst-case approximation ratio, which renders the standard approach impractical. Therefore, we introduce an extension, in which we leverage the power of nowadays branch-and-cut (B&C) algorithms. We refer to this as the *RaR B&C framework*. Although this mechanism does not have polynomial-time runtime guarantees anymore, it maintains truthfulness in expectation and is a viable alternative to the other mechanisms. More importantly, it increases the versatility of the RaR framework substantially. Most real-world allocation problems are not simple applications of well-known optimization problems such as set packing or knapsack problems. They have additional constraints and variables and typically no approximation algorithms are known for these specific problems. The RaR B&C framework can still be used due to the versatility of branch-and-cut algorithms, while the RaR framework relies on the existence of an approximation algorithm.

The RaR framework and the RaR B&C framework both share a property of the VCG mechanism, namely that the outcomes might not be in the core with respect to the submitted bids (Ausubel and Milgrom 2006). *Core-selecting auctions* are not incentive-compatible, but one can argue that manipulation is very hard in larger markets as participants would need to possess an unrealistic amount of information in order to manipulate strategically. Core-selecting prices can also be approximated in cases where the allocation problem is not solved exactly Goetzendorff et al. (2015). One way to analyze incentives deviate from truthful bidding is ex post regret, which measures the payoff increase if a single bidder deviated unilaterally from truthful bidding. We show that core-selecting auctions minimize ex post regret and analyze the level of ex post regret in our experiments.

We evaluate the different coordination mechanisms using extensive numerical experiments and draw on field data in our experiments. These experiments reflect important real-world characteristics of problems in the field, but allow for causal inference as we can explore the impact of different mechanisms on the runtime and solution quality under ceteris paribus conditions. We analyze (i) the potential for waiting time reduction compared to uncoordinated arrivals, (ii) the empirical hardness of the computational problem for scenarios of varying size, and (iii) the relation of VCG and core payments to those of the RaR (B&C) mechanisms. For our experiments, we focus on the coordination problem of carriers in retail transportation logistics and use data sets from the field: one data set with the distribution network from a German retailer and a data set about unloading and waiting times from a time-slot management system.

In summary, our results provide evidence that auction mechanisms provide a very effective means for coordination in retail logistics and that they can alleviate the waiting times significantly. In the RaR framework, the (expected) average-case solution quality is very low and there are no significant savings in waiting times. The RaR B&C framework with a branch-and-cut algorithm solving the problems to near-optimality is an alternative, however, which can achieve truthfulness and tractability for realistic problem sizes. Core-selecting auctions with near-optimal solutions achieve the best solution quality, they are stable with respect to submitted bids, but we do not get provable incentive-compatibility. Both approaches are viable, and the choice depends on the relative importance of design desiderata such as efficiency and truthfulness.

The remainder of this paper is structured as follows. Section 2 presents the auction mechanisms along with basic assumptions of the model. We present the experimental design and results in Section 3 and discuss our findings, implications for practice, as well as limitations and future research in Section 4.

2. The auction mechanisms

We describe the auction design problem by first introducing the basic assumptions and the winner determination problem. Then we briefly introduce the different mechanisms, the VCG auction, the RaR framework, and core-selecting auctions. We aim for a succinct but accessible description of the mechanisms in the context of the considered logistics problem.

An *auction* allocates time slots at warehouse loading docks to carriers. That is, routes are not considered directly but encoded in bids on packages of time slots at different warehouses. Carriers constitute the *bidders* in the auction. Carriers are allowed to bid on alternative routes on a daily basis, the *package bids* on bundles of time slots. Carriers need to get reservations for all relevant time slots on a route to ensure they do not have to wait and arrive in time at every warehouse. Hence, there are complementarities of time slots and carriers have super-additive valuations for packages of time slots. The value for a shorter route is proportional to the time saved compared to long routes and the time freed up for truck drivers considering the usual waiting times. The *payments* for the auction are collected by a coordinator who is organizing the auction on behalf of the warehouses. This coordinator might just provide a software application, which computes allocation and payments once per day upon bid entry by the carriers. Such an auction replaces the time-slot management systems that are already used today, but are based on fixed pricing and a simple first-come first-served principle.

2.1. The allocation problem

We consider a single-period coordination problem with $|K|$ warehouses, $|I|$ carriers, and $|T|$ intra-day time slots. The locations of warehouses and carriers are given within the transportation network with known (average) travel distances and travel times. The service capacity of warehouses (loading docks) is modeled as a multidimensional knapsack problem. Carriers have to deliver freight to a warehouse, pick it up there, or both. We assume that each *carrier* has a truck with sufficient capacity to fulfill the orders. The truck starts at the depot and returns to the depot again after (un)loading his freight at the retailers' warehouses. Within the reserved time slots a carrier can unload his freight. In our simulations, we compute a route to visit all warehouses on the tour solving a TSP. Note that every individual truck that can be processed without long waiting times leads to time savings and is beneficial for carriers, i.e., the savings per tour matter primarily. Thus, we consider trucks individually, which also keeps the experimental design simpler. In a design with multiple trucks per carrier we need a number of additional assumptions and treatment variables (e.g., numbers of trucks, costs for using one or multiple trucks). In trial experiments with multiple trucks, we did not see changes in the ranking of auction formats. Hence, we only report the setting with one truck per carrier.

In each time slot $t \in T = \{1, 2, \ldots, |T|\}$, each warehouse $k \in K = \{1, 2, \ldots, |K|\}$ has a capacity of c_{kt}. That is, warehouse k can service up to c_{kt} trucks in parallel in time slot t, but servicing single trucks can require multiple time slots in a sequence. We define the set of items to be allocated as reservations for time slots at warehouses $h_{kt} \in H$. Let A be the set of all possible allocations. Each outcome $(S_1, \ldots, S_{|I|}) \in A$ allocates the set of items S_i to carrier $i \in I = \{1, 2, \ldots, |I|\}$. For example, reservations for carrier $i = 4$ at warehouse $k = 8$ for three time slots starting at $t = 12$ would yield $S_4 = \{h_{8,12}, h_{8,13}, h_{8,14}\}$.

We drop the subscript i for brevity and described item sets $S \in 2^H$, which encode sequence of visited warehouses (routes) as well as the respective time slots, and carriers are allowed to bid on alternative routes S. Each carrier has a valuation function $v_i : 2^H \to \mathbb{R}_{\geq 0}$ with $v_i \in V$, the set of all player types.

Now, the winner determination problem of the coordinator can be formulated as follows.

$$\omega(I) = \max \sum_{i \in I} \sum_{S \neq \emptyset} v_i(S) x_{iS} \tag{WDP}$$

$$\text{s. t.} \sum_{i \in I} \sum_{S : h_{kt} \in S} x_{iS} \leq c_{kt} \qquad \forall k \in K, t \in T, S \neq \emptyset, \tag{1a}$$

$$\sum_{S \neq \emptyset} x_{iS} \leq 1 \qquad \forall i \in I, \tag{1b}$$

$$x_{iS} \in \{0, 1\} \qquad \forall i \in I, S \neq \emptyset. \tag{1c}$$

The objective is to maximize the valuation of accepted bids in WDP. Constraint (1a) ensures that the warehouse capacities are not exceeded for accepted bids for every time slot. Constraint (1b) models the XOR relation of the bids and ensures that at most one bid can be accepted per carrier. Carriers who win in the auction have reservations for the respective time slots at the loading docks, while losing carriers have to queue for service with lower priority.

2.2. The VCG payment rule

Let us first revisit the Vickrey-Clarke-Groves mechanism and its application to WDP as it is fundamental to all other mechanisms discussed below. We use $v \in V$ to denote the vector of true valuations, and $w \in V$ some reported valuations. A mechanism consists of a social choice function $f: V \to A$, mapping valuations to allocations A, and a payment or pricing rule $p_i^{VCG}: V \to \mathbb{R}$ for each player i.

Definition 1 (Vickrey-Clarke-Groves mechanism). A mechanism $\mathcal{M}^{VCG} = (f, p^{VCG})$ is a VCG mechanism if the social choice function f maximizes the social welfare with respect to the reported valuations $w \in V$.

$$f(w) \in \arg\max_{a \in A} \sum_{i \in I} w_i(a)$$

and for each agent $i \in I$ there exists a function $h_i: V_{-i} \to \mathbb{R}$ such that the pricing scheme p_i^{VCG} is computed as

$$p_i^{VCG}(w) = h_i(w_{-i}) - \sum_{j \in I \setminus \{i\}} w_j(f(w)).$$

The VCG payment scheme aligns an agent's utility with the objective of optimizing the social welfare. To achieve this, the second term of the VCG pricing scheme pays each agent i the sum over all valuations reported by the other agents. Together with the agent's personal valuation, this sum equals the social welfare. Furthermore, since no agent has direct influence on the first term of his pricing scheme h_i, his utility is indeed optimized if he reveals his valuation truthfully.

Proposition 1 **(Vickrey (1961), Clarke (1971), Groves (1973))**. *The VCG mechanism* $\mathcal{M}^{VCG} = (f, p^{VCG})$ *is strategy-proof.*

The proof is a basis for the subsequent discussion and can be found in Appendix A together with all other proofs. Individual rationality is an important requirement to ensure that no agent has negative utility if he reveals his true valuation. Clarke's pivot rule specifies a VCG payment scheme which is individually rational and avoids charging negative prices.

Definition 2 (Clarke's pivot rule). A pricing scheme for a VCG mechanism implements Clarke's pivot rule if for all agents $i \in I$ and valuation functions $w_{-i} \in V_{-i}$ it holds that

$$h_i(w_{-i}) = \max_{a \in A} \sum_{j \in I \setminus \{i\}} w_j(a).$$

The resulting pricing scheme $p_i^{VCG} = (\max_{a \in A} \sum_{j \in I \setminus \{i\}} w_j(a)) - \sum_{j \in I \setminus \{i\}} w_j(f(w))$ charges each agent the difference between the optimal social welfare with and without his participation. In other words, his payments reflect the externalities he causes the other agents, which are neither negative nor do they exceed $w_i(f(w))$.

2.3. The RaR framework

Truthfulness of the VCG mechanism hinges on the optimality of the computations as the proof to Proposition 1 shows. Many real-world allocation problems are NP-hard optimization problems and therefore the VCG mechanism is not strategyproof anymore if we cannot guarantee optimality. Approximation mechanisms give up on optimality, but make sure that the mechanism still incentivizes truthful bidding. The literature in this field is still young. In a seminal contribution, Lavi and Swamy (2011) introduce a framework that allows to turn any packing problem for which an approximation algorithm exists, into a truthful mechanism. We briefly introduce relevant terminology and refer to it as the RaR framework as it relies on relaxation and rounding.

Randomized approximation mechanisms $\mathcal{M}^R = (f^R, p^R)$ are a probabilistic extension of deterministic mechanisms which compute outcomes and prices according to some internal random process. As a result, $f^R(w)$ and $p_i^R(w)$ are random variables. Strategyproofness is defined for deterministic mechanisms and we need new notions of truthfulness for randomized mechanisms. A randomized mechanism is called *truthful in expectation* if each player optimizes his expected utility by declaring his valuation truthfully.

Definition 3 (truthfulness in expectation). A mechanism $\mathcal{M}^R = (f^R, p^R)$ is truthful in expectation if, for every bidder i, true valuation function v_i, reported valuation function w_i, and (reported) valuation functions w_{-i} of the other bidders,

$$E[v_i(f_i^R(v_i, w_{-i})) - p_i^R(v_i, w_{-i})] \geq E[v_i(f_i^R(w_i, w_{-i})) - p_i^R(w_i, w_{-i})].$$

The basic idea of the RaR framework by Lavi and Swamy (2011) is to move to a fractional domain and use the VCG mechanism to obtain a *truthful fractional mechanism* $\mathcal{M}^F$. Let $\mathcal{P}$ denote the feasible region of the linear program (LP) relaxation of (WDP) and $\mathbb{Z}(\mathcal{P}) \subseteq \mathcal{P}$ the set of integer solutions of (WDP). To assign valuations to fractional solutions $x \in \mathcal{P}$, we define $v_i(x) = \sum_{S \neq \emptyset} v_i(S) x_{iS}$. The LP relaxation of the underlying allocation problem can be solved exactly in polynomial time returning an optimal fractional solution x^* and prices p^F. We can scale down allocation and prices of the fractional solution by α, the approximation ratio of an approximation algorithm $\mathcal{A}$, while the mechanism will remain truthful.

A deterministic support mechanism $\mathcal{M}^D$ takes the scaled down fractional solution and generates a probability distribution D over integer solutions one of which will be randomly selected. This distribution D is computed via a convex decomposition of the scaled-down fractional solution into polynomially many integer solutions. Scaling down the fractional solution x^* by the integrality gap of the WDP ensures that this scaled fractional solution can be described as a convex decomposition of feasible integer solutions x^l. The integrality gap is defined as

$$\sup_{v \in V} \frac{\max_{x \in P} \sum_{i,S} v_i(S) x_{i,S}}{\max_{x \in \mathbb{Z}(\mathcal{P})} \sum_{i,S} v_i(S) x_{i,S}}$$

An α-approximation algorithm $\mathcal{A}$ needs to verify an integrality gap of at most α, i.e., for any valuation vector $v \in V$ the algorithm can produce an integer solution of at least $\frac{1}{\alpha}$ times the welfare of the fractional solution.

The convex decomposition yields a probability distribution with a probability λ_l for feasible integer solutions x^l with, $x^* = \lambda_l x^l$, $\lambda_l \geq 0$ and $\sum_{l \in \mathbb{I}} \lambda_l = 1$. Finally, we randomly pick one of these integer solution; $\{\lambda_l\}_{l \in \mathbb{I}}$ describes the probabilities for selecting each of the solutions. This yields the randomized mechanism $\mathcal{M}^R$. In summary, the framework consists of the following three steps:

1. Compute a solution $x^* \in \mathcal{P}$ of the LP relaxation of the WDP $\max\{\sum_{i \in I, S \neq \emptyset} w_i(S) x_{iS} \mid x \in \mathcal{P}\}$ over some $|I| \times 2^{|H|}$-dimensional polytope $\mathcal{P} \subseteq [0,1]^{|I| \times 2^{|H|}}$. Then, we determine the corresponding VCG prices p_i^F for each agent i. This yields a truthful mechanism $\mathcal{M}^F$ for the fractional problem.
2. Use an α-approximation algorithm $\mathcal{A}$ to decompose the scaled-down fractional solution $\frac{x^*}{\alpha}$ into a convex combination $\lambda_l \in [0,1]^{\mathbb{I}}$ of integral outcomes $x_l \in \mathbb{Z}(\mathcal{P})$. The distribution of λ_l is independent of the valuations and only depends on the constraints of WDP such that bidders cannot manipulate the distribution $\{\lambda_l\}_{l \in \mathbb{I}}$. The resulting distribution is deterministic and the mechanism generating it is referred to as deterministic support mechanism $\mathcal{M}^D$.
3. Choose outcome $x_l \in \mathbb{Z}(\mathcal{P})$ at random with probability λ_l and charge each agent i a price of $p_i^R(w) = w_i(x^l) \cdot \frac{E\left[p_i^R(w)\right]}{E\left[w_i(f^R(w))\right]}$ if $E\left[w_i(f^R(w))\right] > 0$ and $p_i^R(w) = 0$ if $E\left[w_i(f^R(w))\right] = 0$. This yields an α-approximation mechanism $\mathcal{M}^R$ that is truthful in expectation.

The first step can be implemented via standard linear programming algorithms. The implementation of the third step is also straightforward from a computational point of view. The central piece in this framework is step 2, where a distribution over outcomes needs to be computed, i.e., a convex combination of integral solutions, matching the relaxed solution in expectation. For this task, Lavi and Swamy resort to a decomposition technique by Carr and Vempala (2000), which models the convex decomposition as an LP. However, since the LP has exponentially many variables, one for each outcome, it is not solved directly. The dual, which has exponentially many constraints, is solved via the ellipsoid method. The use of the ellipsoid method is instrumental as it does not need a complete enumeration of the constraints provided that a suitable separation oracle exists (Bland et al. 1981). As separation oracle in the ellipsoid method, we use the approximation algorithm $\mathcal{A}$. We provide a description of the convex decomposition of the scaled down fractional solution in the context to our logistics problem in Appendix B.

Lavi and Swamy (2011) prove that mechanisms following the RaR framework are truthful in expectation. Although the RaR framework is rather generic, there are a few important prerequisites and the framework cannot be applied to any allocation problem modeled as an integer linear program.

1. Polytope $\mathcal{P}$ must satisfy the *packing property*: $\forall x, x' \in \mathcal{P} : x' \in \mathcal{P} \wedge x' \geq x \implies x \in \mathcal{P}$.

2. The integrality gap of $\mathcal{P}$ must be bounded by an efficiently *verifiable* value $\alpha \geq 1$. More precisely, a polynomial-time approximation algorithm $\mathcal{A} : V \to \mathbb{Z}(\mathcal{P})$ must exists such that $\alpha \sum_{i \in I, S \neq \emptyset} w_i(S)\mathcal{A}(w)_{iS} \geq \max\{\sum_{i \in I, S \neq \emptyset} w_i(S) x_{iS} \mid x \in \mathcal{P}\}$ for any $w \in V$.

The two requirements make sure that a scaled down fractional solution is feasible in $\mathcal{P}$ and there is at least one decomposition into integer solutions that are feasible for $\mathbb{Z}(\mathcal{P})$. In the following proposition, we show that the retail logistics problem actually satisfies the necessary requirements for the RaR framework.

Proposition 2. *The WDP can be implemented as a truthful-in-expectation mechanism following the RaR framework.*

There are two properties that render the RaR framework impractical in most environments.

1. Many real-world allocation problems differ from the textbook optimization problems, where approximation algorithms are known. There are additional constraints or terms in the objective function and typically no polynomial-time approximation algorithm with a worst-case approximation guarantee α is known for these specific problems. While there is an approximation algorithm for the allocation problem as we show below, such an algorithm may not be applicable with additional constraints, e.g., on the number of winners, required for problems in practice.

2. The expected social welfare is the optimal social welfare scaled down by α in order to ensure that for any valuation profile $v \in V$ we find a feasible integer solution. It is important to observe that scaling down the fractional solution by α turns the worst-case solution quality of an algorithm into the average-case solution quality of the mechanism, as we show as a corollary to Lemma 3.2 by Lavi and Swamy (2011).

Corollary 1. *The expected approximation ratio of a randomized mechanism $\mathcal{M}^R$ in the RaR framework is equivalent to the worst-case approximation ratio α of the approximation algorithm $\mathcal{A}$ employed.*

This is a significant difference to worst-case guarantees of approximation algorithms otherwise, which mark a lower bound to the solution quality of an algorithm, not the average-case. The average-case solution quality of approximation algorithms is typically much better than the worst-case. More importantly, worst-case approximation ratios of algorithms for most combinatorial optimization problems are very low. This is also true for the WDP of our logistics problem. We draw on the approximation algorithm by Briest et al. (2011) with a worst-case approximation ratio of

$O((K \cdot T)^{\frac{1}{c+1}})$ where $c \geq 1$ is the minimal capacity or number of loading docks available in parallel in one time slot and at one warehouse. With 90 time slots per day and 10 warehouses with capacity 1, the fractional solution needs to be scaled down by a factor of more than 240 leading to a very low solution quality. As a consequence, the mechanism often selects the null allocation as a feasible allocation. This is impractical for any application of the RaR framework with many items for sale.

2.4. The RaR B&C framework

The RaR B&C framework addresses the two problems of the RaR framework. First, we replace the approximation algorithm $\mathcal{A}$ by a branch-and-cut algorithm in order to find integer solutions during the ellipsoid method. We restrict the runtime of the algorithm by a certain time limit that yields near-optimal allocations on average. Second, rather than scaling down the fractional optimal solution by α, we scale it down a parameter δ, which can be much lower.

Let us discuss this scaling in more detail. The only reason to scale down the fractional solution by the worst-case approximation ratio α is to guarantee that there can never be a scaled fractional solution that cannot be represented as convex decomposition of integer solutions. The price for this guarantee is that the solution quality is forced down to impractical levels. Branch-and-cut algorithms are very powerful today and they can also be used to sample new feasible solutions for the convex decomposition in the RaR framework. Near-optimal solutions can almost always be found very quickly for real-world problems and problem sizes, even though finding a provably optimal solution may turn out to be intractable. This is a well-known phenomenon in combinatorial optimization (Goetzendorff et al. 2015), which is particularly relevant for RaR. As a result, we can use branch-and-cut algorithms with a limit on the time per run or the duality gap as a general-purpose algorithm to sample integer solutions within the ellipsoid algorithm.

Such a randomized mechanism remains truthful in expectation, but has an expected solution quality of δ, which makes the RaR framework practical and applicable to a broad range of allocation problems that satisfy the packing property. Determining an appropriate parameter δ requires some prior experimentation with data that are representative for the real data, since it directly determines the expected approximation ratio. If it is set too small, it might be that no convex decomposition of the scaled down fractional solution $\frac{x^*}{\delta}$ is possible (δ then has to be increased and computations have to be restarted). One can use bisection to determine the optimal δ in a logarithmic number of steps. A few initial experiments on synthetic bid data can provide very good estimates and reduce the range of δ-values significantly. More importantly, this step does not impact truthfulness.

Let us now show that the RaR B&C framework remains truthfulness in expectation as a corollary to Theorem 3.1 in Lavi and Swamy (2011) and that it computes a δ-approximation with δ being the duality gap for the WDP.

COROLLARY 2. *A randomized mechanism in the RaR B&C framework is truthful in expectation.*

PROPOSITION 3. *A randomized mechanism in the RaR B&C framework computes a δ-approximation to the maximum social welfare and is truthful in expectation.*

Formally, a randomized mechanism in the RaR B&C framework is not polynomial-time. However, our experiments show that RaR B&C provides an effective solution to a broad range of allocation problems. The result is appealing: So far, truthful mechanisms to intractable allocation problems such as the WDP have not been known. The expected solution quality RaR framework is impractical, but the RaR B&C framework can strike a balance between solution quality and incentives to bid truthful while still allowing to solve realistic problem sizes.

2.5. Core-selecting auctions

One of the main problems of the VCG mechanism is the fact that the outcome is not always stable (Ausubel and Milgrom 2006). More precisely, VCG payments can be such that a sub-coalition of bidders would have an incentive to deviate together with the auctioneer, i.e. the solution is not in the core. This basic tension between incentive-compatibility and stability is an issue whenever the VCG payment rule is used, i.e. also in mechanisms within the RaR and the RaR B&C frameworks. Stability is considered crucial in many domains and for this reason core-selecting auctions have been adopted for the sale of spectrum licenses world-wide (Bichler and Goeree 2017). In this subsection, we briefly introduce core-selecting auctions, how they are adapted for near-optimal solutions to the allocation problem, and characterize their robustness against manipulation in the context of our logistics problem.

The lack of stability in the VCG payment rule is best understood if the auction is modeled as a coalitional game. Now $\omega(\cdot)$ is referred to as the coalitional value function, i.e., the welfare of the auction game with a certain set of bidders. Let N denote the set of all bidders I plus the auctioneer with $i \in N$ and let $M \subseteq N$ be a coalition of bidders with the auctioneer. Let $\omega(M)$ denote the coalitional value for a subset M, equal to the objective value of the WDP with all bidders $i \in M$ involved. (N, ω) is the coalitional game derived from trade between the auctioneer and all bidders. A core payoff vector Π, i.e., the payoffs of the bidders in this auction and the revenue of the auctioneer, is then defined as follows

$$Core(N,\omega) = \{\Pi \geq 0 | \sum_{i \in N} \pi_i = \omega(N), \sum_{i \in M} \pi_i \geq \omega(M) \quad \forall M \subseteq N\}.$$

This means, there is no coalition $M \subset N$ that has a strictly greater coalitional value than the winning coalition. In other words, no group of bidders can make a counteroffer that leaves themselves at least as well and the auctioneer better off. Note that the definition is based on payoffs

π_i not prices p_i. Unfortunately, in the VCG auction there can be outcomes which are not in the core. An example with three bidders and two items x and y should make the case (Ausubel and Milgrom 2006). Bidder 1 bids $b_1(x) = \$0$, $b_1(y) = \$2$ and $b_1(x, y) = \$2$. Bidder 2 bids $b_2(x) = \$2$, $b_2(y) = \$0$ and $b_2(x, y) = \$2$. Finally, bidder 3 only has a bid of $b_3(x, y) = \$2$, but no valuation for the individual items. In this situation, the net payments of the winners (bidder 1 and 2) are zero, and bidder 3 could find a solution with the auctioneer that makes both better off. Outcomes, which are not in the core, lead to a number of problems, such as low seller revenues or non-monotonicity of the seller's revenues in the number of bidders and the amounts bid. To see this, just omit bidder 1 from the example. Also, such auction results are vulnerable to collusion by a coalition of losing bidders. Therefore, it has been argued that the outcomes of combinatorial auctions should be in the core (Day and Raghavan 2007). Also, in the allocation of tours to carriers it might be hard to defend that some carriers need to pay less for a set of time slots than what losing bidders were willing to pay.

Computing whether the core is empty for a combinatorial optimization game or not can be done in polynomial time (Deng et al. 1999). Bikhchandani and Ostroy (2002) showed that in a single-sided combinatorial auction the core is always non-empty, while in a combinatorial exchange the core might be empty. The following linear program (Core) describes conceptually how minimal core payments with respect to the submitted bids can be computed once the optimal allocation is known. Here, π_i is the payoff of a bidder i and π_a the revenue of the auctioneer. The coalitions $M \subset N$ include all coalitions with the auctioneer a included.

$$\min \sum_{i \in \mathcal{I}} \pi_i \qquad \text{(CORE)}$$

$$\text{s.t.} \sum_{i \in M} \pi_i + \pi_a \geq \omega(M) \qquad \forall M \subset N \qquad \text{(2a)}$$

$$\sum_{i \in \mathcal{I}} \pi_i + \pi_a = \omega(N) \qquad \text{(2b)}$$

$$\pi_i, \pi_a \in \mathbb{R} \qquad \forall i \in \mathcal{I} \qquad \text{(2c)}$$

Conitzer and Sandholm (2006) pointed out that it requires many constraints for all the sub-coalitions. Rather than enumerating core constraints for all possible coalitions, Day and Raghavan (2007) propose an effective iterative algorithm that generates core constraints on the fly, which we also use in our experiments. The procedure to generate core constraints is described in Appendix C. This procedure can easily be adapted to cope with near-optimal solutions to the WDP as we show and it yields prices $p_i^{core}(w_i)$ for bidder i.

Unfortunately, core-selecting auctions cannot be incentive-compatible unless the VCG outcome is in the core (Goeree and Lien 2016). To see this, in our earlier example with three bidders and two items x and y, both winners would have to pay a price of \$2 in total, e.g. \$1 each. Now, one of them could try to free-ride on the other by claiming his valuation is strictly less than \$1. While such types of manipulations are obvious in small toy examples with complete information, they become much harder under incomplete information in large markets with many items and many bidders. A bidder would need to have prior distributional information about exponentially many packages, which is unrealistic. Often bidders do not even know the number of competitors in such an auction and the observation that such mechanisms cannot be incentive-compatible according to established equilibrium solution concepts might be too pessimistic. The volume of information needed to manipulate profitably and the computational hardness of the allocation problem are two arguments that make strategic manipulation in the context of our logistics problem very difficult to say the least.

It is useful to compute quantitative measures of susceptibility. Lubin and Parkes (2012) introduced ex post regret as such a measure for mechanisms that are not truthful. The ex post regret an agent has for truthful reporting in a given game is the amount by which its utility could be increased through a misreport. More formally, ex post regret to an agent i is $\sup_{w \in V}(u_i(w_i, w_{-i}) - u_i(v_i, w_{-i}))$, where $u_i(w_i, w_{-i}) = v_i(f(w_i, w_{-i}) - p_i(w_i, w_{-i})$ when bidder i reports w_i. If there is no ex post regret, then obviously a mechanism is dominant-strategy incentive compatible. The ex post regret can be computed explicitly for core-selecting auctions.

Proposition 4. *The ex post regret of bidder i in a core-slecting auction is $p_i^{core}(v_i, w_{-i}) - p_i^{VCG}(v_i, w_{-i})$.*

The bidder-optimal core computations described in Appendix C minimize total payments and therefore minimize ex post regret among the set of core-selecting auctions. It is interesting to understand the level of ex post regret in a realistic environment. In our experiments, this regret was actually low. One can argue that this sets incentives for truthful bidding, because a bidder needs to trade off the potentially low gains from an unilateral deviation with his risk of losing due to this strategic bidding behavior.

3. Experiments

In this section, we provide details of our experimental evaluation. We first describe the experimental design and the main treatments. Then we present and analyze the results regarding carrier waiting times, computation times, the numbers of submitted and winning bids, and payments.

3.1. Experimental design

In what follows, we describe the auction mechanisms implemented, the simulation environment, and the treatment variables of the full factorial design.

3.1.1. Auction mechanisms The main treatment variable in our experiments is the type of auction or coordination mechanism used. We assume that carriers submit bids based on their opportunity costs to the mechanism. That is, carriers base their bids on estimated time savings, since they can use their resources for other tasks if waiting time is reduced. The mechanism then determines an assignment of bidders to time slots on a given day and a payment for these slots.

In our baseline treatment without coordination (no coord.), there is no information available to the carriers. Carriers do their own planning independently and then start processing their plans in the morning. Arriving at a warehouse, the carrier has to queue up if the loading dock is occupied. Carriers plan their route (solving a TSP) with respect to the travel times while visiting all warehouses on the tour.

Apart from this baseline treatment, we use a VCG mechanism (VCG) and a mechanism computing core payments (Core), which are based on optimal allocations (OPT). Both result in the same allocations, but compute different payments. Further, we also compute near-optimal allocations with a 10% MIP gap (Core0.1). The RaR framework treatment (RaR) uses an approximation algorithm $\mathcal{A}$ as described in Section 2.3. The RaR B&C framework replaces the approximation algorithm and respective integrality gap with a branch-and-cut algorithm and a MIP gap parameter δ. In one treatment, we set this MIP gap to 100% (RaR B&C1.0) and in another one to 10% (RaR B&C0.1). A 100% MIP gap means that the branch-and-cut algorithm stops as soon as the best feasible integer solution found is half the value of the best LP relaxation ($\delta = 2$).

We analyze problem sizes where we can, with enough computation time, determine the optimal solution to be able to compare these results with the other treatments.

3.1.2. Transportation network and vehicle routing We draw on data about the distribution network of a German retailer. The transportation network also provides the distances and typical travel times between the locations. Overall, we have 65 locations of retail warehouses with distances and average travel times between all the sites in an adjacency matrix.

The locations of 10 warehouses and the depots of the carriers are determined randomly in each simulation by drawing from nodes of a transportation network from the field without replacement. Each carrier has 4 or 6 warehouses that he has to visit on a tour. The set of 4 or 6 warehouses on a tour is randomly selected from the set of 10 warehouses without replacement in each experiment.

Each carrier determines all routes without waiting time through the warehouses and bids on those that do not take longer than 110% of the optimal round trip time (RTT). The allocation problem sizes are such that we can compute exact solutions. This provides us with an estimate for the efficiency loss that we experience in larger settings, where we need to restrict to approximations or near-optimal computations.

After the routes for each carrier have been determined, we compute a bid price for each route, which is a package bid in the auction. The bid price is proportional to the time saved per day. We compute the difference between the full day and the time needed for the optimal route. For example, if there are 8 hours on a working day and the optimal route takes 6 hours without waiting times, we assume a willingness to pay of 3 monetary units per warehouse reservation for this ideal route. This number for the willingness to pay is based on empirical observations (Bundesamt für Güterverkehr 2011, p. 25). If the time savings are less than for the ideal route, a carrier is also willing to pay less per reservation. For example, for an alternative route that takes 7 hours (i.e. half of the time savings in the optimal route), the willingness-to-pay is only 1.5 monetary units per reservation.

Of course, carriers might determine prices differently in the field and we do not attach meaning to the absolute numbers for payments. However, the relative ranking of payments might give some additional insights beyond the waiting time reductions and computation times that we report. All bid data generated is available upon request for replication studies and comparisons.

3.1.3. Simulation framework Our experiments are conducted as a discrete event simulation with six types of events. The first event is the *departure event* which takes place when a carrier leaves his depot. It results in a *travel event* describing the carrier's travel from his current location to his next one. The travel event is succeeded by an *arrival event*. It represents the carrier arriving at a warehouse and queuing up in order to be loaded or unloaded respectively. Carriers who won one of their routes and therewith respective reservations are serviced with strict priority at loading docks if they arrive during the time slots, while losing carriers have to queue for service with lower priority (first-come first-served (FCFS) among carriers without reservations). Thus, the arrival event is either followed by a *wait event*, or a *reservation event*. The latter is directly triggered only if the truck arrives at the time when the reservation is valid. If the truck arrives early or does not have a reservation, the wait event is triggered. Travel times are estimated from the field data and stay constant over the day. While this simplifies the environment compared to the field, it makes the comparison between the relative solution quality of the mechanisms easier.

The *service event* is triggered by the warehouse, whenever the loading or unloading process is started for a carrier. Loading and unloading services may require multiple time slots. Having completed the service process, a carrier drives to the next warehouse or his depot if he already completed his tour. The simulation ends when every carrier has reached his depot again.

3.1.4. Parameters and treatment variables The values for each parameter and treatment variable are summarized in Table 1. We assume different *numbers of carriers* ($|I| = \{20, 30, 40\}$) and 10 warehouses ($|K|$) in each experiment. A carrier needs to visit 4 or 6 warehouses on a tour.

A time slot in our simulation is 10 minutes and we consider 90 time slots per day ($|T|$). The unloading time is drawn from a distribution based on data about reservations and unloading times from a time slot management system. We assume a "typical" service time of $\hat{\mathcal{S}} = 30$ minutes per warehouse and travel times $\hat{\mathcal{T}}$ based the historical field data available. We use expected service and travel times for carriers' route planning and bidding, but vary these times in the event simulation described. We assume service times and travel times to be normally distributed with a standard deviation of 5 minutes for service times and 10% for travel times. The Normal distributions are truncated to ±25% of the mean, in order to avoid extremely short or long service and travel times.

	name	values
Parameters	number of locations	65
	number of warehouses ($|K|$)	10
	travel time	distribution from field data
	service time	distribution from field data
	time slot length	10 min
	number of time slots per day ($|T|$)	90
	number of simulations / treatment	20
Treatment variables	warehouse capacity	$\{1, 2\}$
	warehouses per carrier	$\{4, 6\}$
	number of carriers ($|I|$)	$\{20, 30, 40\}$
	auction mechanisms	{no coord., OPT, Core0.1, RaR, RaR B&C1.0, RaR B&C0.1}

Table 1 Parameters and treatment variables

For the RaR (B&C) treatments, we limit execution of the ellipsoid method to 10,000 iterations. If we cannot decompose the fractional solution by a convex combination of integer solutions within this time, we increase the duality gap and restart until we are able to implement the fractional solution into a convex combination of feasible integer solutions. The results of the RaR (B&C) treatments are probability distribution over integer solutions. Hence, we conduct simulations for every integer solution used in the convex combination and calculate the *expected* waiting and round trip times by weighting the results with the corresponding probabilities. Analogously, we report expected values for the numbers of winning bids, prices, revenue, and welfare for the RaR (B&C) treatments.

We use a fully factorial experimental design with $2 \cdot 2 \cdot 3 \cdot 6 = 72$ treatment combinations, for which we randomly draw locations for warehouses and carriers 5 times each resulting in *360* individual simulation scenarios. We draw travel and service times for these scenarios 20 times each resulting in *7,200* simulation experiments in total. In each simulation, all auctions are computed involving a multitude of optimization problems that needed to be solved in each experiment. The total computation time for all optimization problems was approximately *272* hours.

The simulation was implemented in the Java programming language and the commercial mathematical programming solver Gurobi Optimizer v6.5 was used for all optimization problems. Experiments were executed on computing nodes with two 10-core CPUs (Intel Xeon E5-2660 v2) and 240GB of RAM each.

3.2. Results

There is a trade-off between the computational costs of a mechanism and the waiting time reduction that can be achieved. In addition, some mechanisms satisfy strong game-theoretical solution concepts, while others do not, and truthful bidding in these mechanisms relies on the assumption that bidders do not have sufficient information about others to manipulate strategically in our environment. The experiments should shed light on these trade-offs. In what follows, we report the savings in the waiting times of carriers due to the use of certain coordination mechanisms, the computation times they incur, and the payments bidders have to make. More detailed results for different treatment combinations about waiting times, computation times, seller revenue, payments, and the numbers of submitted and winning bids can be found in Appendix D in Tables 11 to 13.

Result 1 *The optimal and the near-optimal allocation of tours led to a significant reduction of carrier waiting times for all auction mechanisms except RaR. The waiting time reduction of the RaR B&C framework depends on the MIP gap one can achieve for a specific problem size. Across all experiments we get a ranking of the auction formats by waiting time reduction: OPT > Core0.1 > RaR B&C0.1 > RaR B&C1.0 > RaR.*

The average round trip time was 482 minutes (567 minutes in treatment combinations with a warehouse capacity of 1 and 397 minutes with a warehouse capacity of 2) across all participants and treatments. The mean waiting time for all bidders without coordination was 140.14 minutes (st. dev. 123.89) per tour. This is significantly higher than the mean waiting time of 113.16 minutes (st. dev. 113.18) per tour for *all* bidders when an optimal allocation can be computed. This means, coordination mechanisms lead to a better allocation and they provide a significant improvement for all bidders, not just for the winners. Table 2 shows the mean waiting time for all mechanisms and the difference to the treatments with no coordination. The differences between OPT, Core0.1, RaR B&C, and the RaR mechanisms to the experiments without coordination were all significant at $p < 0.05$ using a Wilcoxon signed rank test. Table 11 in Appendix D provides details of waiting times per loading ramp for all treatment combinations.

We also report the results of a linear regression analyzing the impact of different factors on the average waiting times. This allows us to control for the treatment variables (number of carriers,

warehouse capacity	1–2		1		2	
mechanism	mean	mean diff. no coord.	mean	mean diff. no coord.	mean	mean diff. no coord.
no coord.	140.14		229.03		51.25	
OPT	113.16	26.98	195.89	33.14	30.43	20.82
Core0.1	115.45	24.69	198.02	31.01	32.88	18.37
RaR B&C0.1	121.37	18.77	207.02	22.01	35.72	15.53
RaR B&C1.0	127.01	13.13	211.89	17.14	42.12	9.13
RaR	139.83	0.31	228.87	0.16	50.78	0.47

Table 2 Waiting time (in minutes) per tour and mean saving in waiting times compared to no coordination.

warehouses per carrier, warehouse capacity). With no coordination as a baseline treatment, the waiting time reductions are provided in Table 3, and they are in line with the averages reported in Table 2. However, according to the regression analysis RaR B&C1.0 and RaR do not improve the waiting time significantly.

	Per Tour	p-value
intercept	21.72	$p < 1.000$
number of carriers	7.17	$p < 0.001$
warehouses per carrier	32.04	$p < 0.001$
warehouse capacity	-171.26	$p < 0.001$
Mechanisms		
OPT	-26.98	$p < 0.001$
Core0.1	-24.68	$p < 0.010$
RaR B&C0.1	-18.77	$p < 0.050$
RaR B&C1.0	-13.13	$p < 1.000$
RaR	-0.31	$p < 1.000$
Adjusted R^2	0.86	

Table 3 Waiting time reduction (in minutes) per tour and per loading dock for all bidders based on a regression analysis.

To study the effects of increased competition between carriers, we ran additional simulation experiments with “hotspot” warehouses. These hotspots are warehouses which have to be visited by every carrier. The resulting waiting times for no coordination and optimal allocations are shown in Table 4. The results illustrate that depending on the competition for particular time slots the potential waiting time reductions vary. We do not report the detailed results for these extra treatment combinations, but the ranking of auction mechanisms by waiting time reduction remained. This was also the case for experiments where we varied the number of competing carriers.

The waiting time reductions reported are based on averages across all carriers. Winners in the auction would ideally have no waiting time at all. Due the fact that there are stochastic travel and

no. of hotspots	mechanism	mean	mean diff.
1	no coord.	239.37	
1	OPT	220.72	18.65
2	no coord.	307.71	
2	OPT	265.17	42.55
3	no coord.	338.75	
3	OPT	283.70	55.05

Table 4 Waiting time (in minutes) per tour for different numbers of hotspots.

service times also winners have some waiting time in the simulation. For example, if the carrier with a reservation has not yet arrived, but another carrier without a reservation is at the loading dock, then the one without reservation would be processed. If the carrier with reservation arrives later, he needs to wait until the one without reservation is finished. These waiting times of winners were very low, however.

It is important to understand why the RaR framework leads to minimal savings in waiting time only. The RaR framework turns the integrality gap into an average-case approximation ratio. Each combination of a warehouse and time slot constitutes an object in the auction, resulting in $K \cdot T$ objects. The integrality gap for the multi-unit combinatorial auction is $\mathrm{e} \cdot (2\gamma + 1) \cdot (K \cdot T)^{\frac{1}{c+1}}$ for any minimal non-zero warehouse capacity $c = min_{k \in K, t \in T}\{c_{kt} : c_{kt} \geq 1\}$, where γ is the approximation ratio of a polynomial time demand oracle which returns the most valued bundle for a bidder and a given price vector (Briest et al. 2011). With $T = 90$ time slots, $K = 10$ warehouses, and a (minimal) warehouse capacity of 1 loading dock (a single unit combinatorial auction) this leads to $\mathrm{e} \cdot 3 \cdot 900^{\frac{1}{1+1}} = 244.65$, i.e., the fractional solution is scaled down by this factor. Only this scaling guarantees that there will always be a feasible integer solution, which does not violate a constraint.

As a consequence, the outcome of the approximation is often an allocation without any winner. Table 5 provides an overview of the percentage of allocations, where the RaR framework or the RaR B&C framework do not allocate a single tour (i.e., an empty allocation, "% empty"), and the percentage of non-empty allocations with only one winning bidder ("% non-empty with 1 winner"). This explains the low performance of the standard RaR framework. The problem is substantial for the RaR framework and renders it impractical. In contrast, allocations with no or only one winning bidder never occur in the OPT and Core0.1 treatments.

Result 2 *For the problem sizes in our experiments, a computation of allocation and core payments solved to optimality took at most 28.84 hours (52.06 minutes on average), whereas the RaR framework took at most 2.17 seconds, and the RaR B&C with 100% MIP gap took at most 7.15 minutes (1.34 minutes on average) to compute. The RaR B&C framework with a very low MIP gap of 10% was more expensive, because many optimization problems need to be solved to near-optimality. In*

		mean	std. dev.
RaR	% empty	82.11	13.80
	% non-empty with 1 winner	96.48	13.01
RaR B&C1.0	% empty	2.11	9.31
	% non-empty with 1 winner	8.38	7.50
RaR B&C0.1	% empty	0.16	1.18
	% non-empty with 1 winner	0.16	0.59

Table 5 Percentages of RaR and RaR B&C allocations with no winner and percentages of allocations with only one winner given that the allocation was non-empty.

the worst case, the computation took 14.06 hours (3.40 hours on average). The average computation time for core computations with near-optimal solutions only took 64.96 seconds on average.

Table 6 provides the mean and maximum of the computation times for the different mechanisms and pricing rules in seconds. To better understand the average problem sizes solved, Table 7 provides the average total number of bids and the average number of bids per carrier in the experiments. We also report the average number of winning bids across all treatments. On average an allocation problem included 255.88 bids from carriers on various routes. Carriers submitted 8.52 bids on average (for details see Table 13 in Appendix D).

In the RaR treatment, for most of the problems a feasible decomposition was found already with the initial integer solutions that are known to be valid either due to the packing property or because only a single assignment is made. This explains the very low average computation times for this treatment. Computing core payments based on optimal solutions was computationally most expensive. In this case, the computation time could be reduced substantially, when we restricted all computations to near-optimality with a 10% MIP gap (see line “Core (Core0.1)” in Table 6). This makes it a very practical approach even for large problem sizes.

	mean	max.
Pay-as-bid (OPT)	34.80	914.03
VCG (OPT)	1037.23	27836.04
Core (OPT)	3123.47	103834.17
Core (Core0.1)	64.96	663.66
RaR B&C1.0	80.40	428.78
RaR B&C0.1	12252.50	50618.55
RaR	0.69	2.17

Table 6 Computation times (seconds).

Result 3 *In expectation carriers paid only a small proportion of their bids in the RaR and the RaR B&C auctions. Core and VCG payments were significantly higher in the experiments.*

Auction	∅ no bids	∅ no. win. bids	∅ no. bids p. carrier	∅ no. win. bids p. carrier
OPT	255.88	19.27	8.52	0.67
RaR	255.88	0.20	8.52	0.01
RaR B&C0.1	255.88	15.28	8.52	0.53
RaR B&C1.0	255.88	10.34	8.52	0.36
Core0.1	255.88	18.60	8.52	0.65

Table 7 Average numbers of (winning) bids.

Table 8 provides an overview of the average auctioneer revenue for the different auction mechanisms. The first line in the table describes the welfare, this means the sum of the values in the optimal allocation. Note that in the RaR framework average revenue is low, because there is a much lower number of winners. The revenue raised via core payments is higher than that in the VCG mechanism, which is to be expected.

	mean	std. dev.
Welfare	224.87	84.36
Core (OPT)	159.06	94.08
VCG (OPT)	122.34	83.21
Core (Core0.1)	173.74	88.39
RaR B&C0.1	97.19	62.12
RaR B&C1.0	68.54	38.66
RaR	1.18	1.16

Table 8 Average auctioneer revenue (in monetary units).

The ex post regret for the two core pricing treatments is shown in Table 9, i.e., the difference of core and VCG prices, which corresponds to the amount bidders could benefit by shading their bids unilaterally. The ex post regret in the core-selecting auctions was at around 2 monetary units while a truthful bid was at around 11.72 monetary units on average. Remember that ex post regret assumes that all other bidders are truthful, and it needs to be traded off against the risk of losing when a bidder is shading his bid too aggressively. In contrast, the core-price computation makes sure that a bidder only pay as much as necessary for a stable solution but he cannot pay too much as it can easily happen with a pay-as-bid rule.

	mean	st. dev.	max.
Core (OPT)	1.87	2.02	7.74
Core (Core0.1)	2.18	2.66	9.48

Table 9 Ex post regret for the core pricing treatments (in monetary units).

In Table 10, we report the results of a regression analysis with overall revenue of the auctions and average payments of the winners as dependent variables. We again control for the number of carriers, the number of warehouses per carrier, and the warehouse capacity. Table 10 shows the respective regression coefficients with truthful valuations or welfare based on the truthful valuations as a baseline. All covariates in the regressions were significant at $p < 0.01$.

Remember that revenue maximization is not a goal in this environment and we want to use payments only to find an efficient allocation of time slots to those with the highest opportunity costs. The average payment in a pay-as-bid auction with truthful bidders as the baseline was 11.72 monetary units. The estimated expected payment in the RaR framework was only 11.72-11.68=0.04 monetary units, which is largely due to the low number of winners in expectation (high probability of the empty allocation). In contrast, the estimated payoff in the Core (OPT) treatment was only 3.13 monetary units, and 2.18 monetary units in the core-selecting auction with near-optimal solutions.

The average Vickrey discount in the VCG mechanism was higher than in the core-selecting auction (5.01 monetary units), as one would expect. The expected payments of winners and also the revenue in the RaR framework was lower than in VCG and core-selecting auctions. This can be explained by the fact that the fractionally optimal solution is scaled down, leading to a lower number of expected winners and therewith lower welfare, which then translates into longer waiting times compared to the optimal solution.

	⌀ Revenue	⌀ Payments	Significance
intercept	-54.42	6.52	$p < 0.010$
number of carriers	4.62	0.12	$p < 0.001$
warehouses per carrier	14.08	0.79	$p < 0.001$
warehouse capacity	46.77	-1.54	$p < 0.001$
Mechanisms			
Core (OPT)	-65.81	-3.13	$p < 0.001$
VCG (OPT)	-102.53	-5.01	$p < 0.001$
Core (Core0.1)	-51.13	-2.18	$p < 0.001$
RaR B&C0.1	-127.68	-8.49	$p < 0.001$
RaR B&C1.0	-156.33	-9.41	$p < 0.001$
RaR	-223.69	-11.68	$p < 0.001$
Adjusted R^2	0.70	0.78	

Table 10 Mean revenue and payments per bidder compared to the truthful valuations as a baseline (in monetary units).

While our experiments are representative for a relevant class of transportation problems, there are many different environments in transportation logistics. It is impossible to cover all possible logistics scenarios. However, our setting nicely illustrates the differences among the coordination mechanisms and the level of waiting time savings one can expect for these selected environments.

The experimental results yield a robust ranking of auction mechanisms in terms of waiting time reductions and computation times across all experiment parameters.

4. Conclusions

Congestion at loading docks of retail warehouses are a substantial problem in retail transportation logistics and an example of coordination problems as they often arise in supply chain management. The problem is due to the fact that carriers optimize locally, but there is no coordination among the carriers leading to globally suboptimal allocations of warehouse capacities. Adding capacity with additional loading docks at warehouse sites requires substantial investments and can also be infeasible in urban and densely populated areas.

Auction-based coordination mechanisms allow to elicit information from carriers about their preferences for different routes. If truthful bidding is an equilibrium strategy, this avoids high bid preparation costs and inefficiencies and it makes bidding easy for participants. Carriers cannot benefit from manipulating their reported preferences. This is also important for the acceptance and perceived fairness of a mechanism in the field. The VCG mechanism is the unique strategy-proof mechanism, but it assumes that the allocation problem can be solved to optimality. The allocation problem in retail logistics is a computationally hard problem. Although we cannot expect to always solve such problems to optimality, near-optimal solutions with a low MIP gap are typically achievable for realistic problem sizes.

In this paper, we explore recent approaches from the literature on mechanism and market design, which provide simple truthful bidding strategies even in the presence of computationally hard allocation problems. This includes VCG and core-selecting payment rules with near-optimal solutions, as well as the RaR framework by Lavi and Swamy (2011). The latter can be seen as a significant contribution in the literature on algorithmic mechanism design, as it allows to turn any approximation algorithm for a packing problem into a mechanism, which is truthful in expectation. While the polynomial runtime and the worst-case solution quality have been proven, the average-case solution quality and applicability of the whole framework remained unexplored.

There is a fundamental trade-off between computational costs of a mechanism and reductions in the waiting times. While the computations could be done overnight, a computation time of more than 10 hours (for even small problem instances) might not be acceptable for a coordination mechanism in retail logistics. We ran extensive simulation experiments to better understand these trade-offs. We tried to make our numerical experiments as realistic as possible. Of course, every application in the field will have different parameters. However, the main results of this study generalize and provide clear guidance as to which approaches are applicable to larger markets, and which are not.

The results of the optimal computations and also the results of the various mechanisms show that waiting time reduction via a coordination mechanism can be substantial. It can significantly improve the efficiency of retail transportation logistics in practice. Independent of the problem specifics, our results suggest that the standard RaR framework suffers from low performance, because it essentially turns the worst-case approximation ratio of an approximation algorithm into its average-case ratio. We introduce an extension where the approximation algorithm in the RaR framework is replaced by a branch-and-cut algorithm, which terminates the computations after a certain MIP gap is reached. While the mechanism does not run in polynomial time anymore, we leverage the fact that for many realistic problem-sizes, modern MIP solvers can find near-optimal solutions very fast. More importantly, this extension allows to use the RaR framework in all types of packing problems, while the standard approach is bound to the limited set of problems where an approximation algorithm is known.

The determination of an appropriate MIP gap requires some experiments with realistic bid sets. In our experiments, we used a very ambitious MIP gap of 10% and one of 100%. The MIP gap of 10% achieved a high solution quality, but the computational costs turned out to be high, because many optimization problems need to be solved in a sequence. A MIP gap of 100% could be solved very fast, at the expense of solution quality. Still, the advantage of the RaR B&C framework is that truthful bidding is a simple equilibrium strategy.

If one is willing to give up truthfulness in expectation as a solution concept, then the computation of VCG or core payments for near-optimal allocations can be a practical solution. We show that allocation and payments can be computed very effectively even for large instances and that the prices satisfy a certain notion of fairness, the core property. Arguably, the possibilities for manipulation would require so much information about the valuations of competitors that it becomes impractical in our scenario.

Overall, recent advances in optimization and mechanism design allow for effective coordination and provide new and promising approaches to solve daunting coordination problems in retail logistics and possibly other applications in supply chain management that have not been available only a few years ago.

References

Adomavicius, G., S. Curley, A. Gupta, P. Sanyal. 2012. Effect of information feedback on bidder behavior in continuous combinatorial auctions. *Management Science* **58**(4) 811–830.

Ausubel, L. M., P. Milgrom. 2002. Ascending auctions with package bidding. *Frontiers of Theoretical Economics* **1** 1–42.

Ausubel, L. M., P. Milgrom. 2006. The lovely but lonely vickrey auction. *Combinatorial auctions*, chap. 1. MIT Press, 57–95.

Banker, Rajiv D., Robert J. Kauffman. 2004. 50th anniversary article: The evolution of research on information systems: A fiftieth-year survey of the literature in management science. *Management Science* **50**(3) 281–298.

Bichler, Martin, J Goeree. 2017. *Handbook of Spectrum Auction Design*. Cambridge University Press.

Bichler, Martin, Alok Gupta, Wolfgang Ketter. 2010. Research commentary–designing smart markets. *Information Systems Research* **21**(4) 688–699.

Bikhchandani, S. l., J. M. Ostroy. 2002. The package assignment model. *Journal of Economic Theory* **107**(2) 377–406.

Bland, Robert G, Donald Goldfarb, Michael J Todd. 1981. The ellipsoid method: A survey. *Operations Research* **29**(6) 1039–1091.

Briest, P., P. Krysta, B. Vöcking. 2011. Approximation techniques for utilitarian mechanism design. *SIAM Journal on Computing* **40**(6) 1587–1622.

Bundesamt für Güterverkehr. 2011. Sonderbericht zur Situation an der Laderampe. URL `http://www.bag.bund.de/SharedDocs/Downloads/DE/Marktbeobachtung/Sonderberichte/SB_Rampe.pdf`.

Bundesverband Güterkraftverkehr Logistik und Entsorgung e.V. 2013. Jahresbericht 2012/2013. URL `http://www.bgl-ev.de/images/downloads/ueber/jahresbericht/bgl-jahresbericht.pdf`.

Carr, Robert, Santosh Vempala. 2000. Randomized metarounding. *Proceedings of the thirty-second annual ACM symposium on Theory of computing*. ACM, 58–62.

Clarke, E.H. 1971. Multipart pricing of public goods. *Public Choice* **11** 17–33.

Conitzer, V., T. Sandholm. 2006. Complexity of constructing solutions in the core based on synergies among coalitions. *Artificial Intelligence* **170** 607–619.

Day, R. W., P. Milgrom. 2008. Core-selecting package auctions. *International Journal of Game Theory* **36**(3) 393–407.

Day, R. W., S. Raghavan. 2007. Fair payments for efficient allocations in public sector combinatorial auctions. *Management Science* **53**(9) 1389–1406.

Deng, Xiaotie, Toshihide Ibaraki, Hiroshi Nagamochi. 1999. Algorithmic aspects of the core of combinatorial optimization games. *Mathematics of Operations Research* **24**(3) 751–766.

Elbert, R., D. Thiel, D. Reinhardt. 2016. Delivery time windows for road freight carriers and forwarders - influence of delivery time windows on the costs of road transport services. *Commercial Transport - Proceedings of the 2nd Interdisciplinary Conference on Production, Logistics and Traffic 2015*. Springer, 255–274.

Goeree, J., Y. Lien. 2016. On the impossibility of core-selecting auctions. *Theoretical Economics* **11**(1) 41–52.

Goetzendorff, A., M. Bichler, P. Shabalin, R. W. Day. 2015. Compact bid languages and core pricing in large multi-item auctions. *Management Science* **61**(7) 1684–1703.

Groves, T. 1973. Incentives in teams. *Econometrica* **41** 617–631.

Lavi, Ron, Chaitanya Swamy. 2011. Truthful and near-optimal mechanism design via linear programming. *Journal of the ACM* **58**(6) 1–24.

Lubin, Benjamin, David C Parkes. 2012. Approximate strategyproofness. *Current Science* 1021–1032.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* **16**(1) 8–37.

Appendix A: Proofs

Proof of Proposition 1: For any agent $i \in I$, true valuations $v \in V$, and reported valuation $w \in V$, let $a = f(v_i, w_{-i})$ denote the outcome for the agent's true valuation while $b = f(w)$ denotes the outcome with respect to the reported valuation. Since the social choice function maximizes over the reported valuations, it holds that $v_i(a) + \sum_{j \in I\setminus\{i\}} w_j(a) \geq v_i(b) + \sum_{j \in I\setminus\{i\}} w_j(b)$. Thus, agent i cannot gain utility by misrepresenting the valuation

$$\begin{aligned} v_i(f(v_i, w_{-i})) - p_i^{VCG}(v_i, w_{-i}) &= v_i(a) - \Big(h_i(w_{-i}) - \sum_{j \in I\setminus\{i\}} w_j(a)\Big) \\ &\geq v_i(b) - \Big(h_i(w_{-i}) - \sum_{j \in I\setminus\{i\}} w_j(b)\Big) \\ &= v_i(f(w)) - p_i^{VCG}(w). \end{aligned}$$

Q.E.D.

Proof of Proposition 2: We show that the WDP in the logistics problem satisfies the requirements of the RaR framework introduced in section 2.3. First, we begin with the packing property. Let $x, x' \in [0,1]^{|I| \times 2^{|H|}}$ be two, possibly fractional, allocations. We say that x' dominates x if $x_i \leq x'_i$ for all carriers i. Allocation x can be created by partially reducing x'. Since this neither violates the warehouse capacities nor the XOR relation of bids, WDP satisfies the packing property. Secondly, the integrality gap can be bounded by the polynomial-time approximation algorithm by Briest et al. (2011) with a worst-case approximation ratio in $O((K \cdot T)^{\frac{1}{c+1}})$. Q.E.D.

Proof of Corollary 1: The allocation rule $f^D(v)$ of the deterministic support mechanism $\mathcal{M}^D$ results in a distribution $\{\lambda_l\}_{l \in \mathbb{I}}$. In order to get feasible integer solutions $\{x^l\}_{l \in \mathbb{I}}$ that satisfy all constraints of the WDP, we need to scale down the fractional solution by at least α, the approximation ratio of the approximation algorithm $\mathcal{A}$ employed. Observe that the value of a bidder $v \in V$ increases linearly in x, and a scaled down version of the fractional VCG mechanism $\mathcal{M}^F$ is still truthful. The distribution $\{\lambda_l\}_{l \in \mathbb{I}}$ is determined via a convex combination of a scaled down fractional solution to WDP into $\sum_{l \in \mathbb{I}} \lambda_l x^l$ of integral solutions where $\lambda_l(v) \geq 0$, and $\sum_l \lambda_l = 1$. Lemma 3.1 by Lavi and Swamy (2011) proves that the support mechanism $\mathcal{M}^D$ is truthful if the convex decomposition exactly implements the scaled down fractional solution. As a result, $v_i(\sum_{l \in \mathbb{I}} \lambda_l(v) x^l) = v_i(\frac{x^*(v)}{\alpha}) = \frac{v_i(x^*(v))}{\alpha}$ such that the expected welfare of $\mathcal{M}^D$ is equivalent to the outcome of the scaled down fractional VCG solution $\frac{x^*}{\alpha}$, which is truthful. Consequently, also the expected total welfare achieved by the randomized mechanism $\mathcal{M}^R$ is exactly equivalent to the one of the fractional VCG mechanism scaled down by α. Q.E.D.

Proof of Corollary 2: The proof follows Theorem 3.1 by Lavi and Swamy (2011). Note that if the time slots could be allocated fractionally, then one could use the VCG payment rule directly and the resulting mechanism was dominant-strategy incentive-compatible. We refer to the allocation rule of this fractional VCG mechanism $\mathcal{M}^F$ as f^F and the payments as p^F. Now, we can define a δ-scaled fractional VCG mechanism with an allocation function that outputs the outcome $\frac{f^F(v)}{\delta}$ and prices $\frac{p^F(v)}{\delta}$. Such a scaled version of the fractional VCG mechanism would and still be truthful, because the scaling of the outcome, valuations, and payments does not impact truthfulness of the VCG mechanism in a fractional domain. Now, we distinguish two cases.

First, we assume that δ is large enough such that we can decompose $\frac{x^*(v)}{\delta}$ into a convex combination of feasible integer solutions $\{x^l\}_{l\in\mathbb{I}}$. If this is the case, we compute $v_i(\sum_{l\in\mathbb{I}}\lambda_l(v)x^l) = v_i(\frac{x^*(v)}{\delta}) = \frac{v_i(x^*(v))}{\delta}$. This yields the deterministic support mechanism $\mathcal{M}^D$ with allocation rule $f^D(v) = \{\lambda_l(v)\}_{l\in\mathbb{I}}$ with prices $p_i^D(v) = \frac{p_i^F(v)}{\delta}$. Truthfulness follows from the equivalence of $\mathcal{M}^D$ and the scaled-down $\mathcal{M}^F$, which is achieved through the exact decomposition of the scaled-down fractional solution $\frac{x^*(v)}{\delta}$ and the truthfulness of $\mathcal{M}^F$. If we randomly select a feasible integer solution x^l according to the distribution $\{\lambda_l\}_{l\in\mathbb{I}}$ we get a randomized support mechanism $\mathcal{M}^R$ with the same expected outcome, which is truthful in expectation.

Second, if the scaled-down fractional solution $\frac{x^*(v)}{\delta}$ cannot be decomposed, we increase δ and repeat the procedure until a convex decomposition is possible. For an increased δ, the mechanism is still equivalent to the outcome of a scaled version of the VCG mechanism in expectation. Hence, increasing δ does not impact truthfulness. Q.E.D.

Proof of Proposition 3: Corollary 2 shows truthfulness in expectation of $\mathcal{M}^R$. What remains to be shown is that we get the desired δ-approximation. In order to maintain truthfulness, we need to compute the convex decomposition such that it exactly implements the optimal fractional allocation scaled down by the achievable duality gap δ as shown in Corollary 2. The decomposition is computed via the linear program (P) which minimizes $\sum_{l\in\mathbb{I}}\lambda_l$ subject to the exact decomposition $\sum_l \lambda_l x_{iS}^l = \frac{x_{iS}^*}{\delta}$, $\lambda_l \geq 0$, and $\sum_l \lambda_l \geq 1$ (see Appendix B). The dual (D) of this linear program has exponentially many constraints, but can be solved efficiently using the ellipsoid method.

In each iteration of the ellipsoid method, we need to determine whether a vector x^l is in the constraint polytope or not. The RaR B&C framework uses a branch-and-cut algorithm $\mathcal{B}$ solving the WDP up to a given duality gap δ as a separation oracle for this purpose. The ellipsoid method solves a problem equivalent to (D) with a polynomial number of constraints. These constraints represent violated inequalities returned by the separation oracle. The current center of the ellipsoid (ν, z) is treated as valuation vector for $\mathcal{B}$. If the objective value of (D) $\frac{1}{\delta}\sum_{(i,S)\in E} x_{iS}^*\nu_{iS} + z > 1$, then we can find a violated constraint for an x^l using Claim 3.3 by Lavi and Swamy (2011); the objective value is used to cut the current ellipsoid otherwise. When we can solve (D) with a polynomial number of constraints, we have a solution to the primal (P) with a polynomial number of variables $\lambda_l > 0$, which now yields the exact convex decomposition of the optimal solution x^* scaled down by δ.

Q.E.D.

Proof of Proposition 4: Regret describes how much a bidder can gain by deviating unilaterally, i.e. $(u_i(w_i, w_{-i}) - u_i(v_i, w_{-i}))$. We show that in a bidder-optimal core-selecting auction, this benefit can at most be $p_i^{core}(v_i, w_{-i}) - p_i^{VCG}(v_i, w_{-i})$ and prove by contradiction. Suppose, there is a report w_i with ex post regret $(u_i(w_i, w_{-i}) - u_i(v_i, w_{-i}))$ strictly larger than $p_i^{core}(v_i, w_{-i}) - p_i^{VCG}(v_i, w_{-i})$. We can rewrite the term as

$$(u_i(w_i, w_{-i}) - u_i(v_i, w_{-i})) \tag{3a}$$

$$= (v_i(f(w_i, w_{-i})) - p_i^{core}(w_i, w_{-i})) - (v_i(f(v_i, w_{-i})) - p_i^{core}(v_i, w_{-i})) \tag{3b}$$

$$> p_i^{core}(v_i, w_{-i}) - p_i^{VCG}(v_i, w_{-i}). \tag{3c}$$

Next, we can rearrange terms and get $v_i(f(w_i, w_{-i})) - p_i^{core}(w_i, w_{-i}) > v_i(f(v_i, w_{-i})) - p_i^{VCG}(v_i, w_{-i})$. The core computation is such that we can never have a core payoff larger than the VCG payoffs, i.e., $v_i(f(w_i, w_{-i})) - p_i^{VCG}(w_i, w_{-i}) \geq v_i(f(w_i, w_{-i})) - p_i^{core}(w_i, w_{-i})$. Consequently, $v_i(f(w_i, w_{-i})) - p_i^{VCG}(w_i, w_{-i}) > v_i(f(v_i, w_{-i})) - p_i^{VCG}(v_i, w_{-i})$. However, we know from Proposition 1 that this cannot be true. Q.E.D.

Appendix B: Convex Decomposition

In this section, we describe the implementation of the convex decomposition, which is arguably the central piece of the RaR framework. For our logistics problem, we relax the integrality constraint of WDP to $0 \leq x_{iS} \leq 1$. Remember, a variable $x_{iS} \in \{0, 1\}$ is used for each set of items S of a bidder i, and a bidder can win only one of his bids. The LP relaxation can be solved using any algorithm which finds an exact solution and this is possible in polynomial time. An optimal fractional solution to the relaxed WDP is obtained which is maximizing social welfare. Next, the convex decomposition of the fractional solution scaled down by α is computed. We denote $\mathcal{P}$ as the feasible region of the WDP and $\mathbb{Z}(\mathcal{P}) = \{x^l\}_{l \in \mathbb{I}}$ as the set of all feasible integer solutions of the relaxed WDP. The index set for the integer solutions is represented by $\mathbb{I}$. The optimal fractional solution is denoted as x^*. We introduce $E = \{(i, S) : x^*_{iS} > 0\}$ as the set of all allocations with positive value that are part of the optimal solution x^*. This fractional solution x^* is scaled down by the integrality gap of the problem α. Then the α-approximation algorithm $\mathcal{A}$ is used to sample integer solutions x^l around $\frac{x^*}{\alpha}$, and express $\frac{x^*}{\alpha}$ as a convex combination of these integer solutions, $\sum_{l \in \mathbb{I}} \lambda_l x^l$, with $\lambda_l \geq 0$ and $\sum_{l \in \mathbb{I}} \lambda_l = 1$. Since the factors λ_l are treated as probabilities over the integer solutions, it is necessary that they add up to 1. The solution of the following LP provides the desired convex combination:

$$\begin{aligned}
\min \quad & \sum_{l \in \mathbb{I}} \lambda_l & & \text{(P)} \\
\text{s.t.} \quad & \sum_{l} \lambda_l x^l_{iS} = \frac{x^*_{iS}}{\alpha} & \forall (i, S) \in E & \\
& \sum_{l} \lambda_l \geq 1 & & \\
& \lambda_l \geq 0 & \forall l \in \mathbb{I} &
\end{aligned}$$

Unfortunately, the objective function is summing over all integer solutions. Thus, we need an exponential number of variables and cannot solve the LP efficiently. To overcome this problem, one can convert the variables to an exponential number of constraints by taking the dual. The dual LP can be solved with the ellipsoid method in polynomial time, because it can handle an exponential number of constraints using a separation oracle. This uses the α-approximation algorithm $\mathcal{A}$ to find integer solutions by treating the current center of the ellipsoid as valuation vector. The dual is defined as follows:

$$\max \quad \frac{1}{\alpha} \cdot \sum_{(i,S) \in E} x^*_{iS} \nu_{iS} + z \tag{D}$$

$$\text{s.t.} \quad \sum_{(i,S)\in E} x^l_{iS}\nu_{iS} + z \le 1 \qquad \forall l \in \mathbb{I}$$

$$z \ge 0$$

$$\nu_{iS} \text{ unconstrained} \qquad \forall (i,S) \in E$$

After solving the dual, a polynomial number of integral solutions is available and therefore the solution to the primal LP can be obtained in polynomial time (Lavi and Swamy 2011). The result is the desired convex combination or probability distribution over integer solutions, respectively. Finally, it is necessary to compute the payments for each of the integer solutions using the expected VCG payments $E\left[p_i^R(v)\right]$ as follows. For every possible outcome $\{x^l : \lambda_l > 0\}_{l\in\mathbb{I}}$, we set $p_i^R(v) = v_i(x^l) \cdot \frac{E\left[p_i^R(v)\right]}{E\left[v_i(f^R(v))\right]}$ if $E\left[v_i(f^R(v))\right] > 0$ and $p_i^R(v) = 0$ if $E\left[v_i(f^R(v))\right] = 0$. The expected payments and valuations can be obtained by scaling down the respective results of the LP relaxation by α.

Appendix C: Core price computation (ONLINE ONLY)

Let $W \subseteq I$ denote the set of winners, $S^* = (S_1^*, \ldots, S_{|I|}^*)$ the allocation that maximizes social welfare, $p^{VCG} = (p_1^{VCG}, \ldots, p_{|I|}^{VCG})$ the VCG payment vector, $p^\tau = (p_1^\tau, \ldots, p_{|I|}^\tau)$ the payment vector, and $C^\tau \subseteq I$ the blocking coalition of bidders in iteration τ.

Following Day and Raghavan (2007), we define the core separation problem, which yields the most violated core constraint, if any. S_i^* denotes the set of items that is allocated to bidder i from the set of winners $W \subseteq I$. To calculate equitable bidder-Pareto optimal (EBPO) core payments iteratively, the procedure is as follows. Solve the core separation problem (SEP$^\tau$):

$$z(p^\tau) = \max \sum_{i \in I} \sum_{S \neq \emptyset} v_i(S) x_{iS}^\tau - \sum_{i \in W} (v_i(S_i^*) - p_i^\tau) \gamma_{iS}^\tau \qquad \text{(SEP}^\tau\text{)}$$

$$\text{s. t.} \sum_{i \in I} \sum_{S: r_{kt} \in S} x_{iS}^\tau \leq c_{kt} \qquad \forall k \in K, t \in T, S \neq \emptyset, \qquad (4a)$$

$$\sum_{S \neq \emptyset} x_{iS}^\tau \leq 1 \qquad \forall i \in I \setminus W, \qquad (4b)$$

$$\sum_{S \neq \emptyset} x_{iS}^\tau \leq \gamma_{iS}^\tau \qquad \forall i \in W, \qquad (4c)$$

$$x_{iS}^\tau \in \{0, 1\} \qquad \forall i \in I, S \neq \emptyset, \qquad (4d)$$

$$\gamma_{iS}^\tau \in \{0, 1\} \qquad \forall i \in W, S \neq \emptyset. \qquad (4e)$$

SEP$^\tau$ yields the most violated core constraint, if any. That is, it finds coalitions of bidders C^τ that block the current outcome (who "would pay more" than the current payments). γ_{iS}^τ is a binary variable, which ensures that any winning bidder will be compensated his opportunity cost if selected as part of the optimal solution to SEP$^\tau$. Let $p^{core,\tau} = (p_1^{core,\tau}, \ldots, p_{|I|}^{core,\tau})$ denote the (temporary) core payment vector in iteration τ. Then, the minimal payments in the core satisfying the core constraints found (EBPOCORE) are calculated in EBPO$^\tau$:

$$\theta^\tau(\epsilon) = \min \sum_{i \in W} p_i^{core,\tau} + \epsilon m^\tau \qquad \text{(EBPO}^\tau\text{)}$$

$$\text{s. t.} \sum_{i \in W \setminus C^{\tau'}} p_i^{core,\tau} \geq z(p^{\tau'}) - \sum_{i \in W \cap C^{\tau'}} p_i^{\tau'} \qquad \forall \tau' \leq \tau, \qquad \text{(EBPOCORE)}$$

$$p_i^{core,\tau} - m^\tau \leq p_i^{VCG} \qquad \forall i \in W, \qquad (5a)$$

$$p_i^{core,\tau} \leq v_i(S_i^*) \qquad \forall i \in W, \qquad (5b)$$

$$p_i^{core,\tau} \geq p_i^{VCG} \qquad \forall i \in W. \qquad (5c)$$

The minimal payments minimize potential gains from deviation and EBPO minimizes the maximum deviation from VCG payments as a secondary objective. Using a sufficiently small value of ϵ, the calculated sum of payments is minimal in the core. This procedure is repeated until no further core constraint violation is found using (SEP$^\tau$), i.e., it is repeated while $z(p^\tau) > \theta^{\tau-1}(\epsilon)$ with $\theta^0(\epsilon) := \sum_i p_i^{\text{VCG}}$.

Now, Algorithm 1 shows how core payments are computed in pseudo code. First, the winner determination problem is solved ($\omega(I)$) and the VCG prices are calculated for the winners. Next, the core separation problem is solved to find the most violated core constraint. Then, the minimal payments in the core satisfying the core constraints found are calculated. This procedure is repeated until no further core constraint violation is found.

Algorithm 1: Core constraints generation (following Day and Raghavan (2007, p. 1398))

$W, S^* \leftarrow$ solve the winner determination problem $\omega(I)$
foreach $i \in W$ **do**
  $p_i^{\text{VCG}} \leftarrow$ compute the VCG price $v_i(S_i^*) - \big(\omega(I) - \omega(I_{-i})\big)$
$p^1 \leftarrow p^{\text{VCG}}$
$\theta^0(\epsilon) \leftarrow \sum_i p_i^{\text{VCG}}$
$\tau \leftarrow 1$
while *true* **do**
  $C^\tau \leftarrow$ solve the separation problem SEP^τ
  if $z(p^\tau) > \theta^{\tau-1}(\epsilon)$ **then**
    add constraint $\sum_{i \in W \setminus C^\tau} p_i^{core,\tau} \geq z(p^\tau) - \sum_{i \in W \cap C^\tau} p_i^\tau$ to EBPO^τ and solve
    $p^{\tau+1} \leftarrow p^{core,\tau}$ from EBPO^τ
  else
    $p \leftarrow p^\tau$
    break
  $\tau \leftarrow \tau + 1$

Core payments suffer from sub-optimal computations of the WDP. For example, it might well be that $\omega(I_{-i}) > \omega(I)$, which would then lead to $p_i^{VCG} > v_i(S_i^*)$. The algorithms by Goetzendorff et al. (2015) trim payments or updates the winning allocation, in cases a better one is found during the computation of payments. It was proposed for markets where it is possible to compute near-optimal allocations quickly (measured by the integrality gap of the MIP solver), but where proving optimality turns out to be intractable. We use the TRIM algorithm by Goetzendorff et al. (2015) in our experiments to deal with near-optimal allocations.

Appendix D: Detailed Results (ONLINE ONLY)

Each line in Table 11 describes a treatment combination consisting of the number of carriers, the warehouse capacity, and mechanism. We average over number of warehouses per tour. Computation times, seller revenue, winning bidder payments, and the numbers of submitted and winning bids are reported in a similar way in Tables 12 and 13.

Carriers	Capacity	Auction	⌀ Waiting Time per tour	⌀ RTT
20	1	OPT	89.98	449.09
20	1	RaR	111.90	463.48
20	1	RaR B&C MG0.1	99.55	456.33
20	1	RaR B&C1.0	103.02	458.65
20	1	Core0.1	90.64	449.25
20	1	no coord.	113.09	464.69
20	2	OPT	4.50	361.58
20	2	RaR	17.40	369.00
20	2	RaR B&C0.1	8.64	363.80
20	2	RaR B&C1.0	12.31	366.06
20	2	Core0.1	5.48	362.35
20	2	no coord.	18.27	369.88
30	1	OPT	195.72	552.10
30	1	RaR	227.42	578.90
30	1	RaR B&C0.1	206.49	561.92
30	1	RaR B&C1.0	211.33	565.98
30	1	Core0.1	198.48	554.41
30	1	no coord.	225.86	577.31
30	2	OPT	22.06	382.96
30	2	RaR	45.99	397.57
30	2	RaR B&C0.1	29.55	389.14
30	2	RaR B&C1.0	36.55	393.09
30	2	Core0.1	26.14	387.17
30	2	no coord.	45.64	397.09
40	1	OPT	301.97	657.97
40	1	RaR	347.29	699.48
40	1	RaR B&C0.1	315.01	670.67
40	1	RaR B&C1.0	321.32	676.37
40	1	Core0.1	304.95	660.83
40	1	no coord.	348.13	700.32
40	2	OPT	64.71	426.12
40	2	RaR	88.95	441.24
40	2	RaR B&C0.1	68.97	428.56
40	2	RaR B&C1.0	77.52	434.41
40	2	Core0.1	67.03	426.91
40	2	no coord.	89.82	442.01

Table 11 **Mean waiting round trip times per tour(in minutes).**

Carriers	Capacity	Auction	⌀ Time	Max. time	⌀ Revenue	⌀ Payment
20	1	CORE	15.85	76.66	123.86	10.34
20	1	Pay-as-bid	0.60	1.78	140.68	11.69
20	1	RaR	0.33	0.58	0.44	0.02
20	1	RaR B&C0.1	7681.99	18608.53	68.10	3.63
20	1	RaR B&C1.0	24.73	74.78	55.45	2.95
20	1	Core0.1	5.39	20.54	124.08	10.56
20	1	VCG	7.28	23.42	89.60	7.61
20	2	CORE	2.23	7.05	6.75	0.34
20	2	Pay-as-bid	0.09	0.27	236.74	11.84
20	2	RaR	0.22	0.49	0.05	0.00
20	2	RaR B&C0.1	2805.83	7836.91	3.29	0.16
20	2	RaR B&C1.0	7.62	22.83	1.93	0.10
20	2	Core0.1	0.29	0.54	65.60	3.44
20	2	VCG	1.14	2.74	5.62	0.28
30	1	CORE	121.68	439.96	139.93	10.78
30	1	Pay-as-bid	6.04	20.62	152.28	11.70
30	1	RaR	0.65	1.19	0.51	0.02
30	1	RaR B&C0.1	12188.96	17611.05	78.29	2.81
30	1	RaR B&C1.0	78.94	202.53	62.77	2.25
30	1	Core0.1	68.15	246.63	141.22	11.11
30	1	VCG	77.60	257.68	104.12	8.11
30	2	CORE	501.92	2677.93	246.75	9.24
30	2	Pay-as-bid	8.65	46.60	313.73	11.73
30	2	RaR	0.60	1.10	2.52	0.08
30	2	RaR B&C0.1	10826.56	27630.60	159.25	5.33
30	2	RaR B&C1.0	55.46	118.26	101.41	3.39
30	2	Core0.1	17.84	128.23	259.03	10.15
30	2	VCG	236.35	1453.47	180.72	6.80
40	1	CORE	275.63	858.31	153.77	11.22
40	1	Pay-as-bid	18.92	108.09	160.43	11.71
40	1	RaR	1.18	2.17	0.59	0.02
40	1	RaR B&C0.1	17666.27	37750.68	88.40	2.49
40	1	RaR B&C1.0	194.83	428.78	72.31	2.04
40	1	Core0.1	173.64	663.66	153.92	11.30
40	1	VCG	196.06	702.52	138.89	10.16
40	2	CORE	17823.52	103834.17	283.29	9.56
40	2	Pay-as-bid	174.49	914.03	345.35	11.63
40	2	RaR	1.16	2.14	2.98	0.08
40	2	RaR B&C0.1	22345.39	50618.55	185.83	4.92
40	2	RaR B&C1.0	120.84	220.31	117.35	3.10
40	2	Core0.1	124.47	592.48	298.57	10.62
40	2	VCG	5704.92	27836.04	215.12	7.29

Table 12 **Computation times (seconds), auctioneer revenue, and winning bidder payments (in monetary units). "Pay-as-bid" describes the values for the winner determination based on true valuations.**

Carriers	Capacity	Auction	∅ no bids	∅ no. win. bids	∅ no. bids p. carrier	∅ no. win. bids p. carrier
20	1	OPT	170.20	12.10	8.51	0.60
20	1	RaR	170.20	0.06	8.51	0.00
20	1	RaR B&C0.1	170.20	8.51	8.51	0.43
20	1	RaR B&C1.0	170.20	6.68	8.51	0.33
20	1	Core0.1	170.20	11.80	8.51	0.59
20	2	OPT	170.00	20.00	8.50	1.00
20	2	RaR	170.00	0.25	8.50	0.01
20	2	RaR B&C0.1	170.00	16.45	8.50	0.82
20	2	RaR B&C1.0	170.00	10.00	8.50	0.50
20	2	Core0.1	170.00	19.50	8.50	0.97
30	1	OPT	255.20	13.10	8.51	0.44
30	1	RaR	255.20	0.06	8.51	0.00
30	1	RaR B&C0.1	255.20	9.65	8.51	0.32
30	1	RaR B&C1.0	255.20	7.64	8.51	0.25
30	1	Core0.1	255.20	12.80	8.51	0.43
30	2	OPT	255.30	26.80	8.51	0.89
30	2	RaR	255.30	0.35	8.51	0.01
30	2	RaR B&C0.1	255.30	22.02	8.51	0.73
30	2	RaR B&C1.0	255.30	13.84	8.51	0.46
30	2	Core0.1	255.30	25.60	8.51	0.85
40	1	OPT	342.40	13.80	8.56	0.35
40	1	RaR	342.40	0.07	8.56	0.00
40	1	RaR B&C0.1	342.40	10.39	8.56	0.26
40	1	RaR B&C1.0	342.40	8.37	8.56	0.21
40	1	Core0.1	342.40	13.70	8.56	0.34
40	2	OPT	342.20	29.80	8.56	0.74
40	2	RaR	342.20	0.39	8.56	0.01
40	2	RaR B&C0.1	342.20	24.65	8.56	0.62
40	2	RaR B&C1.0	342.20	15.51	8.56	0.39
40	2	Core0.1	342.20	28.20	8.56	0.70

Table 13 Average number of bids and winners.

Acknowledgments

The financial support from the TUM Institute for Advanced Study is gratefully acknowledged.