



Ô[{ àã æ | ãé &ç } Å^• ã }
{ !Á^æË•cæ^Á æ\^o

Ö:a•Ö[[••^}•ËÙæ å^!Á} å^!•cæËæ Áü æ\^!Áæ åÁü•Á] ã\•{ æ



DEPARTMENT OF DECISION SCIENCES AND INFORMATION MANAGEMENT (KBI)

Combinatorial auction design for real-estate markets*

Dries R. Goossens^(1,2), Sander Onderstal⁽³⁾,
Jan Pijnacker⁽⁴⁾, Frits C.R. Spieksma⁽²⁾

(1) PostDoc researcher for Research Foundation - Flanders

Corresponding author. E-mail: Dries.Goossens@econ.kuleuven.be

(2) Center for Operations Research and Business Statistics (ORSTAT),
Faculty of Business and Economics, KU Leuven, Belgium

(3) University of Amsterdam, The Netherlands

(4) Stadgenoot, The Netherlands

Abstract

This paper describes a combinatorial auction design for real-estate markets where floor space (in a multi-storeyed building) is to be allocated to various interested parties. Our design is based on lab experiments that provide guidance in choices regarding pricing, feedback, and activity rules. In addition, the allocation needs to satisfy several municipal and building regulations. We show how these regulations can be included in an integer program that is used to solve the winner determination problem, and discuss its computational complexity. Finally, we report on a practical application of this design, where over one hundred bidders took part in the first combinatorial auction for housing space, in a newly erected building in Amsterdam (the Netherlands).

Keywords: combinatorial auction, real estate, auction design,
integer programming

*An extended abstract based on a preliminary version of this paper is published in the proceedings of the Second Conference on Auctions, Market Mechanisms and Their Applications (AMMA), Lecture Notes of the Institute for Computer Science (2012), vol. 80:76-87

1 Introduction

Traditionally, real estate is sold as follows: the characteristics of the property are announced; the owner sets an ask price, and waits for an interested buyer to show up; the actual selling price is determined in a private negotiation. In the last decades, auctions have gradually become more popular in real-estate markets across the world. At first, real-estate auctions were most commonly used for disposing of real estate in bankruptcy cases, but more recent research shows that real-estate auctions have earned their place alongside private negotiation as a market mechanism (see e.g. Eklöf & Lunander (2003)).

In this work, we focus on markets where floor space in a multi-storeyed building is to be allocated to various interested parties. These interested parties can be very diverse, each with their own preferences and aspirations. They may well include real-estate brokers, but we assume the presence of a substantial number of non-professional, occasional participants. Usually, residences, business accommodations, shops, and offices are treated as separate markets, with different rules and target groups. In our setting, however, they can all compete for the same floor space. The floor space is divided into so-called *lots*, and buyers are interested in one or more adjacent lots, which we call a *space*. By choosing an appropriate set of lots (on a bidder's favorite floor, according to a preferred orientation), a bidder specifies the resulting space. However, there are various constraints on the allocation that must be taken into account. For instance, spaces may share utility shafts (electricity, water, heating, ...) and capacity limits on each shaft may exclude some allocations. Municipal regulations may also constrain the allocation, for instance when a given percentage of the floor space needs to be allocated to social housing. A detailed problem description is provided in Section 2.

Real-estate auctions typically involve a sequential and individual auctioning of

properties (Quan 1994). Occasionally, these auctions are followed by a second auction, where bidders can make a single bid on all the properties that were offered in the previous auctions; the seller then selects the auction with the highest payoff. Bayers (2000) describes a case where 14 units in a San Francisco condominium complex are auctioned simultaneously, but bidders are not allowed to submit package bids on sets of units. In our setting, a combinatorial auction was used. A combinatorial auction is an auction where multiple items are sold simultaneously, and bidders can express their preferences for (sub)sets of the items. This is quite useful, because typically, bidders are interested in multiple lots, and may value some sets of lots higher than the sum of the values of the lots individually. Moreover, some lots have no value at all as a single item, because e.g. they are not directly accessible from the hallway. These so-called complementarity effects may be bidder-specific, since bidders have different needs and preferences with respect to the space they want to rent. For instance, some bidders may want a spacious apartment with sun in the evening (west side), some bidders want a small but practical working space next to their shop on the ground level, and others prefer a cheap place for installing offices. A combinatorial auction allows bidders to express their preferences (including different budgets) to a greater extent than for individual items only. Combinatorial auction design forces trade-offs between desirable properties such as allocative efficiency, revenue maximization, transparency, fairness, and computational tractability (Pekeç & Rothkopf 2003). The Federal Communications Commission's (FCC's) request for auction proposals to sell spectrum rights triggered research on practical combinatorial auction design (see e.g. Goeree & Holt (2010), McMillan (1994)). Later, combinatorial auctions were designed and successfully implemented for procuring e.g. school meals (Epstein et al. (2002)), inputs for Mars (Hohner et al. (2003)), and transportation services (Ledyard et al. (2002)). For a thorough discussion on combinatorial auctions, we refer to the book edited by Cramton, Steinberg & Shoham (2005).

The contribution of this paper is a specific format for combinatorial auctions in real-estate markets. After reviewing several auction formats, we describe how we use lab experiments to specify precise auction rules (Section 3). We show how we can efficiently determine which bidder should be allocated which part of the building, taking into account various regulations (Section 4). As far as we are aware, the combinatorial auction we designed is the first to be used in practice for allocating real estate. Indeed, on May 7, 2011, over one hundred bidders participated in the auction of a building called Solid 11. This building offers 7000 square meters of floor space in the west part of Amsterdam (The Netherlands). We report on the specifics and the outcome of the auction of Solid 11 in Section 5.

2 The Concept

The problem is to develop an auction for allocating space in a multi-storeyed building to interested buyers, taking into account various constraints. A first step is to divide the building into lots; by choosing an appropriate set of lots (on a bidder's favorite floor, according to a preferred orientation), a bidder specifies the resulting space. As an example, Figure 1 shows how one floor of a building with two wings is divided into 22 lots; the surface area of each lot is given. Drawing a good allotment, however, is not a trivial task. Clearly, the allotment has an influence on the space a bidder can bid on, and hence, on the eventual allocation. There should be enough lots, with various characteristics and sizes, such that each bidder can compose a number of spaces that satisfy his preferences (throughout this paper, he can be replaced by she, and his by her). For instance, to accommodate bidders looking for a space not larger than 60 square meters, there should be enough small lots to form a space that meets their needs. It is also important to have enough substitutability across the lots, such that bidders can adapt the composition of their spaces depending on the

course of the auction. Having the same allotment on several floors is one way to accomplish this. On the other hand, having too many lots will result in a number of combinations that is no longer surveyable for the bidders, and increases the complexity the auction in every aspect. Therefore, if a pair of lots is likely to be seen as complementary by the vast majority of the bidders, it is better to combine these lots. Typically, the allotment will also be heavily constrained by building regulations and technical issues.




Figure 1: Allotment of the second floor of Solid 11

Bidders taking part in this auction may be quite diverse: residential bidders planning to live in the building, real-estate brokers looking for a bargain, commercial bidders who plan to open a business (hotel, offices, shop...), or even people with a low income looking for social housing. We assume that there are no externalities present in bidder's valuations. This means that the valuation of a bidder for a space depends only on the lots it contains. Thus, a bidder does not care which particular bidder will be his neighbor, or what this neighbor intends

to do with his space. We also assume that bidders have a good idea of what floor space in the building is worth, e.g. by comparing prices for similar real-estate in the neighborhood. A bidder cannot bid on just any subset of the lots, indeed, we only allow bidders to bid on sets of adjacent lots (and consequently, located on the same floor). This bears resemblance to the market for spectrum, oil drilling, and wind rights where adjacent geographic areas are often complementary. Apart from adjacency, there may be various other bid constraints: enough doors giving access to the central hallway, access to at least one utility shaft, sufficient light incidence, etc. For the remainder of this paper, when we use the term “space”, we mean in fact a set of lots satisfying these requirements.

Apart from restrictions on the sets of lots bidders can bid on, there are also constraints on the allocation itself. A bidder is allowed to provide multiple bids, however, he can win at most one bid. Indeed, most bidders can afford only one space, and do not want to end up with more. However, bidders who do want multiple spaces should be allowed to register in the auction using multiple identities (which allows them to win at most one space per identity). As an exception, bidders should be able to indicate that they want to see either all their bids allocated or none. This allows for instance a hotel owner to specify that he wants the top 3 floors, or nothing at all. Another type of constraints that the auction should accommodate is a given percentage of the surface that should (at least/at most) be allocated to a certain type of bidders (e.g. residential bidders), or similarly, limits to the number of spaces intended for a given purpose (e.g. bars). Other constraints originate from municipal and building regulations. Indeed, each floor above ground level has a rescue capacity per wing that should not be exceeded. The rescue capacity needed for each bid depends on the surface area of the space and the function that the bidder has in mind. Spaces on several floors are linked through utility shafts, providing ventilation, electricity, heating, etc. Typically, each space has to get its entire utilities through a single shaft, however, if a space contains multiple shafts, the

choice is open which one to use. This choice can be made by the bidder, but can also be left open for the auctioneer in case the bidder has no preference. Each bid has some requirement for utilities (again depending on the surface area and the intended function), and bids can only be allocated insofar the capacity of the shaft that they use is not exceeded. A lack of bidders, or unpopular lots, may cause some parts of the building to be unoccupied after the auction. In that case, the unoccupied lots should also form one or more spaces which can still be sold in the future. Similarly, in order to make sure that a vacant space can be occupied later, sufficient rescue capacity should be reserved on its wing and floor, and there should be at least one shaft available with enough spare capacity to provide basic utilities.

In our setting, we assume that the seller is a private company and that its goal is to find a feasible allocation, maximizing the total revenue. However, there are a number of other considerations. We want the auction to be perceived as fair by the participants. Therefore, it is important that the auction is as transparent as possible. Second, the auction should be accessible to everyone. Indeed, a key concern in the design of any auction is participation: whether the design is able to attract enough bidders. Learning the auction rules, certainly for a non-professional bidder, can pose quite a challenge. Therefore, the auction rules should be as simple as possible, and the bidding process should be user-friendly. Finally, the auction must be completed within a single day. Although buying real-estate is an important decision for each bidder, bidders generally prefer a shorter auction, which reduces participation costs. Indeed, we assume that once a bidder has made up his mind to participate, he does not want the process to drag along for too long, and wants to know the results on a relative short term.

3 Auction design

Clearly, every auction needs a set of rules that determine the course of the auction, the actions the bidders can take, and the feedback they will get. First, we evaluate a number of designs suggested in the literature, and explain the general principles we opted for in our real-estate auction (see Section 3.1). Next, the design is fine-tuned based on a series of experiments with human bidders in which we studied the effect of various auction rules (see Section 3.2).

3.1 Auction formats

Several well-known auction mechanisms do not offer an answer to all the concerns mentioned in the previous section. The *simultaneous multi-round auction* (Milgrom 2000) was used extensively by the FCC for their spectrum auctions. It consists of multiple rounds, where bidders can raise the highest bid on any item by a bid increment specified for each lot. After each round, provisional winners and prices are announced, and bidders are also subject to an activity rule which limits the number of items a bidder can bid on in a round as a function of the bidder's behavior in previous rounds. Since bids on sets of items are not allowed, it is possible that the bidder wins some items that he needs, but not all (i.e. the exposure problem). Later, the simultaneous multi-round auction was enhanced to allow package bids. Nevertheless, as far as we are aware, this upgraded design was only used in experimental labs, and replaced by the clock auction, which is at least as efficient (Brunner et al. (2010)).

In the *clock auction* (Porter et al. (2003)), the auctioneer announces prices for each item individually, and bidders respond with demands at the specified prices. Variants of the clock auction involve bids on packages, rather than quantities for individual items. In subsequent rounds, prices are raised incrementally as long as there is excess demand for any item. The prices of the packages are

assumed to be the sum of the prices of the constituting items. In a setting like ours, where most lots have no value individually, this procedure would not be successful. Moreover, when bidders pull out, a large drop in demand may lead to unsold lots, despite the fact that there are bidders willing to pay the reserve price. Furthermore, there is no way to ensure that the unsold lots form one or more solid spaces that are suitable for selling on a secondary market.

The *PAUSE* procedure (Kelly & Steinberg 2000) also treats the items separately in its first phase (a simultaneous ascending auction). Moreover, in its second phase, where bids on sets of items are allowed, it puts the computational burden with the bidder: each bidder is required to submit his bid as a part of a global bid (including bids of other bidders) that covers all items in the auction. Given the inexperienced bidders in our setting, this is not appropriate for our real-estate auction.

The *resource allocation design* (Kwasnica et al. (2005)) is iterative, forces a minimum bid increment, and allows package bids. It uses an activity rule that is quite similar to the one used in the simultaneous multi-round auction. The key contribution in this auction design is the pricing rule. It computes linear and anonymous prices, such that they exactly match the bid price for all winning bids, and (approximately) indicate that losing bids should be raised in order to win. Despite a good performance with respect to efficiency in lab experiments, the resource allocation design is not well suited for our setting, since it is based on an additive bidding language, that allows multiple bids to be accepted from one bidder.

Quan (1994) reports the *winner's choice auction* for auctioning multiple real-estate properties. This auction is a classic open outcry auction, where the winner has the right to pick any number of properties he wishes to buy at the unit price of his bid. If any properties remain, a second auction is held fol-

lowing the same rules, etc. Although appealing for its simplicity, this auction mechanism neglects the numerous allocation constraints described in Section 2. In fact, none of the above described auction designs seems able to handle these allocations constraints.

We opt for a sealed-bid first-price auction with binding bids. This is easy to conduct and explain. In a sealed-bid auction, the bidder communicates his bids directly and solely to the auctioneer. Sealed-bid auctions encourage participation: since a bidder with the highest valuation is uncertain about the best competitive bid, he may bid too greedily and lose. Thus, bidders with a lower valuation still have chance to win, and consequently, an incentive to participate (Pekeç & Rothkopf 2003). One disadvantage of a (single-round) sealed-bid auction is that it does not encourage price discovery. However, for our real-estate auction, this is less important, since we assume that bidders have a pretty good idea of what the lots are worth, by comparing with prices of nearby real estate. In a first-price system, a bidder pays exactly the amount he bids for the space he wins. This is not an obvious choice, since e.g. second-price auctions (i.e. the highest bidder pays the price of the second highest bid) makes it a dominant strategy for each bidder to bid his true valuation (Vickrey 1961). Nevertheless, the transparency of the first-price concept, taking into account the inexperience of the bidders in our real-estate auction, and the computational complexity of the Vickrey-Clarke-Groves (VCG) auction (i.e. the combinatorial variant of the second price auction) lead us to adopt a first-price rule. Binding bids mean that each bid should remain valid until the end of the auction: withdrawing or lowering bids is not allowed. Again, the simplicity of this rule is a reason for adopting it.

The computational complexity of the winner determination problem directs us to an iterative auction, with well-defined rounds and some opportunity for the bidders to reflect between rounds. Indeed, only for a small number of bidders

and bids per bidder, computation times allow a continuous mechanism, where bidders can bid at any time, and the provisional allocation is updated (quasi) in real-time (see Section 4.3).

Ideally, an auction ends when each winning bidder is happy with what he wins, and each losing bidder realizes that he is unwilling or unable to pay the amount needed to obtain what he wants. In a first-price auction, it suffices that no bidder bids more than he is willing to pay to realize the first part; the second part, however, cannot be guaranteed if the number of rounds is fixed before the start of the auction. Nevertheless, announcing the number of rounds beforehand makes sense for practical reasons, and reduces participation costs. A fixed deadline requires strong activity rules to prevent that serious bids are delayed until the final round (see Section 3.2).

An interesting issue is how to deal with a target group that should be favored (e.g. people with a low income looking for social housing). One idea is to give those bidders virtual money (i.e. a subsidy) in order to be able to compete with the other bidders. Another idea is simply to reserve a percentage of the lots exclusively for bidders in the target group to bid on (i.e. a set-aside). A recent study on the pros and cons of these two possibilities is Pai & Vohra (2012). We opted for a constraint in the winner determination problem stating that the outcome should allocate at least a given percentage of the total floor space to bidders from the target group (see Section 4.1). This approach provides a stricter guarantee than a subsidy may give, and does not force the seller to decide beforehand which parts of the building should be allocated to the target group.

Since there are many ways of partitioning a building into spaces, a bidder for some space benefits from bids on complementary spaces. This is inherent to any combinatorial auction, but may encourage bidders to collude (e.g. “don’t

bid on my space, and I won't bid on yours"), which can go at the expense of the total revenue. In our auction design, we provide the following measures to limit collusion: (i) sealed bids, which - at least in the last round - make violations of the collusive agreement undetectable until it's too late to react, (ii) binding bids, with guaranteed bidder solvency, and (iii) reserve prices, which reduce the maximum gain of collusive bidding.

Another issue is communication complexity (see e.g. Nisan & Segal (2006)). In general, in a combinatorial auction with n items, as many as $2^n - 1$ bids may have to be communicated to the auctioneer by each bidder. Simply keeping track of all these bids during the auction poses a considerable burden for the bidder. In our auction, the space constraints described in Section 2 drastically reduce this maximal number of bids per bidder. Bidders usually know very well what characteristics a space should have, and how they are going to use it. Thus, we believe a bidder will rarely have more than a dozen different spaces in his mindset, which makes the communication process manageable.

3.2 Lab experiments

We ran a laboratory experiment to gain insight into the effects of various auction design factors: the number of rounds, activity rules, anchor rules, and between-round feedback. We wanted to find out how particular choices for these settings influenced the behavior of real-life and inexperienced bidders, and how, consequently, the resulting auction performed on total revenue. The experiments took place at the CREED laboratory of the University of Amsterdam. At the start of the experiment, all subjects received a starting capital of 7 euros, which could be raised with additional payoffs (8 euros on average) depending on their performance in the auctions.

In all auctions in our experiment, 20 lots were offered for sale, corresponding to

two identical floors of a building. The lots were sold in various combinatorial auction designs, each with a fixed number of rounds. After each round, a winning allocation was determined according to the formulation in Section 4.1; no allocation constraints were taken into account other than that a bidder can win at most one space. Notice that the other constraints are not very meaningful in a setting with a limited number of lots and bidders, and hence would only cloud the view on the effect of the various auction design factors, which we will specify below.

In each auction, 8 to 10 subjects competed. We did not rematch between auctions so that each subject competed with the same group of subjects in all three auctions. Each subject was assigned one of 4 roles (“big business”, “small businesses”, “residential couple”, “residential single”), each with their own target surface. Note that we ensured that demand always exceeded supply. A subject made a bid by expressing a price for a set of adjacent lots; this price had to be at least 6 euros/m². In all auctions, each bidder could have bids on at most 8 different spaces at any point, and bids could never be revoked or lowered. Furthermore, bids had to satisfy the activity rule and the anchor rule (see below).

Together with his target, each subject received a value equal to price/m² x target surface. The values were private information for each subject, i.e., before the auction, each subject learned his own value, but the only information he had about the value of the other subjects was that all prices/m² were drawn independently from the uniform distribution on the interval [6,12]. A value could be raised with a bonus of (40) 20 euros per lot, if the space in the bid consisted of a number of (highly) “interesting” lots; the bonus per lot was the same for each subject. This ensured that preferences of the subjects were not uniform over the building, and simulated that some lots are more popular than others (which is likely to hold true in practice). When a subject won his target surface (or more) in the final round, his payoff equalled $3.5 \cdot (\text{value} - \text{bid}) / (\text{target}$

surface). At the end of the session, we paid the subject his starting capital and the sum of the payoffs collected in the auctions in which he participated. After each auction, we drew new values and assigned new roles to the subjects. For the sake of comparability between auctions, we kept the draws of the values constant across treatments.

We ran 9 sessions in which we varied (1) the number of auction rounds, (2) the activity rule, (3) the anchor rule, and (4) the feedback mechanism. In each session, the subjects interacted in 3 subsequent auctions. Sessions had auctions with 1, 5, and 8 rounds. We varied the order of the number of rounds from one session to the next. In all rounds, subjects had a fixed amount of time to submit their bids. The first round of each auction took 10 minutes while all the following rounds took 3 minutes.

An activity rule restricts subjects in terms of the spaces they can bid on. This forces them to remain active in early rounds of the auction because otherwise they lose the right to bid in later rounds. In the first auction round, subjects are free to bid on any space they want. In later rounds, the eligibility to bid on certain spaces depends on the activity rule. The activity rules do not apply on 'new bids', i.e. bids on spaces on which the subject did not bid in previous rounds. We include two types of activity rule:

- A1: In all rounds, subjects are free to bid on any space the anchor rule allowed them to. They can freely choose the amount with which they increase their bid; not increasing a bid is a valid option as well.
- A2: In a certain round, a subject is active on a space if in the last round in which he is not a provisional winner, he raises his previous bid on this space with at least a minimum bid increment of 0.5 euro/m². A subject is only allowed to raise bids on spaces on which he is active. Note that the activity of a subject always refers to the last round in which he was

not a provisional winner. In other words, provisional winners in a certain round do not have to bid in order to keep their eligibility to bid in the next round. Also note that when a subject raises a bid with less than the minimum bid increment, his bid is valid and taken into account, but he is not allowed to raise this bid again in the next rounds.

Our design involves the following three anchor rules:

- B1: In any round, a subject can submit a bid on any space.
- B2: In the first half of the auction, a subject can submit a bid on any space. In the second half of the auction, a subject can no longer bid on new spaces, and hence, he can only raise previous bids on spaces, provided that this is allowed by the activity rule.
- B3: In the first round, the subject is free to bid on any space. In the following rounds, the subject can only raise these bids (provided that this is allowed by the activity rule), and he is not allowed to introduce bids on new spaces.

Feedback is reported after each round. Our design includes two forms of feedback:

- F1: After each auction round, a subject only learns which of his bid(s) are currently in the winning allocation (if any).
- F2: After each auction round, a subject learns which spaces are allocated, and for what prices. Thus, each subject sees the winning allocation, however, information about individual bids is made anonymous.

Notice that in settings with only one auction round, the two activity rules coincide as well as the three anchor rules. Similarly, feedback is irrelevant in the case of single-round auctions. Still, a full $2 \times 2 \times 3 \times 2 + 1$ factorial design was not feasible from a practical point of view. Therefore, we used a design in which we varied the four treatment factors both between-subjects and within-subjects as

shown in Table 1. This table summarizes our experimental design. The second column presents the number of rounds in the three auctions of each experimental session. The second, third, and fourth columns include the relevant activity rules, the anchor rules, and the feedback rules respectively. If multiple rules are included, the first (second) refers to the first (second) multi-round auction. The final column shows the number of subjects that participated in the session (the distribution of subjects over the two auctions is between brackets).

Table 1: Summary of the experimental design

Session	# rounds	Activity rule	Anchor rule	Feedback	# subjects
1	1 - 5 - 8	A1 - A2	B1	F1	19 (10 + 9)
2	1 - 5 - 8	A1 - A2	B2	F2	18 (9 + 9)
3	1 - 5 - 8	A2 - A1	B3	F1	18 (9 + 9)
4	5 - 8 - 1	A2 - A2	B1	F2	16 (8 + 8)
5	5 - 8 - 1	A2 - A2	B2	F1	18 (9 + 9)
6	1 - 8 - 1	A2	B3	F2	18 (9 + 9)
7	8 - 1 - 5	A2 - A1	B1	F1	16 (8 + 8)
8	8 - 1 - 5	A2 - A1	B2	F2	16 (8 + 8)
9	8 - 1 - 5	A1 - A2	B3	F1 - F2	17 (9 + 8)

We evaluated the auctions' performance on the basis of their revenue as the fraction of the maximum total value subjects could obtain, i.e., the value that would be generated in the optimal feasible allocation of lots over subjects. To disentangle the effect of the four auction design factors on performance, we ran a random effects panel regression in which dummies for the factors, as well as the number of subjects, acted as explanatory variables. Table 2 presents the regressions results of a random effects model on auction revenue as a fraction of maximum surplus. The regression analysis takes bidding group specific effects into account where a bidding group is the set of 8, 9, or 10 subjects that competed against each other in three consecutive auctions. The first column indicates the independent variables and the goodness-of-fit measures. The first five variables are dummies. Five rounds, activity rule A2, anchor rule B2, and feedback rule F2 are used as benchmarks. The second column contains estimates

with (robust) standard errors between brackets, and ** (*) denoting significance at the 1% (5%) level.

Table 2: Random effects model on revenue as a fraction of maximum surplus

Variable	Estimate
1 round	-0.1781 (0.0557)**
8 rounds	-0.0598 (0.0447)
Activity rule A1	-0.0657 (0.0460)
Anchor rule B1	-0.1182 (0.0595)*
Anchor rule B3	-0.0785 (0.0530)
Feedback F1	0.0132 (0.0487)
Number of bidders	0.0137 (0.0273)
Constant	0.8401 (0.2382)
R^2 within	0.27
R^2 between	0.17
R^2 overall	0.23

The regression results suggest that if we restrict our attention to the design features tested in the lab, the auction should have 5 rounds, activity rule A2, anchor rule B2, feedback rule F1, and as many bidders as possible. In the case of 8 bidders, the model estimates that this auction’s revenue equals about 96% of the maximum bidder surplus. Note that according to this regression, the number of rounds and the anchor rule have a significant impact on the auction’s performance. If the auction takes only 1 round instead of 5, the regression results indicate a revenue loss of about 18 per cent points. If anchor rule B1 is used instead of B3, the predicted loss in revenue is about 12 per cent points. Both effects are not only statistically significant but also economically significant. The results with respect to feedback seem to contradict findings by Adomavicius et al. (2012), claiming that (for a continuous auction) without price feedback, bidders are unable to formulate effective bids, resulting in 30 to 40% of dead bids (i.e. bids that have no chance of winning). On the other hand, the effects of the feedback rule (and of the activity rule) in our lab experiments are not very pronounced, both statistically and economically.

Of course, we should be cautious interpreting our results. The number of (independent) observations is not particularly high. Moreover, “real” bidders might behave differently than the subjects in the laboratory, who were mainly students from the university of Amsterdam and perhaps not the most typical bidders in an actual real-estate auction. Still, we believe that we can safely advise against running a single-round auction and against using anchor rule B1. Note that anchor rule B1 allows a bidder to “stay below the radar” until the final round of the auction so that this rule practically amounts to an auction where all the action takes place in its final round. It appears that bidders tend to hold back if they have to make their final bids without (much) information about other bidders’ preferences.

4 Winner determination problem

Clearly, the design of the auction must be computationally feasible. In particular, the problem of deciding which bidders should get what items in order to maximize the auctioneer’s revenue, i.e. the winner determination problem, must be tractable. Although fast exact algorithms for this problem exist (e.g. CABOB, see Sandholm et al. (2005)), they are unable to cope with various allocation constraints, as needed in our setting. In Section 4.1, we develop a mathematical formulation for our winner determination problem. The computational complexity is discussed in Section 4.2, and computational experiments are reported in Section 4.3.

4.1 Mathematical formulation

The integer programming formulation that we develop for our winner determination problem is based on a set partitioning formulation, and uses the following

notation. The building, with a total surface area of A square meters, is divided into a set of lots denoted by L , and has a number of utility shafts $s \in S$. Each shaft s has a capacity C_{su} related to utility $u \in U$. Per floor $f \in F$, there are a number of wings $w \in W$, which have a rescue capacity for $O_{f,w}$ persons. We use B to denote the set of all bidders; subsets $B_c \subseteq B$ should be awarded at least (or at most) a fraction f_c (or F_c) of the building's surface area. Similarly, we use r_c (R_c) to denote the minimum (maximum) number of spaces that should be allocated to bidders of type c . The set $B_a \subseteq B$ includes those bidders who indicated that they want to win either all their bids, or none at all. Each bid $t \in T$ belongs to one bidder $b(t)$, and is characterized by the following parameters. We use $L(t)$ to represent the set of lots included in bid t , a_t for the surface area of this space, and p_t for the price that the bidder is willing to pay. The space is situated on wing $w(t)$ and floor $f(t)$, and the utility shaft that is to be used is given by $s(t)$. If the bidder does not explicitly mention which shaft he wishes to use, we duplicate his bid for each shaft that is contained in his space: the set $D(t)$ contains bid t and its duplicates. We use o_t to denote the number of persons for which rescue capacity is needed for bid t , and c_{tu} for the required capacity of utility u . We also add a dummy bidder $d \notin B$, that bids on all (valid) spaces. The dummy bidder always bids the reserve price (i.e. the price that reflects the seller's preference for leaving spaces unallocated), and requires some minimal utility capacities. We use the decision variable x_t which is 1 if bid t is allocated, and 0 otherwise, and y_b for each bidder b in B_a which is 1 if all bidder b 's bids are allocated, and 0 if none is allocated.

maximize

$$\sum_{t \in T} p_t x_t \tag{1}$$

subject to

$$\sum_{t \in T: l \in L(t)} x_t = 1 \quad \forall l \in L \quad (2)$$

$$\sum_{t \in T: b(t)=b} x_t \leq 1 \quad \forall b \in B \setminus B_a \quad (3)$$

$$f_c A \leq \sum_{t \in T: b(t) \in B_c} a_t x_t \leq F_c A \quad \forall c \quad (4)$$

$$r_c \leq \sum_{t \in T: b(t) \in B_c} x_t \leq R_c \quad \forall c \quad (5)$$

$$\sum_{t \in T: w(t)=w, f(t)=f} o_t x_t \leq O_{f,w} \quad \forall w \in W, \forall f \in F \quad (6)$$

$$\sum_{t \in T: s(t)=s} c_{tu} x_t \leq C_{su} \quad \forall s \in S, \forall u \in U \quad (7)$$

$$\sum_{t \in D(t)} x_t = y_b \quad \forall b \in B_a, \forall t \in T : b(t) = b \quad (8)$$

$$y_b \in \{0, 1\} \quad \forall b \in B_a \quad (9)$$

$$x_t \in \{0, 1\} \quad \forall t \in T \quad (10)$$

The objective function (1) states that the total rent should be maximized. The first set of constraints (2) enforces that each lot needs to be allocated. Indeed, this is necessary to ensure that unoccupied lots (i.e. lots allocated to a dummy bidder) form a valid space. The second set of constraints (3) ensures that each bidder wins at most one space (except for those bidders that want to win all their bids or none at all, and the dummy bidder). Constraints (4) make sure that each type of bidders acquires at least (at most) a given percentage of the total surface area, the next constraint quite similarly puts a limit on the number of spaces that is allocated to each type of bidders. Constraints (6) guarantee that the rescue requirement is satisfied for each wing, and constraints (7) enforce that the utility capacities are respected for each shaft. The fact that some bidders want all their bids allocated, or none at all, is settled with constraints

(8) and (9). The final set of constraints makes sure that bids are fully accepted or not at all.

Since we have dummy bids on all valid spaces, constraints (2) can always be satisfied by allocating the entire building to the dummy bidder. However, due to constraints (4) and (5), there may not be a feasible solution. If this happens in the first round, the round can be recomputed with f^c and/or r_c lowered, if necessary even to 0. In the latter case, the model will produce a feasible solution, for which it should be announced to the bidders that in the current setting, the conditions for a valid auction result are not satisfied. Bidders can use the next round to bid on other spaces, such that after that round, constraints (4) and (5) are satisfied. Otherwise, the model returns infeasibility, and the auction is canceled.

4.2 Theoretical results

In general, the winner determination problem is NP-hard (van Hoesel and Müller (2001)), and does not allow good approximation results (Sandholm (2002)). However, our winner determination problem is based on a restricted topology. Indeed, in our auction, the lots in each wing and on each floor are *arranged in rows* (see e.g. Figure 1). Notice that the lots need not have the same size, or be aligned over all rows. Moreover, bidders can only bid on sets of adjacent lots. Thus, the resulting problem is the following: given lots arranged on rows, find an allocation that maximizes the total rent, provided that each bidder can win at most one bid, and each lot can be allocated at most once. We refer to this problem as the real-estate winner determination problem where each bidder is allowed to bid up to b bids (REWDP-b). This problem corresponds to the formulation in Section 4.1, when ignoring constraints (4)-(9), and relaxing '=' to ' \leq ' in constraints (2). Problem REWDP-b is related to a problem known in literature as the job interval selection problem (JISP1): given n pairs of

intervals on the real line, select as many intervals as possible such that no two selected intervals intersect, and at most one interval is selected from each pair. The fact that JISP1 is MAX SNP-hard (Spieksma 1999) immediately leads to the following result:

Theorem 1 *REWDP- b for $b \geq 2$ is NP-hard, even when all lots are arranged on a single row.*

An interesting case arises when $b = 1$, the real-estate winner determination problem with a single bid per bidder (REWDP-1). Notice that this is equivalent to a setting where bidders have multiple bids, and are allowed to win more than one bid. It is well-known that if all lots are arranged on a single row, this problem is polynomially solvable. Indeed, Rothkopf et al. (1998) found that the winner determination problem can be solved in polynomial time if a linear order exists among the items, and bidders can only bid on subsets of consecutive items, even when the first item in the ordering is considered the successor of the last (i.e. a circular order). It is not difficult to prove that REWDP-1 is NP-hard if the lots are arranged on n rows. Orlin (2011) showed us that REWDP-1 remains easy if the lots are arranged on two rows; we describe an algorithm for REWDP-1 in Theorem 2.

Theorem 2 *REWDP-1 with all lots arranged on two rows can be solved in polynomial time.*

Proof.

In REWDP-1 with n bidders, each bidder k expresses one bid b_k on a set of contiguous lots. All $m = p + q$ lots are arranged on two rows: we use l_1, l_2, \dots, l_p to denote the p lots in the lower row; the lots in the upper rows are denoted with u_1, u_2, \dots, u_q . A block of consecutive lots in the lower row including l_i, l_{i+1}, \dots, l_j is denoted as $[l_i, l_j]$. Similarly, a block of consecutive lots in the upper row including u_i, u_{i+1}, \dots, u_j is denoted as $[u_i, u_j]$. If $i > j$, then $[l_i, l_j] = [u_i, u_j] = \emptyset$. A space consists of a contiguous set of lots, however, if the lots in the upper

row are not consecutive, or the lots in the lower row are not consecutive, we say that the space *has a gap*. Figure 2 shows a space with a gap in the upper row. Notice that a space can have multiple gaps.




Figure 2: A bid with a gap in a 2-row problem

In a preprocessing step, we replace each bid on a space B with one or more gaps, by a bid on this space that takes into account bids on spaces within the gaps (say B'). We do this by filling each gap optimally with one or more other bids. Thus, the bid on B' is a combined bid consisting of the bid on B and one or more single row bids. Notice that optimally filling a gap corresponds to solving REWDP-1 on a single row, which can be done in polynomial time ($O(m^2)$, see Rothkopf et al. (1998)). Since a space can have at most $O(m)$ gaps, filling all gaps can be done in $O(nm^3)$.

We now show how we can solve REWDP-1 with bids on gap-free spaces and lots arranged on two rows with dynamic programming. In order to simplify the presentation, we write the dynamic program as a longest path problem on a graph $G = (V, A)$. The node set is $V = \{s\} \cup \{\langle i, j \rangle : i \in \{0, 1, 2, \dots, p\}, j \in \{0, 1, 2, \dots, q\}\}$, where node s is the source node, and $\langle i, j \rangle$ designates the collection of lots $[l_1, l_i] \cup [u_1, u_j]$. The arc set A includes the following arcs:

- an arc from s to $\langle 0, 0 \rangle$ with length 0,
- for every bid b_k on a gap-free space $[l_i, l_{i'}] \cup [u_j, u_{j'}]$, there is an arc from $\langle i - 1, j - 1 \rangle$ to $\langle i', j' \rangle$ with a length b_k ,
- for every bid b_k on a space $[l_i, l_{i'}]$ and each $j \in \{0, 1, 2, \dots, q\}$, there is an arc from $\langle i - 1, j \rangle$ to $\langle i', j \rangle$ with a length b_k ,

- for every bid b_k on a space $[u_j, u_{j'}]$ and each $i \in \{0, 1, 2, \dots, p\}$, there is an arc from $\langle i, j - 1 \rangle$ to $\langle i, j' \rangle$ with a length b_k ,
- an arc from $\langle i, j \rangle$ to $\langle i + 1, j \rangle$ with length 0, for $i \in \{0, 1, 2, \dots, p - 1\}$ and $j \in \{0, 1, 2, \dots, q\}$, provided that no arc from $\langle i, j \rangle$ to $\langle i + 1, j \rangle$ has been specified above,
- an arc from $\langle i, j \rangle$ to $\langle i, j + 1 \rangle$ with length 0, for $i \in \{0, 1, 2, \dots, p\}$ and $j \in \{0, 1, 2, \dots, q - 1\}$, provided that no arc from $\langle i, j \rangle$ to $\langle i, j + 1 \rangle$ has been specified above.

The graph G has $O(m^2)$ nodes and $O(m^2 + nm)$ arcs. By construction, G is such that the length of the longest path from node s to node $\langle i, j \rangle$ is the maximum profit that one can obtain by assigning the lots in $[l_1, l_i] \cup [u_1, u_j]$ to bidders. Therefore, an optimal solution for REWDP-1 will be given by a maximum length path from node s to node $\langle p, q \rangle$. Notice that a longest path can efficiently be found since G is acyclic. ■

4.3 Computational experiments

We implemented the formulation described in Section 4.1, and solved it using IBM ILOG Cplex, version 12.3. In order to evaluate the performance of our algorithm, we carried out a number of computational experiments on randomly generated instances. These instances involve a building with 125 lots, spread over 7 floors, and with which 1000 valid spaces can be formed. We included 5 allocation constraints (heating, electricity, rescue capacity, gas, ventilation), and 3 groups of bidders, with constraints on the surface to be allocated in total to each of these groups. We generated instances with 100, 500, 1000, and 2500 bidders, where each bidder expressed 5, 10, 20, and 40 bids. All instances are available online (Goossens 2012).

The main goal of these experiments is to evaluate whether the computation times are reasonable. Therefore, only one single round was considered. All instances were solved on a Windows XP based system, with 2 Intel Core 2.8GHz processors. The computation times are summarized in Table 3; for each line, 10 instances were solved. It should be noted that these computation times only account for solving the model; processing all bids and constructing the model may also take several minutes.

Table 3: Average and maximum (between brackets) computation times, in seconds

Bidders	Bids per bidder			
	5	10	20	40
100	0.5 (0.9)	0.7 (1.1)	1.2 (2.0)	1.2 (2.3)
500	1.5 (2.4)	2.5 (4.8)	5.7 (11.8)	20.2 (55.2)
1000	2.2 (3.8)	5.2 (11.5)	39.9 (89.1)	37.5 (118.8)
2500	13.0 (26.0)	23.2 (55.9)	157.6 (362.2)	491.4 (836.1)

Clearly, these instances can be solved very efficiently, although there are considerable differences in computation time between instances with the same number of bidders and bids per bidder. The computation times do not generally allow a continuous auction, but we may expect our model to compute an optimal allocation within 15 minutes for realistic problem sizes, which makes a discrete iterative auction tractable.

5 Practical application: Solids

To raise buildings that will last for at least 200 years, with its tenants deciding how to use the building. The Dutch housing association Stadgenoot aims to make this happen in Amsterdam with so-called “Solids”. A Solid is a sustainable building without predefined purpose; in fact, a Solid should accommodate any (legal) functionality. The Solids concept is inspired by Amsterdam’s squatting tradition, particularly the freedom and flexibility with which squatters make

use of the space they occupy (Bijdendijk 2006). The main idea is that it is up to the tenants to decide on the use, the size, the configuration and even the rent of their space in the Solid. Stadgenoot sees Solids as highly suitable for a large variety of tenants: for (large) families, with everyone getting their own area, all linked up to a shared family room and kitchen, for entrepreneurs who are looking for a living area with a work space, for students, restaurants, etc. Stadgenoot is, like other housing associations in The Netherlands, a non-profit private organization with strong links with local authorities. The first housing associations were founded in the 19th century to do something about the housing shortage and deplorable housing conditions for working class families. Given this background, it is Stadgenoot's position that Solids should be open for everyone, including people with a small budget.

Space in a Solid is delivered as a shell. This means that within the building, there are walls between the solid spaces, and each solid space has access to a shaft with ventilation provision, drainage, electricity, etc. However, within a solid space, it is the tenant who decides where to place partition walls, interior doors, etc. This enables the rented space to be designed for a whole range of purposes: living, working, culture, or a any combination of these. Stadgenoot remains the owner of the shell; the tenants rent the solid space, and own the interior. If a tenant leaves, he can sell the interior to a next tenant. Over time, solid spaces can grow (when they are merged with another solid space) or shrink (when they are split up), and be used in very different ways, depending on the needs of the time. We refer to Stadgenoot (2011) for more information about the various Solids.

5.1 The Solid auction

The first Solid to be auctioned is Solid 11. Solid 11 accounts for over 7000 square meters of floor space, and features a spacious roof terrace with a splendid view

over the city. The solid is divided into 125 lots (i.e. items), distributed over 7 floors; Figure 1 shows how the second floor of Solid 11 is divided into 22 lots. Since the allotment is heavily constrained by building regulations and technical issues, we were not primarily involved in drawing the allotment, and from the point of view of the auction we can consider it given.

Any interested tenant can specify solid spaces, together with the price he is willing to pay as a monthly rent (paying rent by month is the usual setting in the Netherlands). Stadgenoot distinguishes three types of bidders: residential, commercial, and social bidders (i.e. people with a low income). For social bidders, Dutch law imposes an upper bound on the monthly rent of 650 euro. There is a reserve price of 6 euro per square meter. Our assumption that there are no externalities in the valuations is in line with the philosophy of the Solid project. Indeed, if key concepts like flexibility and freedom appeal to a bidder, he should accept that this applies to his neighbor as well. Nevertheless, activities that could severely disturb other tenants (e.g. a night club or heavy industry) are not allowed.

There are several constraints on the composition of the spaces. A space must consist of either all lots of an entire floor, or adjacent lots on a single floor. Figure 1 shows that floors above ground level consist of two wings, separated by an open space in the middle, such that lots in different wings are not adjacent. On the ground level, however, the wings are connected by additional lots, and consequently, valid solid spaces spanning both wings are possible. Furthermore, a valid subset of lots needs to have enough doors to the central hallway (denoted by black triangles in Figure 1), and have access to at least one appropriate utility shafts. Indeed, Dutch building regulations prescribe that separate utility shafts are provided for commercial and non-commercial purposes. There are also bounds to the surface area of the solid space, depending on the type of bidder. These constraints allow only 1214 valid spaces to be formed with the

125 lots.

Apart from the allocation constraints discussed in Section 2, Stadgenoot's historical background and philosophy lead to a number of additional restrictions. The allocation should reserve at least 15% of the surface to social bidders. Residential bidders should get at least 25% of the surface, whereas commercial bidders should be allocated at least 30%. If there are not enough social bidders, there should be enough unoccupied and suitable solid spaces, such that when they are rented to social bidders later, the required 15% of the total surface of the solid can still be met. Furthermore, at most 3 restaurants are allowed in the Solid. Stadgenoot's goal is to find a feasible allocation, maximizing the total rent.

Eventually, Stadgenoot implemented the auction design that is optimal according to our lab experiments, except that they provide the bidder feedback on which spaces were allocated, and for what prices, after each round (i.e. feedback rule F2). A user interface has been developed by an IT company for the bidders to express their preferences. This system shows them the allotment of the solid, including the positioning of the shafts, doors, and windows, and allows them to select lots to form a solid space. If the selected lots do not form a valid solid space, the system immediately provides feedback. Furthermore, a user can provide the user interface with a number of desired characteristics (e.g. surface area, orientation, etc.) and receive a list of solid spaces that satisfy them. This can be done weeks before the start of the auction; interested bidders also have the opportunity to visit the Solid beforehand to actually see the space(s) that they intend to bid on.

5.2 Outcome and evaluation

On May 7, 2011, 114 bidders showed up to take part in the combinatorial auction for Solid 11. The bidding started after the gong was sounded by the mayor of Amsterdam, and ended 8 hours and 725 bids later with 39 winning bidders and 95% of the Solid rented out (the remaining 5% were rented out in a more traditional fashion in the weeks after the auction). All bids are available on our website (Goossens 2012).

Table 4: Outcome statistics

	Round 1	Round 2	Round 3	Round 4	Round 5
Bids	528	725	752	725	725
Winners	38	39	43	41	39
New winners	38	18	16	6	4
Total rent (indexed)	100	114	137	141	167
Computation time (s)	0.27	2.28	0.23	0.91	1.50

The outcome of the auction is summarized in Table 4. It shows that in the first round, 528 bids were expressed, followed by another 197 in the next round (recall that our anchor rule prevents bids on new solid spaces after round 2). In each round, there were about 40 winners, and as many as 82 bidders were winners at some point in the auction. There were 4 bidders who managed to fly under the radar and strike only in the last round, but on the other hand, 9 bidders were winning bidders from the first till the last round. A look at the (indexed) total rent in the winning allocation that was computed after each round shows that the auction closed with a total rent that was 67% higher than after the first round. The most substantial increase occurred in the third and fifth round, where 50 bidders increased at least one bid. The winner determination problem was solved very efficiently, with all computation times below 3 seconds.

Our evaluation focusses on two aspects: participation and design. The number of bidders was borderline for a decent competition. One consequence for

the outcome was that similar solid spaces went for varying prices. There are several reasons that influenced participation. First of all, the auction was held in unfavorable economic circumstances. Indeed, the housing market in general, and the rental market in particular follows a downward trend in recent years (Francke 2010). The economic and financial crisis certainly did not help either. Second, there is the threshold that comes with a pioneer project with a unique philosophy as Solids. Indeed, the concept of renting space and having to build the entire interior design, and the idea of maximal freedom (also for ones neighbors), requires open-minded tenants with some sense of initiative and adventure. Finally, the auction itself inevitably also forms a threshold. Indeed, the solid auction is new, and certainly for a non-professional bidder, learning this mechanism poses quite a challenge. For instance, consider an elderly couple, both not “computer-literate”, who want to participate. Simply mastering the process of entering a bid is a hurdle to take.

With respect to the auction design, the substantial increase in round 3 suggests that our anchor rule provided an incentive for people to raise their bids. Moreover, the fact that almost 90% of the winning bidders had already made a provisionally winning bid in one or more of the previous rounds shows that the activity rule worked as well. Most bidders found the feedback they received very useful. Some bidders, though, did not see their bid being part of the winning allocation, despite the fact that it was higher than the winning bid on the same solid space. In those cases, the feedback did not always provide a proper explanation for the role played by constraints (2)-(10), but thanks to several training sessions, the bidders were perfectly aware of the existence of these constraints. Furthermore, few of these constraints were binding in the final allocation. Overall, the auction worked adequately, and Stadgenoot was pleased with the outcome. Consequently, our design was used again one month later for the auction of two other Solids.

6 Conclusions

We described a combinatorial auction for allocating floor space in a building to non-professional bidders. Based on experiments with human bidders, a multi-round auction format was chosen, where bidders are encouraged to start bidding early in the auction. The auction uses a fair and transparent allocation mechanism, that guarantees that the outcome is in accordance with building and municipal regulations, and the seller's policy.

The main managerial implications of this work are threefold. The primary contribution is that a combinatorial auction in a real-estate market, where various allocation constraints need to be respected, is practicable. Moreover, whereas combinatorial auctions were usually considered as too complicated for non-professional bidders, we showed that they can be made accessible to bidders whose auction experience does not exceed eBay. Indeed, our design was successfully implemented in practice, culminating in two auctions where over 13000 square meters (in total) were allocated at a satisfactory price in a very short period to many bidders. A key assumption, however, is that bids can only be expressed on adjacent lots, which reduces the number of combinations, and the complexity of the winner determination problem.

A secondary implication is the understanding that preparing the bidders is immensely important. It is highly recommendable to organize a mock auction before the real auction, to let the bidder familiarize with the bidding system interface, and perhaps even experience the consequences of certain bid strategies. In a combinatorial auction, it may happen that a lower bid is allocated at the expense of a higher bid. When allocation constraints are present, this phenomenon will happen more frequently. Bidders should be informed beforehand that this may happen, preferably with specific examples that add to their insight. Ignoring this will lead to frustration with some bidders, and possibly

result in legal cases.

Finally, when designing an auction, it is not enough to develop a set of simple and effective rules, and carefully explain them to the bidders. Indeed, we should focus on the so-called *cognitive complexity* of formulating combinatorial bids (Porter et al. (2003)). In this regard, it is the complexity of participating in the auction, and formulating sensible bids that matters. The importance of the auction software cannot be overestimated in that regard. It should for instance be able to give an overview of spaces that satisfy the needs of the bidder (surface area, orientation, shape, ...), well before, but also during the auction. Furthermore, it should alert the bidder of the consequences of his action (e.g. with respect to the activity rule). User-friendly auction software is the key to a positive experience for the bidders.

References

- Adomavicius, G., Curley, S. & Gupta, A. (2012). Effect of information feedback on bidder behavior in continuous combinatorial auctions, *Management Science*, published online before print, DOI:10.1287/mnsc.1110.1443 .
- Bayers, C. (2000). Capitalist econstruction, *Wired Magazine* **8**(3): 210–212.
- Bijdendijk, F. (2006). *Met andere ogen*, Het Oosten Woningcorporatie, Amsterdam [in Dutch].
- Brunner, C., Goeree, J., Holt, C. & Ledyard, J. (2010). An experimental test of flexible combinatorial spectrum auction formats, *American Economic Journal: Microeconomics* **2**(1): 39–57.
- Cramton, P., Steinberg, R. & Shoham, Y. (2005). *Combinatorial auctions*, MIT Press.

- Eklöf, M. & Lunander, A. (2003). Open outcry auctions with secret reservation prices: an empirical application to executive auctions of tenant owners apartments in sweden, *Journal of Econometrics* **114**: 243–260.
- Epstein, R., Henríquez, L., Catalán, J., Weintraub, G. Y. & Martínez, C. (2002). A combinational auction improves school meals in Chile, *Interfaces* **32**(6): 1–14.
- Francke, M. (2010). *Casametrie: de kunst van het modelleren en het voorspellen van de marktwaarde van woningen [in Dutch]*, Oratiereeks 353, Vossiuspers, UvA.
- Goeree, J. & Holt, C. (2010). Hierarchical package bidding: A paper & pencil combinatorial auction, *Games and Economic Behavior* **70**(1): 146–169.
- Goossens, D. (2012). Solids: Auction and instances. <http://www.econ.kuleuven.be/public/NDBAE53/solids.html>.
- Hohner, G., Rich, J., Ng, E., Reid, G., Davenport, A. J., Kalagnanam, J., Lee, H. S. & Chae, A. (2003). Combinatorial and quantity-discount procurement auctions benefit Mars, incorporated and its suppliers, *Interfaces* **33**(1): 23–35.
- Kelly, F. & Steinberg, R. (2000). A combinatorial auction with multiple winners for universal service, *Management Science* **46**(4): 586–596.
- Kwasnica, A., Ledyard, J., Porter, D. & DeMartini, C. (2005). A new and improved design for multiobject iterative auctions, *Management Science* **51**(3): 419–434.
- Ledyard, J. O., Olson, M., Porter, D., Swanson, J. A. & Torma, D. P. (2002). The first use of a combined-value auction for transportation services, *Interfaces* **32**(5): 4–12.
- McMillan, J. (1994). Selling spectrum rights, *Journal of Economic Perspectives* **8**: 145–162.

- Milgrom, P. (2000). Putting auction theory to work: The simultaneous ascending auction, *Journal of Political Economy* **108**(2): 245–272.
- Nisan, N. & Segal, I. (2006). The communication requirements of efficient allocations and supporting prices, *Journal of Economic Theory* **129**(1): 192–224.
- Orlin, J. (2011). Personal communication.
- Pai, M. M. & Vohra, R. (2012). Auction design with fairness concerns: subsidies vs. set-asides. Manuscript, available online at <http://faculty.chicagobooth.edu/appliedtheory/workshop/Spring2012/Vohra.pdf>.
- Pekeç, A. & Rothkopf, M. (2003). Combinatorial auction design, *Management Science* **49**(11): 1485–1503.
- Porter, D., Rassenti, S., Roopnarine, A. & Smith, V. (2003). Combinatorial auction design, *Proceedings of the National Academy of Sciences of the United States of America* **100**: 11153–11157.
- Quan, D. (1994). Real estate auctions: A survey of theory and practice, *Journal of Real Estate Finance and Economics* **9**: 23–49.
- Rothkopf, M., Pekeç, A. & Harstad, R. M. (1998). Computationally manageable combinatorial auctions, *Management Science* **44**(8): 1131–1147.
- Sandholm, T. (2002). Algorithm for optimal winner determination in combinatorial auctions, *Artificial Intelligence* **135**(1-2): 1–54.
- Sandholm, T., Suri, S., Gilpin, A. & Levine, D. (2005). CABOB: a fast optimal algorithm for combinatorial auctions, *Management Science* **51**: 374–390.
- Spieksma, F. C. R. (1999). On the approximability of an interval scheduling problem, *Journal of Scheduling* **2**: 215–227.
- Stadgenoot (2011). Solids: The floor is yours. [in Dutch]. <http://www.stadgenoot.nl/aanbod/nieuwbouw/solids>.

Van Hoesel, S. & Müller, R. (2001). Optimization in electronic markets: examples in combinatorial auctions, *Netnomics* **3**(1): 23–33.

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