A Market-Inspired Bidding Scheme for Peer Review Paper Assignment

Reshef Meir
Technion-Israel Institute of Technology
reshefm@ie.technion.ac.il

Jérôme Lang
CNRS, LAMSADE, Université Paris-Dauphine
lang@lamsade.dauphine.fr

Julien Lesca
Université Paris-Dauphine
julien.lesca@gmail.com

Natan Kaminsky
Technion-Israel Institute of Technology
natank@campus.technion.ac.il

Nicholas Mattei
Tulane University
nsmattei@tulane.edu

ABSTRACT

We propose a market-inspired bidding scheme for the assignment of paper reviews in large academic conferences. The primary contribution is an analysis of the incentives of reviewers during the bidding phases, when reviewers have private costs and some information about the demand for each paper, and want to obtain the best possible k papers for a predetermined k.

We show that by assigning budgets to reviewers and a 'price' for every paper that is (roughly) inversely proportional to its demand, the best response of a reviewer is to bid *sincerely*, i.e., on her most favorite papers, and match the budget even when it is not enforced.

Finally, we show via extensive simulations on bidding data from real conferences, that our suggested bidding scheme would substantially improve the assignment, under several common assignment algorithms.

1 INTRODUCTION

Academics spend much of their time and effort (that is, *our* time and effort) on peer-review for journals and conferences. This is an unpaid labor that academics perform out of sense of duty, which serves several important purposes for all involved parties. It helps editors and program chairs make informed decisions on what papers to publish; it provides authors with valuable feedback on their work; and it keeps the reviewer updated about recent advances in their respective fields.

While in journals the assignment of papers to reviewers is typically handled manually by the editors, peer-reviewed conferences and workshops often use automated assignment algorithms whose outcome is based on the stated preferences of the program committee members (reviewers). The program chairs intervene to solve problems in the assignment, such as allocating papers that no one asked to review. The frequency of these 'orphan' papers varies and that this is a recurrent problem in the design of bidding processes (see Table 1 below).

The central problem we consider in this paper is *how to improve the assignment in conferences, to the benefit of all involved parties: reviewers, program chairs, and authors.* Crucially, the improvements we suggest are easy to implement: they only require to reveal certain information to the reviewers about the current paper demands. These suggestions are orthogonal to the assignment algorithm and other design choices which are specific to the conference or the platform in use.

In the remainder of Section 1 we describe the current paper assignment process as it is typically performed in large computer science conferences and demonstrate some of the problematic issues using statistics from recent examples. We briefly argue, using ideas from mechanism design, that creating a market for bids is one way to alleviate these problems, we offer a summary of the primary contributions of our work, and explain how it complements previous work in the area.

1.1 The Paper Assignment Process

Consider a large CS conference such as AAAI, IJCAI, ICML, or NeurIPS which have somewhere between 1,000 – 10,000 submissions and 1,000 – 3,000 reviewers. The assignment process typically proceeds through the following steps:

- The program chair recruits program committee members (PCMs) to serve as reviewers. We denote them by the set N where $|N| = n$. PCMs do not know which papers they are going to review, but rather agree to review some number of papers, which ranges from 4-6 in some conferences to 10-14 in others.
- Authors submit their papers by a certain date. We denote the set of papers by M and $|M| = m$.
- PCMs get access to papers' titles and abstracts via an online platform such as EasyChair or Confmaster, and are asked to "bid" on papers they want to review. It is typically possible to bid one of several levels (e.g. "want to review", "can review if needed" etc.) as well as to report a conflict of interests (COI). There are usually 2-5 days to complete the bidding process.
- Typically PCMs are asked (but not enforced) to bid positively on a minimum number R of papers (e.g., 30), R being much higher than the actual number of expected reviews.
- After the bidding round is complete, the bid matrix is fed as input to an *assignment algorithm*, together with additional constraints such as the minimum and maximum number of

Appears at the 2nd Games, Agents, and Incentives Workshop (GAIW 2020). Held as part of the Workshops at the 19th International Conference on Autonomous Agents and Multiagent Systems., May 2020, Auckland, New Zealand

https://doi.org/

papers per PCM. Each paper should be assigned to a fixed number r of reviewers, typically 2—4.

- Once the program chair is satisfied with the assignment, they forward the assigned papers to PCMs and the reviewing period starts. PCMs who are dissatisfied with their assignment may email the program chair, who handles these issues manually.

On average, each PCM should get $k := \frac{mr}{n}$ papers for review. There are many assignment algorithms in use, which typically solve a constrained optimization problem that maximizes the number of papers that go to PCMs that have bid for them. Conferences use either an off-the-shelf algorithm provided with the platform or written specifically for the conference [8].[1] In other cases, the actual assignment algorithm itself can be a black-box, see discussion by Lian et al. [12].

In Table 1 below we provide some data for a few small, medium, and large conferences.[2] See more information on anonymized conferences in Table 2. Unfortunately, it is common that the bidding across papers is highly skewed, with some papers getting an overwhelming amount of positive bids, while others remain with very few or none at all (we can see that the EC conference, at least in 2018, is a refreshing outlier!)

The standard protocol has several drawbacks: (1) it is hard or impossible to find proper reviewers for 'orphan' (underdemanded) papers; (2) the assignment of overdemanded and underdemanded papers is somewhat arbitrary, and results in additional work for the PCMs and often low quality reviews; (3) PCMs spend time and effort bidding on papers that they are unlikely to get, either because they are overdemanded, or because they already unknowingly bid on enough papers; (4) program chairs typically spend a lot of time and effort coping with underdemanded papers.

1.2 A market for reviews

Imagine that instead of the current bidding system, we had a *market* where papers had prices and reviewers had money. Overdemanded papers would become very expensive, whereas the price of underdemanded papers would become low (or even negative). A reviewer currently holding some expensive (overdemanded) papers, might want to sell it and buy cheap (underdemanded) papers instead. When prices stabilize, papers with few bids will be rare since they will have very low (or negative) price.

Naively applying a market approach to the paper assignment problem may be impractical, as it requires too much interaction among already-busy PCMs. Moreover, the "paper market" differs from other markets in that all PCMs are eventually assigned a similar number of papers, even if they do not bid at all. The lack of quasi-linear utilities (as there is no use of real money) further complicates things.

Most importantly, the assignment may have various constraints and considerations that are difficult to capture formally and handle rigorously. Thus, program chairs are unlikely to adopt a completely new algorithm that may solve some issues but is not tailored for their needs, and PCMs may be deterred by an interface very different from what they are used to.

[1] One of the program co-chairs of AAAI-17, Shaul Markovitch, wrote his own algorithm to handle assignments.

[2] We thank the program chairs of these conferences that have agreed to contribute these statistics.

1.3 Paper Goal and Contribution

Our goal is to improve paper bidding, which in turn will improve paper assignment. We present a simple paper bidding mechanism called the *proportional bidding scheme*: PCMs are each assigned some initial budget R they are expected to exhaust, and the price of every paper changes dynamically, inversely proportional to its demand. Then, papers are assigned based on these bids using any ordinary assignment algorithm. That is, we do not aim to replace existing assignment algorithms, but to *enhance their input*—and thus to improve their output.

In order to compare the suggested bidding scheme to the current one, the main theoretical challenge is to model and analyze bidding incentives, without specifying an explicit assignment algorithm. More specifically, we do the following:

- Define a "mock assignment", as an *abstraction* that captures the key properties of assignment algorithms and aims at capturing the *perceived probability*, for a PCM, to be assigned a paper.
- Prove that in the game induced by the proportional bidding scheme and the mock assignment, there is an incentive for PCMs to make sincere bids that match the requested bidding amount R.
- Show via extensive simulations that if bidders are sensitive to price (less likely to bid on cheap, low probability papers) then our bidding scheme obtains socially better outcomes: the bids are more balanced across papers, and PCMs get better papers on average.

A secondary claim that we make is that the price mechanism can make the bidding process easier for some PCMs under the assumption that there is some effort involved in checking one's own utility for each paper.

We defer most of the longer proofs and some of the simulation results to the appendix.

1.4 Related Work

Assignment Algorithms. A first stream of papers consists of *information retrieval approaches* to the assignment problem: a paper is considered as a query, and each reviewer is considered as a text document. The aim is to retrieve a certain number of reviewers that are relevant to the paper without asking them to make explicit bids. These (numerous) papers are far from our work and focus; we only cite a recommendation-based approach that aims at easing paper bidding, namely the *Toronto Paper Matching System* (TPMS) [6]: for each paper j and each reviewer i, a fitness degree is determined by the proximity between j and some of the papers authored by i. The PCM can e.g. sort papers by their fitness degree during the bidding process to facilitate bidding and read fewer abstracts. Our proposal for a bidding scheme is compatible with TPMS or similar systems as well as advances in assignment and matching algorithms. Default bids may be computed taking into account keywords or TPMS score, if one of these is available, and reviewers would then be able to update their bid taking paper prices into account.

Conference	m	n	r	k	required R	$< r$ bids	0 bids	$\overline{\text{bid}}$ / PCM	$\overline{\text{bid}}$ / paper
IJCAI 18	3470	2035	3	~ 5	40-50	140	5	40.4	29.7
AAAI 17	2414	1321	3	~ 4	30	47	18	24.9	13.66
EC 18	280	193	3	$\sim$5	25	-	-	?	?
KSEM 06	235	88	2	10.7	30	18	5	16	6
KR 08	234	93	3	7.5	20	-	-	21.3	9.7
WINE 11	101	27	3	$\sim$11	?	10	-	?	?
TARK 17	91	20	3	13.5	25-30	32	8	15.85	3.4

Table 1: We recall that m is the number of papers, n the number of PC members, r the number of reviews per paper, k the average number of reviews per PCM, and R (when defined) the bidding target expressed by the program chair ("please bid on at least R papers"). We show the number of papers that received less than r bids, or 0 bids, respectively; we also show the average number of bids per PCM (respectively paper).

A second stream consists of *matching-based approaches*. The aim is to find an optimal assignment of papers to reviewers, given the input data. There are two subcategories of work in this area.

In the first subcategory, the input to the assignment is based on *keyword or topic overlap*: the input data consists of a set of keywords and/or topics for each paper, and for each reviewer. In this category we find Hartvigsen et al. [10], Nguyen et al. [16], Ahmed et al. [1], Long et al. [13], and Conry et al. [7]. The latter two integrate a recommendation stage and a matching stage. Long et al. [13] makes an interesting note echoes our observation from Table 1:

> *(...) it might happen that some popular reviewers are assigned to excessive papers while some other reviewers with very few or even no papers. (...)*

In the second subcategory, the assignment is based on *bids* that PCMs place on papers, as described in step 3 in the process presented in Section 1.1. In this category we find Lian et al. [12], Garg et al. [8], Goldsmith and Sloan [9] (who also consider keyword overlap), and the (unknown) algorithm used by EasyChair.

Bidding Behavior. Very few papers focus on the behavior of bidders within the assignment system. The only empirical work we know of is Rodriguez et al. [18], which analyses the bids of reviewers for a real conference, and studies its correlation with reviewer-paper fit. This fit is measured by a complex combination of techniques involving the co-author network, keyword occurrence, reviewer similarity and submission similarity. They find that the bidding behaviour is only weakly related to the subject of the submission, and that plenty of other (unidentified, conjectured) factors influence the bidding behaviour. They repeat the observation on imbalanced bids, and conclude that

> *Since bidding is the preliminary component of the manuscript-to-referee matching algorithm, sloppy bidding can have dramatic effects on which referees actually review which submissions.*

Game Theory and Mechanism Design. Some papers focus on strategyproof peer review [11, 19]. Especially, [19] consider scenarios where reviewers (who are also authors) may bid strategically so as to influence the assignment. We assume here reviewers bid independently from their interest as authors, but are still self-interested and would like to minimize their effort during both the bidding and reviewing phases. Finally, market approaches (such as using scrip money) have been considered for various allocation problems such as course allocation [3, 5, 17]). Yet we are unaware of a theoretical or practical application to review assignments.

All of the above points, together with our own observations in the past ten years or so (i.e., since paper bidding has been commonly used), lead to the conclusion that *the most effective way to improve the assignment is indeed to improve the input to the assignment algorithms, rather than fine tuning the algorithms themselves.* So without further ado, we get to work.

2 PRELIMINARIES

Throughout the paper, we denote the sets of PCMs (reviewers) and papers by N and M, respectively. Furthermore, we denote by n and m the sizes of these sets, respectively.

Assignments. Each paper j should be ultimately assigned to r_j reviewers, and each PCM i should get exactly k_i papers for review. Note that for an assignment to be possible, we must assume $\sum_{i \in N} k_i = \sum_{j \in M} r_j$.

A valid *fractional partial assignment* is a non-negative matrix $\boldsymbol{X} = (x_{ij})_{i \in N, j \in M}$ that meets the following constraints:

Capacity constraints $\sum_{j \in M} x_{ij} \leq k_i$ (every reviewer gets at most k_i papers);

Quota constraints $\forall i \in N, j \in M,\ x_{ij} \leq q_{ij}$. Typically $q_{ij} \in \{0, 1\}$, where $q_{ij} = 0$ meaning that there is a COI. However for analysis purposes we allow any nonnegative value.

Paper constraints $\forall j \in M,\ \sum_{i \in N} x_{ij} \leq r_j$.

Note that the above constraints are met trivially by the empty assignment. We say that $\boldsymbol{X}$ is *full* for PCM i if $\sum_{j \in M} x_{ij} = k_i$, and *complete* if it is full for all PCMs.

For ease of presentation we assume, unless mentioned otherwise, that $r_j = r$ for all j, that $r \geq 1$, and that $q_{ij} \leq 1$ for all $i \in N, j \in M$. An assignment is *integral* if $x_{ij} \in \mathbb{Z}$ for all i, j.

Bids. A (fractional) bidding profile is a real matrix $\boldsymbol{B} = (f_{ij})_{i \in N, j \in M}$, where $f_{ij} \in \mathbb{R}_+$ is the bid of PCM $i \in N$ on paper $j \in M$. We assume unless explicitly stated otherwise that $f_{ij} \in [0, 1]$, and we always require $f_{ij} \leq q_{ij}$ (in particular, a PCM cannot place a bid on a paper if she has a conflict of interest with it). The bid of PCM i is the vector $B_i := (f_{ij})_{j \in M}$. A bidding profile $\boldsymbol{B}$ is *integral* if $f_{ij} \in \{0, 1\}$ for all i, j. When bids are integral we sometimes abuse notation by writing B_i as the subset of papers that PCM i bids on.[3] We similarly denote by $D_j := (f_{ij})_{i \in N}$ the demand profile for paper j (induced by $\boldsymbol{B}$). In a given profile $\boldsymbol{B}$, we denote by $d_j := \sum_{i \in N} f_{ij}$ the total demand of paper j, and $\boldsymbol{d} := (d_j)_{j \in M}$.

An *assignment algorithm* is a function that takes as input a bidding profile $\boldsymbol{B}$, and outputs a valid partial assignment $\boldsymbol{X}$. We describe several assignment algorithms in use in Section 5.

Proportional Bidding Scheme. We describe steps of the bidding scheme we propose. Note that these steps are aligned with the process outlined in Section 1.1.

(1) For each paper j, the system displays a "price" $p_j := \min\{1, \frac{r}{d_j}\}$ (that is, inversely proportional to the demand of paper j).
(2) Each PCM i that enters the system submits a bid B_i.
(3) The *contribution* of PCM i in profile $\boldsymbol{B}$ is defined as the total price of demanded papers. Formally, $cp_i := \sum_{j \in M} f_{ij} p_j$. Crucially, the prices and thus the contribution will depend on the bidding profile, so they change dynamically as more PCMs enter or change their bids.
(4) The system may announce (or enforce) a bidding requirement R, such that a bid B_i is *sufficient* in profile $\boldsymbol{B}$ only if $cp_i \geq R$. A bid B_i that meets the requirement R exactly (i.e. where $cp_i = R$) is called *exact*.
(5) Like in current bidding systems, reviewers may log in later on and revise their bid. Unlike current systems, they may see different prices each time.
(6) Bidding continues until some stopping condition is met, e.g. all PCMs bid at least once, no bidder wants to revise her bid, or some deadline has passed.

The output of the bidding mechanism is some final bidding profile $\boldsymbol{B}^*$. When the bidding phase is over, we run some assignment algorithm A, with $\boldsymbol{B}^*$ and various allocation constraints as input to get an assignment.

Remark 1. *The price p_j that a PCM i sees is always the price* as if *i is already bidding on paper j, as it should reflect the contribution if a bid occurs. We demonstrate this in Example 2.1.*

Example 2.1. Suppose we have $m = 4$ papers (a, b, c, d) and $n = 6$ reviewers. $r = 2$ which entails each reviewer should get $k = \frac{mr}{n} = \frac{4}{3}$ papers. We only use integral bids in this example.

- Initially, bids are $B_1 = \{a, b\}, B_2 = \{a, c\}$, all other bids are empty. This yields $\boldsymbol{d} = (2, 1, 1, 0)$
- Reviewer 3 logs in. She sees the prices $p_a = \min\{1, \frac{r}{|D_a \cup \{3\}|}\} = \frac{2}{3}, p_b = \min\{1, \frac{r}{|D_b \cup \{3\}|}\} = \min\{1, \frac{2}{2}\} = 1, p_c = p_d = 1$ (the PCM is counted in the demand, cf. Remark 1). Suppose she bids on $\{a, b, d\}$.
- Now Reviewer 4 logs in and sees prices of $p_a = \frac{2}{4} = \frac{1}{2}, p_b = \frac{2}{3}, p_c = 1, p_d = 1$. Suppose he bids on $\{b, d\}$, so now demands are $(3, 3, 1, 2)$.
- Now, Reviewer 2 logs in again, and sees prices of $p_a = \frac{2}{3}, p_b = \frac{1}{2}, p_c = 1, p_d = \frac{2}{3}$.

[3]For simplicity we assume that there is only one level of positive integral bids. We later explain how to extend analysis and simulations to multiple bid levels.

3 PROPORTIONAL MOCK ASSIGNMENTS

In order to analyze the outcome and the bidding behavior, we must also take into account the assignment algorithm. However, since most practical assignment algorithms are based on integer programming, the connection between input (bids) and output (assignment) is quite complicated and sensitive to small changes in the input. Thus a PCM cannot readily use them to derive her beliefs about her assignment. Moreover, the bid-to-assignment (or the probability of assignment) connection may be both counter-intuitive, and requires information not available to the bidder, i.e., knowledge of the full bid matrix. Since we are interested in incentive analysis, it is more important to capture the way that PCMs *perceive* the effect of their actions, especially when the actual effect may vary considerably across assignment algorithms and may not even be known, for example, EasyChair does not reveal what algorithm they use [12].

To make a rigorous theoretical analysis possible, instead of dealing with particular assignment algorithms, we describe a (fractional) *mock assignment* that captures in an intuitive way the connection between the bid of a single PCM and her assigned papers. Intuitively, the PCM gets more of paper j when bidding positively on j, when the demand for j is lower, and when she bids on fewer other papers. We will further assume that the fractional assignment of paper j to PCM i changes *linearly* w.r.t. the above factors. This will allow us to single out a unique mock assignment. Whenever we focus our attention on a particular PCM, we will use the uppercase index I, keeping i for contexts where multiple PCMs are considered. In the remainder of this section and the next one, we consider the assignment from the point of view of a particular PCM I.

We make a distinction between two cases: if $cp_I > k_I$ we say that I *overbids*; and if $cp_I < k_I$ we say that she *underbids*. If $cp_I = k_I$ then we say that the bid is *exact*, and this can be considered both as a weak overbid or weak underbid.

Similarly, a paper can be either *overdemanded* (if $d_j > r$) or *underdemanded* (if $d_j < r$). We thus define:

$$u_j := [r - d_j]_+, \; o_j := [d_j - r]_+. \tag{1}$$

Recall that we also defined $p_j := \min\{1, \frac{r}{d_j}\}$, thus overdemanded papers have price $p_j = 1$.

Definition 3.1 (PMA). $X_I = (x_{Ij})_{j \in M}$ is a *proportional mock assignment (PMA) for PCM I* w.r.t. input $B_I = (f_{Ij})_{j \in M}, \boldsymbol{q}_I = (q_{Ij})_{j \in M}, \boldsymbol{d} = (d_j)_{j \in M}$ if it is:

full $\sum_{j \in M} x_{Ij} = k_I$;
valid For all $j \in M$, $x_{Ij} \leq \min\{r, q_{Ij}\}$;
proportional There is a constant $\alpha \leq 1$ such that:
- **(OB)** If the PCM is weakly overbidding, then $x_{Ij} = \alpha \cdot f_{Ij} p_j$ for all $j \in M$;
- **(UB)** If the PCM is weakly underbidding, then $x_{Ij} = \min\{q_{Ij}, f_{Ij} p_j + \alpha \cdot u_j\}$ for all $j \in M$.

Given an underbidding PMA X_I with its constant α, we denote by $Q_I := \{j \in M : f_{Ij} p_j + \alpha \cdot u_j > q_{Ij}\}$ as the set of *strictly constrained papers*. We similarly define $\tilde{Q}_I$ with a weak inequality, see Example 3.4.

A few explanations are in order. If I is overbidding, then the assignment is simply proportional to bids weighted by paper prices.

If I is underbidding, however, an assignment proportional to bids weighted by prices may both violate validity (if α is too large) and fail to be full (if α is too small). In that case, a proportional assignment will assign first papers for which the reviewer has placed a bid, weighted by their price, and then completes reviewer i's assignment by giving her a fraction of underdemanded papers, proportionally to their degree of underdemand u_j, while making sure that the quota constraints are not violated.

Existence and uniqueness of PMA. A proportional mock assignment is not guaranteed to exist. We give a necessary and sufficient condition for its existence and prove that in this case it is unique.

Definition 3.2 (Initial assignment). Define $\overline{x}_{Ij} := f_{Ij}p_j$ as the *initial (partial) assignment*. We say that $\overline{X}_I = (\overline{x}_{Ij})_{j\in M}$ is *extendable* w.r.t. input $(B_I, \boldsymbol{q}_I, \boldsymbol{d})$ if there is a valid and full X_I such that $x_{Ij} \leq \overline{x}_{Ij} + u_j$ for all j.

The initial partial assignment $\overline{X}_I$ simply assigns each paper proportionally to the bid and the price, and may assign more than k_I papers to I. Considering the leftovers u_j of all underdemanded papers, assuming extendability means that there is *some way* to partially allocate these leftovers, so that I obtains a complete and valid assignment.

Extendability can be violated in extreme cases, e.g. if I does not bid at all and still all papers are overdemanded. A sufficient condition extendability is either that I overbids, or that

$$k_I \leq \sum_{j\in M} min\{q_{Ij}, u_j\}.$$

Since k_I is bounded by 10-15 and the sum on the right hand side tends to grow linearly with the size of the conference, a violation is highly unlikely in any practical situation.[4]

Thus, **for the rest of the paper on we will assume that extendability holds** (for PCM I).

We next show that a PMA always exists as long as the initial partial assignment is extendable. Further, in this case the PMA is unique.

Proposition 3.3. *Given bid B_I, quotas $\boldsymbol{q}_I$ and demands $\boldsymbol{d}$, there exists a PMA if and only if the initial assignment $\overline{X}_I$ is extendable.*

The proof (in Appendix C) is constructive and relies on a simple algorithm that roughly proceeds as follows: it first computes the initial assignment $\overline{X}_I$; then, if I is overbidding then α is a simple normalization factor; if I is underbidding, α is computed in a series of iterations (at most m). The details of the algorithms are relegated to Appendix A. Here we give a small example (the details of how it is computed by the algorithm are in Appendix B).

Example 3.4 (Proportional Mock Assignment). $n = 5, m = 6$ (papers are called a, b, c, d, e, f) $r = 2$, and all quotas are 1. We use PCMs with different k_i for exposition purposes only. The bids f_{ij} are shown in Figure 1 (left) and the initial assignment $\overline{X}$ is shown on Figure 1 (right).

PCM 1 is overbidding, and her PMA is $x_{1j} = \alpha \cdot f_{1j}p_j$ with $\alpha = \frac{60}{67}$. PCM 2 is underbidding, and her PMA is $x_{Ij} = \min\{q_{Ij}, f_{Ij}p_j + \alpha \cdot u_j\}$ with $\alpha = \frac{2}{5}$, and the constrained papers are $Q_I = \tilde{Q}_I = \{e\}$. Thus

$$X_1 = \left(\frac{40}{77}, \frac{30}{77}, \frac{24}{77}, \frac{60}{77}, 0, 0\right) \qquad X_2 = \left(0, \frac{1}{2}, \frac{2}{5}, \frac{2}{5}, 1, \frac{7}{10}\right)$$

See Appendix B for the complete exposition of the example.

Note that computing the PMA independently for all PCMs may not result in a valid assignment, as some papers may be allocated less or more than r times (see Appendix B). Yet, from the viewpoint of a particular PCM that does not know the exact bid matrix, or even the assignment algorithm, this is a reasonable abstraction.

Proposition 3.5. *Consider input $B_I = (f_{Ij})_{j\in M}, \boldsymbol{q}_I = (q_{Ij})_{j\in M}, \boldsymbol{d} = (d_j)_{j\in M}$, and suppose that $\overline{X}_I$ is extendable. Then the PMA is unique.*

The proof is in Appendix C.

PCMs expect that bidding on a paper would increase the chance of getting this paper. While in actual assignment algorithms there may be corner cases that violate this, we want to show that the PMA does follow this natural expectation, which plays a key role in the game-theoretic analysis.

Definition 3.6 (assignment monotonicity). Consider two bids $B_I = (f_{Ij})_j$, $B'_I = (f'_{Ij})_j$ for which a (unique) PMA exist, where for some specific j, $f'_{Ij} < f_{Ij}$ whereas all other bids are the same. An assignment is *monotone in bids* if for any such B_I, B'_I, the respective obtained assignments X_I, X'_I satisfy:

MON1 $x'_{Ij} \leq x_{Ij}$;
MON2 $x'_{Ij'} \geq x_{Ij'}$ for all $j' \neq j$;

That is, PCM I gets no more of paper j and no less of all other papers.

Proposition 3.7. *Both the PMA and the initial assignments are monotone in bids. Moreover, for the initial assignment MON1 holds with a strict inequality.*

Proposition 3.8. *The (unique) PMA is continuous in the bid.*

Both proofs are in Appendix C.

4 INCENTIVES

The bidding process can be thought of as a game, where every time a PCM logs in, she sees the current prices (which reflect current demand) and reacts with her own bid. Given demands $\boldsymbol{d}^{-I}$, every bid B_I induces a PMA X_I (as I does not know — and cannot know — the actual assignment), and therefore some expected utility for PCM I. Crucially, we assume that the quotas $\boldsymbol{q}_I$ and capacity k_I are provided externally, and are not part of the strategy of the PCM.

In order to define this utility, we assign a cost C_{Ij} reflecting the inconvenience of reviewing paper j to PCM I. We assume that costs are *generic*, i.e. no two papers have the same cost. Taking quotas and capacities as fixed, we define

$$c_I(B_I, \boldsymbol{d}^{-I}) := \sum_{j\in M} x_{Ij}C_{Ij},$$

where $X_I = (x_{Ij})_{j\in M}$ is the unique PMA corresponding to bid B_I (recall that we assume the initial assignment is extendable and thus the PMA exists).

[4]Not a single violation occurred in any of our simulations, for example.

f_{ij}	a	b	c	d	e	f	k_i
1	1	1	1	1			2
2		1	1		4/5	1/2	3
3	1		1			1	3
4		1	1				2
5	1	1	1				2
d_j	3	4	5	1	4/5	3/2	
p_j	2/3	2/4	2/5	1	1	1	

$\overline{x}_{ij}$	a	b	c	d	e	f	ub_i	ob_i
1	2/3	1/2	2/5	1				17/30
2		1/2	2/5		4/5	1/2	4/5	
3	2/3		2/5			1	14/15	
4		1/2	2/5				11/10	
5	2/3	1/2	2/5				13/30	
u_j	0	0	0	1	6/5	1/2		

Figure 1: Bids (left) an initial assignment (right) in Example 3.4.

A *best response* of I to $\boldsymbol{d}^{-I}$ is a bid B_I such that $c_I(B_I, \boldsymbol{d}^{-I}) \leq c_I(B'_I, \boldsymbol{d}^{-I})$ for all B'_I. Note that a best response always exists since the strategy sets are compact.

Given $\boldsymbol{d}^{-I}$, two bidding strategies B_I, B'_I are said to be *equivalent* if they induce the same PMA, *i.e.*, $X_I = X'_I$.

Remark 2. *If X_I is not full, then the cost is not well defined—presumably, such a PCM will be asked to review some papers outside the assignment mechanism. However recall that we explicitly assume that the initial assignment is extendable to a full assignment, which is sufficient for the existence of a unique PMA (and in particular full).*

4.1 Sincere and exact bids

By genericity, I has a strict preference order over papers. We say that a bid $B_I = (f_{Ij})_{j \in M}$ is *sincere* if there is j' such that $f_{Ij} = 1$ for all papers that I prefers over j', and $f_{Ij} = 0$ for all papers that I prefers less than j'.

Note that a sincere bid can be characterized by a single number b_I, meaning that the PCM bids $f_{Ij} = 1$ on her favorite $\lfloor b_I \rfloor$ papers and $b_I - \lfloor b_I \rfloor$ on the next paper.

Recall that a bid $B_I = (f_{Ij})_{j \in M}$ is exact under prices $(p_j)_{j \in M}$,[5] if $\sum_{j \in M} f_{Ij} p_j = k_I$, that is, the PCM neither overbids nor underbids. Under the genericity assumption, there is a unique sincere exact bid, where $b_I = k_I$. We denote the sincere exact bid of I (w.r.t. $\boldsymbol{d}^{-I}, \boldsymbol{q}_I, k_I$) by b^*_I.

Observation 4.1. *If B_I is exact, then $x_{Ij} = f_{Ij} p_j$. In particular, if I places a full bid on paper j ($f_{Ij} = 1$) then $x_{Ij} = p_j$, i.e., the price reflects exactly the fraction (or probability) of paper j that PCM I will get if bidding on the paper.*

Observation 4.1 suggests that among all *exact bids* it is always best to be sincere. However, it does not rule out the possibility that the PCM overbids or underbids to improve the set of allocated papers. Also, changing the bid also changes the prices so the bid may not remain exact.

Our main theoretical result is that any best response for I is equivalent to b^*_I. The remainder of this section is used to prove this claim in two steps: first we show that any optimal bid is equivalent to a sincere bid (Prop. 4.2). Then, we show that among sincere bids it is always better to be exact.

Proposition 4.2. *Suppose that B_I is a best response, then B_I is equivalent to a sincere best response.*

[5] Note that the prices are affected by the entire demand, which in turn depends on the bid B_I.

Proof. We measure the *insincerity* of a bid B_I as follows. Let $j' = \text{argmin}_{j \in M: f_{Ij} < 1} C_{Ij}$ and $j'' = \text{argmax}_{j \in M: f_{Ij} > 0} C_{Ij}$. If $j' \geq j''$ then $IN(B_I) := 0$ (bid is sincere). Otherwise, $IN(B_I) := j'' - j' - f_{Ij'} + f_{Ij''}$, We say that strategy B'_i is *more sincere* than B_i if $IN(B'_i) < IN(B_I)$.

By continuity, the set of best responses is compact. Thus let B_I be a best response with minimal $IN(B_I)$. We argue that $IN(B_I) = 0$, otherwise we can construct an equivalent bid that is strictly more sincere.

By monotonicity, increasing the demand of j' results in $x'_{Ij'} \geq x_{Ij'}$, whereas the allocation of all other papers weakly decreases. If $x'_{Ij'} = x_{Ij'}$ then this means B'_I is equivalent to B_I but more sincere. Similarly, if we can decrease $f_{Ij''}$ without affecting $x_{Ij''}$ we are done. Thus assume that both modifications result in a strict increase in $x_{Ij'}$ and a strict decrease in $x_{Ij''}$.

Case I: PCM is weakly overbidding. Let $\delta \in (0, \min\{1 - f_{Ij'}, f_{Ij''}\})$ sufficiently small (will be specified later). Consider the bid $B'_I = (f'_{Ij})_{j \in M}$ defined by $f'_{Ij'} = f_{Ij'} + \delta$. This is clearly an overbid, which results in a new PMA X'_I with some $\alpha' = \alpha - \varepsilon$, where $\varepsilon = \varepsilon(\delta)$. Since $\overline{x}'_{Ij'} > \overline{x}_{Ij'}$, $\varepsilon > 0$.

By continuity, for small ε there is $\mu = \mu(\varepsilon)$ such that setting $B''_I := ((f'_{Ij})_{j \neq j''}), f_{Ij''} - \mu)$ results in a PMA X''_I with $\alpha'' = \alpha' + \varepsilon$, unless ε is too large (again, $\alpha'' > \alpha'$ since $\overline{x}'_{Ij''} < \overline{x}_{Ij''}$). We thus set $\delta > 0$ sufficiently small so that $\mu(\varepsilon(\delta))$ exists, and get $\alpha'' = \alpha' + \varepsilon = \alpha$. Since $\overline{X}''_{Ij} = \overline{X}_{Ij}$ for all $j \neq j', j''$, and I gets strictly more of j' and strictly less of j'', we get $c_I(B''_I, \boldsymbol{d}^{-I}) < c_I(B_I, \boldsymbol{d}^{-I})$, in contradiction to B_I being a best response.

Case II: PCM is strictly underbidding. By assumption, setting $f'_{Ij'} = f_{Ij'} + \delta$ strictly increases $x_{Ij'}$, and thus $\alpha' = \alpha - \varepsilon$ where $\varepsilon = \varepsilon(\delta) > 0$. Note that this entails that $j' \notin \tilde{Q}_I$, and thus we can set δ sufficiently small so that $j' \notin \tilde{Q}'_I$ either. The exact opposite holds for decreasing j''. What we would want is to do both and keep α at the same level. However we need to prove that this is possible—that is, that increasing demand on j' does not eliminate strict monotonicity of $x_{Ij''}$. Indeed, failure of strict monotonicity of $x_{Ij''}$ occurs either if (a) j'' is strictly constrained, or if (b) all papers j with positive excess u_j (except j'' itself) are strictly constrained. In all other cases, decreasing $f_{Ij''}$ increases k_I and thus the assignment of *some* paper $j \neq j''$ strictly increases, which entails strict monotonicity.

However, increasing the demand on j' can only make other papers *less constrained*, i.e. if $j \notin Q_I$ then $j \notin Q'_I$. The only paper whose positive excess may change by increasing $f_{Ij'}$ is j' itself, but again this can only decrease u_j. Thus if we can set $f'_{Ij''} = f_{Ij''} - \tau$

to reach $\alpha'' = \alpha - \mu$ with respect to B_I, then we can still reach $\alpha'' = \alpha' - \mu$ (i.e., w.r.t. bid B'_I) by setting $f'_{Ij''} = f_{Ij''} - \tau'$.

To conclude the proof, we set $\mu = \varepsilon$ sufficiently small so that both $\tau = \tau(\mu)$ and $\delta = \delta(\varepsilon)$ exist (and such that $j' \notin \tilde{Q}'_I$). In the bid $B''_I = (f_{Ij'} + \delta, f_{Ij''} - \tau', (f_{Ij})_{j \neq j', j''})$, we get that $Q''_I = Q_I$ and $\alpha'' = \alpha' - \mu = \alpha + \varepsilon - \mu = \alpha$. Thus X'_I is same as X_I except it has strictly more of j' and strictly less of j''. We get that $c_I(B''_I, \boldsymbol{d}^{-I}) < c_I(B_I, \boldsymbol{d}^{-I})$, in contradiction to B_I being a best response. □

Proposition 4.3. *Any best response is equivalent to a sincere and exact bid.*

Proof. We start from an arbitrary sincere best response $b_I \in \mathbb{R}_+$. If b_I is an overbid, then $b'_I < b_I$ reduces the cost, as long as b'_I is still a weak overbid. This is since by proportionality, the fraction of the least favorite papers in $\overline{X}_I$ is strictly decreasing. Therefore no best response can be overbidding.

Suppose now that b_I is an underbid. Label papers by increasing C_{Ij}, and denote by $j' := \text{argmin}_{j \in M: f_{Ij} < 1} C_{Ij}$ the first paper with less-than-maximal bid (as in the previous proof). Note that $j' = \lfloor b_i \rfloor + 1$ We define $b'_I = b_I + \delta$, where $\delta > 0$ is sufficiently small so that only paper j' is affected, and PCM I is still weakly underbidding in b'_I.

By monotonicity, $x'_{1j'} \geq x_{Ij'}$, whereas $x'_{Ij} \leq x_{Ij}$ for all other papers. We argue that for all $j < j'$, $x'_{Ij} = x_{Ij}$. This would complete the proof, as

$$
\begin{aligned}
c_I(b'_I, \boldsymbol{d}^{-I}) - c_I(b_I, \boldsymbol{d}^{-I}) &= \sum_j (x'_{Ij} - x_{Ij}) C_{Ij} = (x'_{Ij'} - x_{Ij'}) C_{Ij'} + \sum_{j > j'} (x'_{Ij} - x_{Ij}) C_{Ij} \\
&\leq (x'_{Ij'} - x_{Ij'}) C_{Ij'} + \sum_{j > j'} (x'_{Ij} - x_{Ij}) C_{Ij'} \\
&= C_{Ij'} \sum_{j \geq j'} (x'_{Ij} - x_{Ij}) = C_{Ij'} \sum_j (x'_{Ij} - x_{Ij}) \\
&= C_{Ij'} (\sum_j x'_{Ij} - \sum_j x_{Ij}) = C_{Ij'} (k_I - k_I) = 0.
\end{aligned}
$$

where the inequality comes from $x'_{Ij} - x_{Ij} \leq 0$ by monotonicity, and $C_j > C_{j'}$ for all $j > j'$.

Indeed, consider some paper $j < j'$, and note that $f_{Ij} = 1$. If $p_j = 1$ then $\overline{x}_{1j} = f_{1j} p_j = 1$, and this is still the case for $\overline{x}'_{1j}$. Then, since $x_{1j} \in [\overline{x}_{1j}, q_{1j}] = [1, 1]$, we get $x_{1j} = 1 = x'_{1j}$.

If $p_j < 1$, then $u_j = 0$ (paper is overdemanded). Also note that the bid on j has not changed and thus

$$\overline{x}'_{1j} = f'_{1j} p'_j = f_{1j} p_j = \overline{x}_{1j}.$$

Since $x_{1j} \in [\overline{x}_{1j}, \overline{x}_{1j} + u_j] = [\overline{x}_{1j}, \overline{x}_{1j}]$, we have $x'_{1j} = \overline{x}'_{1j} = \overline{x}_{1j} = x_{1j}$, as required. □

4.2 Price-sensitive bids

Consider a bidder facing papers with private costs $(2, 3, 4, 5, 6)$ and respective prices $(0.6, 0.6, 0.15, 1, 0.6)$. Suppose the bidding requirement is $R = 2$, so a sincere strategy would bid either on the first 3 or first 4 papers.

However in practice, inferring a good estimate of one's own private cost requires some effort (e.g. reading the paper's abstract), whereas the price in our suggested mechanism is presented explicitly. Since papers with a low price have low assignment probability, and also contribute little to meet the bidding requirement, the PCM may skip low-price papers such as paper #3 above, and not consider them for bidding at all (or at least not until she exhausted the high-price papers).

We consider two bidding behaviors that may capture this effect.

Sincere bidding + Costly exploration. Preferring high-price papers may be due to uncertainty on their true cost. Consider a model where for each paper j and PCM I, there is a distribution C_{Ij} from which the "true" cost c_{Ij} will be realized. During bidding, the PCM can decide to invest some effort, and then observe the realization c_{Ij} (e.g. by reading the abstract). For simplicity, suppose that the realization is either 0 or 1, meaning that the paper is a good fit or a bad fit for the reviewer. In that simple case, C_{Ij} is the probability of a bad fit.

A myopic PCM would maintain a "current bid" which is optimal according to known costs (explored or expected), and thus exact and sincere. Then, the PCM will gradually explore papers until exploration is no longer beneficial.

Suppose that the current optimal (sincere and exact) bid is some B^*_i, whereas after exploring paper j, the new optimal bid is $B^{(j)}_i$ (which is also sincere, but according to the revealed cost).

Each paper has some "exploration effort" e_j, and some *marginal exploration gain* $mg_j := c_i(B^{(j)}_i, \boldsymbol{d}^{-i}) - c_i(B^*_i, \boldsymbol{d}^{-i})$. The optimal bidding policy is clearly to explore the paper maximizing $mg_j - e_j$ as long as this value is positive, and submit bid B^*_i once all papers yield non-positive improvement.

We argue that papers with a higher price should be explored earlier, ceteris paribus, and thus the PCM is likely to skip more low-price papers.

Proposition 4.4. *The marginal gain mg_j is weakly increasing in p_j.*

Proof. Denote by G the set of all explored papers that turn out to be good, i.e. with a known cost of 0. The optimal bid is to bid on papers from G, and on unexplored paper in increasing order of C_{Ij},until the bidding requirement is reached: $\sum_{j \in M} f_{Ij} p_j = k_I$.

Note that since agent I bids are always exact, she receives under the PMA a fraction of $f_{Ij} p_j$ of each paper j. To maintain an exact bid, if she increases f_j then she has to decrease her bids on other papers (by sincerity, on the ones with the largest expected cost).

Denote by z the worst paper with nonzero bid. Note that $f_{Ij} = 1$ for all $j < z$ and $f_{Ij} = 0$ for all $j > z$ (we assume here that papers are ordered by non decreasing expected costs). There are three cases when exploring the next paper j:

- $f_{Ij} = 0$. That is, j is not in the current bidding set at all. In this case, with probability C_{Ij} the paper turns out to be bad , i.e. $C_{Ij} = 1$, and the bid does not change. With probability $1 - C_{Ij}$ the paper turns out to be good. Then it is added to the bidding set ($f'_{Ij} = 1$) and pushes out of the bidding set a fraction p_j of the highest-cost papers (possibly including other 0-cost papers). Denote by $y(F)$ the total cost of the F most costly papers in B_I , then $c_I(B^*_I, \boldsymbol{d}^{-I}) - c_I(B^{(j)}_I, \boldsymbol{d}^{-I}) = y(p_j)$. Then the marginal gain when exploring j is $mg_j = (1 - C_{Ij}) y(p_j)$, which is monotonically increasing in p_j.
- $f_{Ij} = 1$. That is, $j < z$ is in the current bidding set. With probability $1 - C_{Ij}$ the paper turns out to be good, the bid

does not change, and the expected cost (after receiving a new information on paper j) becomes $c_I(B^*_I, \boldsymbol{d}^{-I}) - F \cdot C_{Ij}$. With probability C_{Ij} the paper turns out to be bad and is removed from the bidding set, making room for the an F fraction of the least costly papers that are not in the bidding set. In this case, $p_j \cdot C_{Ij}$ is eliminated from $c_I(B^*_I, \boldsymbol{d}^{-I})$. Now denote by $\hat{y}(F)$ the total cost of these papers, then the cost of the new bidding set is

$$c_I(B^*_I, \boldsymbol{d}^{-I}) - p_j \cdot C_{Ij} + \hat{y}(p_j) = c_I(B^*_I, \boldsymbol{d}^{-I}) + F(\frac{\hat{y}(p_j)}{p_j} - C_{Ij}).$$

In total

$$mg_j = c_I(B^*_I, \boldsymbol{d}^{-I}) - E_{C_{Ij}}[c_I(B^{(j)}_I, \boldsymbol{d}^{-I})] = -(1 - C_{Ij})(-p_j \cdot C_{Ij}) - C_{Ij}F(\frac{\hat{y}(p_j)}{p_j} - C_{Ij})$$

$$= p_j \cdot C_{Ij}((1 - C_{Ij}) + (C_{Ij} - \frac{\hat{y}(p_j)}{p_j})) = FC_{Ij}(1 - \frac{\hat{y}(p_j)}{p_j})$$

Therefore

$$\frac{\partial mg_j}{\partial p_j} = C_{Ij} - C_{Ij}\frac{\partial \hat{y}(F)}{\partial F} \geq C_{Ij}(1 - 1) = 0,$$

as required, since $\frac{\partial \hat{y}(F)}{\partial F} \leq \max_{j'} C_{ij'} \leq 1$.

- $j = z$. We analyze the marginal gain as a sum of two exploration actions: the last paper in the bidding set with price $p_j f_{Ij}$; and the first paper outside the set with price $p_j(1-f_{Ij})$. Since both parts are non-decreasing in p_j we get the result from linearity of expectation.

□

That is, cost-wise, it is better to explore papers close to the decision threshold (high-cost in the bidding set or low-cost out of the set), and for two papers with a similar uncertainty over their cost, it is better to explore first the one with the higher price.

Greedy price-sensitive bidding. The "attractiveness" of each paper is some function that is increasing (linearly) in price and decreasing (linearly) in private cost. The bidder selects papers greedily by decreasing attractiveness, until reaching the bidding quota.

The greedy heuristic is easy to apply and similar to how a rational decision maker would behave under quasi-linear utilities, or in other contexts of selecting multiple items [4].

4.3 Multiple bid levels

It is both a conceptual and technical question whether a "strong bid" should be considered differently when computing papers' prices.

The technical part is easy: we may allow f_{ij} to take any nonnegative value. Thus a PCM i that bids $f_{ij} = 2$ will get in the initial assignment twice the amount of paper j than bidder i' that bids $f_{i'j} = 1$, at any given price. PCM i also contributes more to the demand.

As for incentive analysis, note that if multi-level bids are allowed, then the PMA is still well defined with slight modifications (we need to make sure that the initial assignment does not violate the quota). Then following almost the same proofs shows that bidders still have an incentive to be sincere, in the sense that they will give the maximal possible bid on some t most preferred papers, and 0 bid on all papers after $t + 1$. The proofs for exact bids also remain almost.

In practice, various considerations may guide the decision of the bid strength, such as how competent the PCM believe they are for this particular review.

4.4 Dynamics and Equilibrium

A *behavior* of bidder i is a mapping from a profile $\boldsymbol{B}$ and private costs C_i to a new bid B'_i. A profile $\boldsymbol{B}$ is an *equilibrium* if for all $i \in N$, B_i is a best response to $\boldsymbol{d}^{-I}$ (note that $\boldsymbol{d}^{-I}$ is induced by $\boldsymbol{B}$).

Increasing Bids. We say that a sequence of bidding profiles $\boldsymbol{B}^0, \dots, \boldsymbol{B}^T$ is *weakly increasing* if $f^{t+1}_{ij} \geq f^t_{ij}$ for all $i \in N, j \in M, t < T$.

One way to justify increasing bids is that reviewers rarely bother to eliminate an existing bid. For the rational behavior we can also deduce this from the properties of the best-response strategy:

Observation 4.5. *If in $\boldsymbol{B}^t$ all bidders are sincere and weakly underbid, and some set of bidders $N' \subseteq N$ update their bids by playing an exact sincere bid, then:*

(1) $f^{t+1}_{ij} \geq f^t_{ij}$ for all $i \in N', j \in M$, and all other bids remain unchanged;
(2) all prices weakly decrease in $\boldsymbol{B}^{t+1}$;
(3) all bidders are sincere and weakly underbid in $\boldsymbol{B}^{t+1}$.

Therefore, if the initial bid is sincere and below the bidding requirement (e.g. the empty profile) then by induction the bidding sequence is weakly increasing. Such a bidding sequence must converge to an equilibrium since bids are bounded. In particular, an equilibrium exists (recall that this is contingent on our assumption that a PMA always exists).

5 IN-SILICO EXPERIMENTS

For space all experiments are in the Appendix.

6 DISCUSSION

In this paper we propose a simple alternative bidding mechanism for reviewers in conferences: for every paper, the reviewers observe a dynamic "price" that reflects the paper's demand and updates throughout the bidding process. We showed via the introduction of a stylized assignment that changes linearly with the bids (the PMA), that bidders have an incentive to follow instructions and sincerely bid on their favorite papers until they reach the bidding requirement (in terms of overall price of papers).

Interestingly, sincere bidding is *not* socially favorable. In contrast, a very simple greedy behavior—where reviewers are less likely to bid on cheap (overdemanded) papers—is cognitively plausible, reduces the effort *during* bidding, and substantially improves the allocation under a broad range of conditions.

Our suggestions can be very easily implemented in existing conference management platforms such as EasyChair and ConfMaster, and allow reviewers to obtain a more preferred lot of papers for less effort. Crucially, we show that the improvement is gradual, and there is no harm if some of the reviewers ignore paper prices altogether and just bid as they are used to.

Privacy and bias. One possible concern is that information on the demand may bias the judgment of the PCMs, or that authors may be unhappy with showing information about the popularity of their paper. We argue that these concerns are unjustified. First, there is no obvious connection between paper popularity and quality, as PCMs may bid on a paper either because it is relevant to their specific interests, because it looks attractive, or because it seems to be an easy reject(see also Rodriguez et al. [18]). Second, the demand information is shown at a very crude accuracy: the PCM only observes some arbitrary snapshot, and even then cannot differentiate e.g. a paper with 0 bids from a paper with r bids. Adding the uniform bootstrap (as we do in the simulations and recommend doing also in practice) demand makes any such inference even less likely, especially if we add some noise to it. Finally and perhaps most importantly, unpopular papers are not damaged by revealing information on their demand. On a contrary, they are more likely to get suitable reviewers in our suggested bidding scheme, whereas in the current system the paper will end up with a possibly disappointed PCM how did not ask for it.

Next steps. We are currently designing lab experiments that will help us understand the actual paper-bidding behavior of people with and without prices. We also plan a field experiment in a medium-size workshop where part of the program committee will bid via the suggested bidding scheme.

REFERENCES

[1] Faez Ahmed, John P. Dickerson, and Mark Fuge. 2017. Diverse Weighted Bipartite b-Matching. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017.* 35–41.

[2] Haris Aziz, Xin Huang, Nicholas Mattei, and Erel Segal-Halevi. 2019. The Constrained Round Robin Algorithm for Fair and Efficient Allocation. *CoRR* abs/1908.00161 (2019). arXiv:1908.00161 http://arxiv.org/abs/1908.00161

[3] Moshe Babaioff, Noam Nisan, and Inbal Talgam-Cohen. 2017. Competitive equilibria with indivisible goods and generic budgets. *arXiv preprint arXiv:1703.08150* (2017).

[4] James R. Bettman, Eric J. Johnson, and John W. Payne. 1991. Consumer Decision Making. In *Handbook of Consumer Behavior*, Thomas S. Robertson and Harold H. Kassarjian (Eds.). Prentice-Hall.

[5] Eric Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.

[6] Laurent Charlin, Richard S. Zemel, and Craig Boutilier. 2011. A Framework for Optimizing Paper Matching. In *UAI 2011, Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence, Barcelona, Spain, July 14-17, 2011.* 86–95.

[7] Don Conry, Yehuda Koren, and Naren Ramakrishnan. 2009. Recommender systems for the conference paper assignment problem. In *Proceedings of the 2009 ACM Conference on Recommender Systems, RecSys 2009, New York, NY, USA, October 23-25, 2009.* 357–360.

[8] Naveen Garg, Telikepalli Kavitha, Amit Kumar, Kurt Mehlhorn, and Julián Mestre. 2010. Assigning Papers to Referees. *Algorithmica* 58, 1 (2010), 119–136.

[9] Judy Goldsmith and Robert H. Sloan. 2007. The AI Conference Paper Assignment Problem. In *Proceedings of the AAAI Workshop on Preference Handling for Artificial Intelligence.*

[10] David Hartvigsen, Jerry Wei, and Richard Czuchlewski. 2007. The Conference Paper-Reviewer Assignment Problem. *Decision Sciences* 30 (06 2007), 865 – 876.

[11] David Kurokawa, Omer Lev, Jamie Morgenstern, and Ariel D. Procaccia. 2015. Impartial Peer Review. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015.* 582–588.

[12] Jing Wu Lian, Nicholas Mattei, Renee Noble, and Toby Walsh. 2018. The Conference Paper Assignment Problem: Using Order Weighted Averages to Assign Indivisible Goods. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018.* 1138–1145.

[13] Cheng Long, Raymond Chi-Wing Wong, Yu Peng, and Liangliang Ye. 2013. On Good and Fair Paper-Reviewer Assignment. In *2013 IEEE 13th International Conference on Data Mining, Dallas, TX, USA, December 7-10, 2013.* 1145–1150.

[14] Nicholas Mattei and Toby Walsh. 2013. PrefLib: A Library for Preferences, HTTP://WWW.PREFLIB.ORG. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT).*

[15] Nicholas Mattei and Toby Walsh. 2017. A PrefLib.Org Retrospective: Lessons Learned and New Directions. In *Trends in Computational Social Choice*, U. Endriss (Ed.). AI Access Foundation, Chapter 15, 289–309.

[16] Jennifer Nguyen, Germán Sánchez-Hernández, Núria Agell, Xari Rovira, and Cecilio Angulo. 2018. A decision support tool using Order Weighted Averaging for conference review assignment. *Pattern Recognition Letters* 105 (2018), 114–120.

[17] Abraham Othman, Tuomas Sandholm, and Eric Budish. 2010. Finding approximate competitive equilibria: Efficient and fair course allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1.* International Foundation for Autonomous Agents and Multiagent Systems, 873–880.

[18] Marko A. Rodriguez, Johan Bollen, and Herbert Van de Sompel. 2007. Mapping the bid behavior of conference referees. *J. Informetrics* 1, 1 (2007), 68–82.

[19] Yichong Xu, Han Zhao, Xiaofei Shi, and Nihar B Shah. 2018. On Strategyproof Conference Peer Review. *arXiv preprint arXiv:1806.06266* (2018).

A THE FRACTIONAL MOCK ASSIGNMENT ALGORITHM

The algorithm specifies the final allocation of PCM I, as a function of her bid (f_{Ij} for all j) and total demands of the other PCMs, denoted $d_j^{-I} := d_j - f_{Ij}$ (for all j).

ALGORITHM 1: Fractional Mock Assignment for PCM I

Input: bid $(f_{Ij})_{j \in M}$; quotas $(q_{Ij})_{j \in M}$; capacity k_I; demands without I $(d_j^{-I})_{j \in M}$; r.
Output: I's assignment $X_I = (x_{Ij})_{j \in M}$
$\forall j \in M$, set $d_j := d_j^{-I} + f_{Ij}$; `// add up the demands of all PCMs`
$\forall j \in M$, set $p_j := \min\{1, \frac{r}{d_j}\}$;
`/* Step IA: Initial Assignment. */`
$\forall j \in M$, set $\overline{x}_{Ij} := f_{Ij} p_j$;
if $\sum_{j \in M} \overline{x}_{Ij} > k_I$ **then**
  `/* Step OB: reduce assignment in case of an OverBid. */`
  Set $ob_I := \sum_{j \in M} \overline{x}_{Ij} - k_I$; `// The overbid of PCM I`
  $\alpha := \frac{k_I}{k_I + ob_I}$;
  $\forall j \in M$, set $x_{Ij} := \alpha \cdot \overline{x}_{Ij}$;
end
else
  `/* Step UB: allocate leftovers of unassigned papers in case of an UnderBid. */`
  $\forall j \in M$, $u_j := [r - d_j p_j]_+$; `// unassigned leftovers of paper j`
  $\forall j \in M$, $x_{Ij} = \overline{x}_{Ij}$;
  repeat
    $Q := \{j \in M : x_{Ij} = q_{Ij}\}$; `// constrained papers`
    $\overline{k}_I := k_I - \sum_{j \in Q} q_{Ij} - \sum_{j \in M \setminus Q} \overline{x}_{Ij}$; `// Current available space of PCM I`
    $\eta := \max\{\overline{k}_I, \sum_{j \in M \setminus Q} u_j\}$; `// normalization factor`
    $\alpha := \frac{\overline{k}_I}{\eta}$;
    $\forall j \in M \setminus Q$, $x_{Ij} := \min\{q_{Ij}, \overline{x}_{Ij} + \alpha \cdot u_j\}$;
    $Q^+ = \{j \in M : \overline{x}_{Ij} + \alpha \cdot u_j > q_{Ij}\}$
  until $Q^+ = \emptyset$;
end
return X_I;

Intuitively, we can think of p_j as the *portion* (or *probability*) that a bidder putting a full bid ($f_{ij} = 1$) on j will get from j in step IA .

The more complicated case is when I is underbidding. Then we allocate unassigned papers to I, but need to be careful to respect quota and capacity constraints. Step UB repeatedly tries to allocate the remaining papers as follows: first, allocate the maximal possible amount of constrained papers $j \in Q$; then, for any $j \in M \setminus Q$, allocate the leftovers of paper j (an amount of u_j) proportionally to the available space that the PCM has (denoted by $\overline{k}_I$). As long as some papers exceed their quota, the algorithm adds them to Q and repeats the process.

Of course, real assignment algorithms have to deal with fully allocating all papers, and in particular re-allocate excess papers from overloaded bidders (even if these papers are overdemanded). Computing the PMA independently for all PCMs may not result in a valid assignment, as some papers may be allocated less or more than r times. Yet, from the viewpoint of a particular PCM that does not know the exact bid matrix, or even the assignment algorithm, this is a reasonable abstraction.

B PROPORTIONAL MOCK ASSIGNMENT: COMPLETE EXAMPLE

We give here full details about how the PMA works on Example 3.4. Recall that $n = 5$, $m = 6$ (papers are called a, b, c, d, e, f) $r = 2$, all quotas are 1, and the bids are as follows (left):

f_{ij}	a	b	c	d	e	f	k_i
1	1	1	1	1			2
2		1	1		4/5	1/2	3
3	1		1			1	3
4		1	1				2
5	1	1	1				2
d_j	3	4	5	1	4/5	3/2	
p_j	2/3	2/4	2/5	1	1	1	

All PCMs: Step IA

We first compute the Step I allocation to all PCMs.

$\overline{x}_{ij}$	a	b	c	d	e	f	ub_i	ob_i
1	2/3	1/2	2/5	1				17/30
2		1/2	2/5		4/5	1/2	4/5	
3	2/3		2/5			1	14/15	
4		1/2	2/5				11/10	
5	2/3	1/2	2/5				13/30	
u_j	0	0	0	1	6/5	1/2		

PCM 1: Step OB

PCM 1 is overbidding because $\sum_{j \in M} f_{1j} p_j = \frac{77}{30} > k_1 = 2$. The final allocation is obtained in Step OB by normalizing so that $x_{1j} = \frac{2}{2 + \frac{17}{30}} \overline{x}_{1j}$, which gives the allocation vector

$$\left(\frac{40}{77}, \frac{30}{77}, \frac{24}{77}, \frac{60}{77}, 0, 0 \right)$$

.

PCM 2: Step UB

PCM 2 is underbidding because $\sum_{j \in M} f_{2j} p_j = \frac{11}{5} < k_2 = 3$.

The PMA tries to allocate the leftovers u_j of papers d, e and f. The remaining available review space for PCM2 after Step I is $\overline{k}_2 = ub_2 = \frac{4}{5}$, while the global degree of underassignment after Step I is $u_d + u_e + u_f = \frac{27}{10}$; therefore, $su = \max\{\overline{k}_2, u_d + u_e + u_f\} = \frac{27}{10}$, and the value of α at this step is $\frac{k_2}{su} = \frac{8}{27}$. We get $z_d = \alpha u_d = 1 \cdot \frac{4/5}{27/10} = \frac{8}{27}$; $z_e = \frac{6}{5} \frac{4/5}{27/10} = \frac{16}{45}$; and $z_f = \frac{1}{2} \frac{4/5}{27/10} = \frac{4}{27}$, which are the additional fractions of papers to be tentatively assigned to 2 at this stage, in addition to those already assigned in Step I. This gives $x_{2d} = \overline{x}_{2d} + z_d = \frac{8}{27}$; $x_{2e} = \overline{x}_{2e} + z_e = \frac{4}{5} + \frac{16}{45} = \frac{52}{45}$; and $x_{2f} = \overline{x}_{2f} + z_f = \frac{1}{2} + \frac{4}{27}$.

Paper e now exceeds its quota ($q_{2e} = 1$). Thus $e \in Q^+$ in this iteration and the algorithm goes into a second iteration.

In the second iteration, $e \in Q$ so $x_{2e} = 1$ and $\overline{k}_2 = 3 - 2 - \frac{2}{5} = \frac{3}{5}$. Also, $\sum_{j \notin Q} u_j = u_d + u_f = \frac{3}{2}$ thus $su = \max\{\frac{3}{5}, \frac{3}{2}\} = \frac{3}{2}$. Therefore $z_d = \frac{u_d \overline{k}_2}{su} = \frac{2}{3}\frac{3}{5} = \frac{2}{5}$; and $z_f = \frac{u_f \overline{k}_2}{su} = \frac{1}{2}\frac{3}{5}\frac{2}{3} = \frac{1}{5}$, and the final assignment is $x_2 = (0, \frac{1}{2}, \frac{2}{5}, 0 + \frac{2}{5}, 1, \frac{1}{2} + \frac{1}{5})$, that is,

$$\left(0, \frac{1}{2}, \frac{2}{5}, \frac{2}{5}, 1, \frac{7}{10}\right)$$

Note that it is a PMA for $\alpha = \frac{2}{5}$, and it sums up to $k_2 = 3$.

PCM 3: Step UB

PCM 3 is underbidding because $\sum_{j \in M} f_{3j} p_j = \frac{31}{15} < k_3 = 3$.

At the first iteration, paper f is already assigned to the quota and the fractional allocation of f to 3 will no longer change. We have $\overline{k}_3 = \frac{14}{15}$ and $\sum_{j \notin Q} u_j = \frac{11}{5}$, therefore $su = \frac{11}{5}$ and $\alpha = \frac{14}{33}$. The fractional assignment at the end of this iteration is

$$\left(\frac{2}{3}, 0, \frac{2}{5}, \frac{14}{33}, \frac{28}{55}, 1\right)$$

No paper exceeds the quota so this assignment is final. It is a PMA for $\alpha = \frac{14}{33}$ and it sums up to 2.

PCM 4: Step UB

PCM 4 is underbidding because $\sum_{j \in M} f_{4j} p_j = \frac{9}{10} < k_4 = 2$. We have $\overline{k}_4 = \frac{11}{10}$ and $\sum_{j \notin Q} u_j = \frac{27}{10}$, therefore $su = \frac{27}{10}$ and $\alpha = \frac{11}{27}$. One iteration is enough and leads to the assignment

$$\left(0, \frac{1}{2}, \frac{2}{5}, \frac{11}{27}, \frac{22}{45}, \frac{11}{54}\right)$$

It is a PMA for $\alpha = \frac{11}{27}$ and it sums up to 2.

PCM 5: Step UB

PCM 5 is underbidding because $\sum_{j \in M} f_{5j} p_j = \frac{47}{30} < k_5 = 2$.

We have $\overline{k}_5 = \frac{13}{30}$, $su = \frac{27}{10}$ and $\alpha = \frac{13}{81}$. One iteration is enough and leads to the assignment

$$\left(\frac{2}{3}, \frac{1}{2}, \frac{2}{5}, \frac{13}{81}, \frac{26}{135}, \frac{13}{162}\right)$$

It is a PMA for $\alpha = \frac{13}{81}$ and it sums up to 2.

If we gather all mock assignments to the five PCMs we find that the number of PCMs assigned to a, b, c, d, e and f are respectively 1.85, 1.88, 1.91, 2.25, 2.11 and 1.98, which shows that the individual mock assignments are not globally consistent.

C PROOFS

Proposition 3.3. *Given bid B_I, quotas $\boldsymbol{q}_I$ and demands $\boldsymbol{d}$, there exists a PMA if and only if the initial assignment $\overline{X}_I$ is extendable.*

Moreover, Algorithm 1 always outputs a valid and proportional assignment. If the initial assignment $\overline{X}_I$ is extendable. Then the output is also full, and thus a PMA. In particular, the set of weakly constrained papers $\tilde{Q}_I$ is the set Q that Alg. 1 reaches in the last iteration.

Proof. First note that the algorithm must terminate since at least one new item is added to Q at each step.

Note that proportionality holds by construction, as the algorithm uses the same α for all papers, namely the value of α at the last iteration.

For validity, note that for any paper j, at most r is allocated since $z_j = \frac{u_j \cdot \overline{k}_I}{\eta} \leq u_j$. If I is underbidding, then $x_{Ij} \leq q_{Ij}$ by construction. If I is overbidding, note that $f_{Ij} \leq q_{Ij}$ and $p_j \leq 1$, so Step IA does not violate the quota. Step OB only reduces the allocation.

It is left to show that the output assignment $X_I = (x_{Ij})_{j \in M}$ is full whenever $\overline{X}_I$ is extendable.

If I is overbidding then

$$\sum_{j \in M} x_{Ij} = \sum_{j \in M} \overline{x}_{Ij} \frac{k_I}{k_I + ob_I} = \frac{k_I}{k_I + ob_I} \sum_{j \in M} \overline{x}_{Ij} = k_I.$$

If I is underbidding, then first we can see that allocation does not exceed k_I. Let x_{Ij} obtained in the last iteration of the algorithm. Because this is the last iteration, $Q^+ = \emptyset$, i.e., $\overline{x}_{Ij} + \alpha u_j \leq q_{Ij}$ holds for all $j \in M \setminus Q$.

$$\begin{aligned}
\sum_{j \in M} x_{Ij} &= \sum_{j \in Q} q_{Ij} + \sum_{j \in M \setminus Q} \min(q_{Ij}, \overline{x}_{Ij} + \alpha u_j) \\
&= \sum_{j \in Q} q_{Ij} + \sum_{j \in M \setminus Q} x_{Ij} + \alpha u_j \\
&= (k_I - \overline{k}_I) + \overline{k}_I \frac{1}{\eta} \sum_{j \in M \setminus Q} u_j \\
&\leq (k_I - \overline{k}_I) + \overline{k}_I \frac{1}{\sum_{j \in M \setminus Q} u_j} \sum_{j \in M \setminus Q} u_j \\
&= k_I - \overline{k}_I + \overline{k}_I = k_I
\end{aligned}$$

Further, if the inequality is strict, this means that $\overline{k}_I > \sum_{j \in M \setminus Q} u_j$. Let $\hat{X}_I$ be an arbitrary extension of $\overline{X}_I$ (in particular valid and full), then for all $j \in M$, $\hat{x}_{Ij} \leq \overline{x}_{Ij} + u_j$. Also, since $\hat{X}_I$ is valid, $\hat{x}_{Ij} \leq q_{Ij}$.

Putting the inequalities together

$$k_I - \sum_{j \in Q} q_{Ij} - \sum_{j \in M \setminus Q} \overline{x}_{Ij} = \overline{k}_I > \sum_{j \in M \setminus Q} u_j \quad \Rightarrow$$

$$k_I > \sum_{j \in Q} q_{Ij} + \sum_{j \in M \setminus Q} (\overline{x}_{Ij} + u_j) \geq \sum_{j \in M} \min\{q_{Ij}, \overline{x}_{Ij} + u_j\} \geq \sum_{i \in M} \hat{x}_{Ij},$$

which means that $\hat{X}_I$ is not full, in contradiction to our assumption that $\overline{X}_I$ is extendable. □

Proposition 3.5. *Consider input $B_I = (f_{Ij})_{j \in M}, (q_{Ij})_{j \in M}, (d_j)_{j \in M}$, and suppose that $\overline{X}_I$ is extendable. Then the PMA is unique.*

Proof. For overbidders this is obvious, as any change in α will violate the full requirement.

For underbidders, we need to show first that there are no two PMAs with different constrained sets. Assume towards a contradiction that there are two distinct PMAs X_I, X'_I.

Recall that $\tilde{Q}_I = \{j \in M : f_{Ij}p_j + \alpha \cdot u_j \geq q_{Ij}\}$ is the set of weakly constrained papers. Suppose there is some paper $a \in \tilde{Q}'_I \setminus \tilde{Q}_I$. Then $x_{Ia} = \overline{x}_{Ia} + \alpha u_a < q_{Ia}$ since $a \notin \tilde{Q}_I$, and $x'_{Ia} = q_{Ia} \leq \overline{x}_{Ia} + \alpha' u_a$. Thus $\alpha < \alpha'$. This in particular entails that $\tilde{Q}_I \subseteq \tilde{Q}'_I$. We then get that

$$k_I = \sum_{j\in M} x_{Ij} = \sum_{j\in\tilde{Q}_I} q_{Ij} + \sum_{j\in\tilde{Q}'_I\setminus\tilde{Q}_I} (\overline{x}_{Ij} + \alpha u_j) + \sum_{j\in M\setminus\tilde{Q}'_I} \overline{x}_{Ij} + \alpha u_j$$
$$< \sum_{j\in\tilde{Q}_I} q_{Ij} + \sum_{j\in\tilde{Q}'_I\setminus\tilde{Q}_I} q_{Ij} + \sum_{j\in M\setminus\tilde{Q}'_I} (\overline{x}_{Ij} + \alpha u_j) \qquad (\text{since } \tilde{Q}'_I \setminus \tilde{Q}_I \neq \emptyset)$$
$$= \sum_{j\in\tilde{Q}'_I} q_{Ij} + \sum_{j\in M\setminus\tilde{Q}'_I} (\overline{x}_{Ij} + \alpha u_j)$$
$$\leq \sum_{j\in\tilde{Q}'_I} x'_{Ij} + \sum_{j\in M\setminus\tilde{Q}'_I} (\overline{x}_{Ij} + \alpha' u_j) = \sum_{j\in M} x'_{Ij} = k_I,$$

which is a contradiction. Then, two PMAs with the same $\tilde{Q}_I$ and different α cannot be both full. □

Proposition 3.7. *Both the PMA and the initial assignments are monotone in bids. Moreover, for the initial assignment MON1 holds with a strict inequality.*

Proof. Without loss of generality, we prove for $j = 1$. Denote $\delta = f_{I1} - f'_{I1} > 0$, then $d'_1 = d_1 - \delta$. Without loss of generality, by continuity of the Step IA allocation, we may assume that PCM I is either weakly overbidding in both B_I, B'_I, or weakly underbidding in both.[6] We similarly assume that either paper 1 is weakly overdemanded in both bids ($d_1 > d'_1 \geq r$), or weakly underdemanded in both ($r \geq d_1 > d'_1$).

For paper 1, consider the case where $r \leq d'_1 < d_1$. Then $p_1 = \frac{r}{d_1}$, $p'_1 = \frac{r}{d'_1}$ and

$$f'_{I1}p'_1 \leq (f_{I1} - \delta)\frac{r}{d'_1} = r\frac{f_{I1} - \delta}{d_1 - \delta} < r\frac{f_{I1}}{d_1} = f_{I1}p_1,$$

where the inequality is since $\frac{a-x}{b-x} < \frac{a}{b}$ as long as $a < b$ and $x > 0$, and since $f_{I1} \leq r < d_1$ Next consider $r > d_1 > d'_1$, then $p'_1 = p_1 = 1$ and

$$f'_{I1}p'_1 = f_{I1} - \delta < f_{I1} = f_{I1}p_1.$$

In either case, we have

$$\overline{x}'_{I1} = f'_{I1}p'_1 < f_{I1}p_1 = \overline{x}_{I1}. \qquad (2)$$

For all $j \neq 1$, we have $\overline{x}'_{1j} = \overline{x}_{1j}$. This already shows that Step IA is monotone.

We are going to consider separately the case where the PCM is overbidding and the case where the PCM is underbidding. Note that since a PMA is full, MON2 entails MON1, and thus it is sufficient to prove MON2.

PCM is overbidding in both B_I, B'_I: i.e., $ob_I > ob'_I \geq 0$ holds. We have that for all $j \neq 1$,

$$x'_{Ij} = \overline{x}'_{Ij}\frac{k_I}{k_I + ob'_I} = \overline{x}_{Ij}\frac{k_I}{k_I + ob'_I} > \overline{x}_{Ij}\frac{k_I}{k_I + ob_I} = x_{Ij}. \qquad (3)$$

PCM is underbidding in both B_I, B'_I: Denote $\varepsilon := \overline{x}_{I1} - \overline{x}'_{I1}$. Note that

$$u_1 = r - p_1 d_1 = r - \min\{r, d_1\} = \max\{0, r - d_1\}.$$

More precisely, if paper 1 is overdemanded, then $u'_1 = 0 = u_1$. If paper 1 is underdemanded, then $\varepsilon = f_{I1} - f'_{I1} = \delta$ and $u'_1 = r - d'_1 = r - (d_1 - \delta) = u_1 + \delta = u_1 + \varepsilon$. In either case, $u'_1 \leq u_1 + \varepsilon$.

We argue that $\alpha' \geq \alpha$. Otherwise, assume by contradiction, that $\alpha' < \alpha$ holds. By the full property and by proportionality of the PMA,

$$k_I = \sum_{j\in M} x_{Ij} = \sum_{j\neq 1} \min\{q_{Ij}, \overline{x}_{Ij} + \alpha u_j\} + \min\{q_{I1}, \overline{x}_{I1} + \alpha u_1\},$$

and

$$k_I = \sum_{j\in M} x'_{Ij} = \sum_{j\neq 1} \min\{q_{Ij}, \overline{x}_{Ij} + \alpha' u_j\} + \min\{q_{I1}, \overline{x}_{I1} - \varepsilon + \alpha' u'_1\}$$
$$\leq \sum_{j\neq 1} \min\{q_{Ij}, \overline{x}_{Ij} + \alpha' u_j\} + \min\{q_{I1}, \overline{x}_{I1} - \varepsilon + \alpha'(u_1 + \varepsilon)\}$$
$$\leq \sum_{j\neq 1} \min\{q_{Ij}, \overline{x}_{Ij} + \alpha' u_j\} + \min\{q_{I1}, \overline{x}_{I1} + \alpha' u_1\}$$
$$= \sum_{j\in M} \min\{q_{Ij}, \overline{x}_{Ij} + \alpha' u_j\} < \sum_{j\in M} \min\{q_{Ij}, \overline{x}_{Ij} + \alpha u_j\} = k_I$$

which is a contradiction. The strict inequality in the last line is since there must be unconstrained papers for which α applies, or else $\alpha' = \alpha$ vacuously. Therefore, for all $j \neq 1$,

$$x'_{Ij} = \min\{q_{Ij}, \overline{x}_{Ij} + \alpha' u_j) \geq \min\{q_{Ij}, \overline{x}_{Ij} + \alpha u_j) = x_{Ij},$$

showing MON2 (and thus MON1).

□

Proposition 3.8. *The (unique) PMA is continuous in the bid.*

Proof. The prices and thus the initial assignment $\overline{X}_I$ change continuously with the demand. Therefore α must change continuously when I is overbidding, and also when I is underbidding as long as the set of constrained papers remains the same. We only need to verify that α does not change abruptly when a new paper becomes constrained.

This occurs exactly when $q_{Ij} = \overline{x}_{Ij} + \alpha u_j$ for some paper j, w.l.o.g. paper 1, i.e., $1 \in \tilde{Q}_I \setminus Q_I$. Suppose first that paper 1 is treated as unconstrained ($1 \notin \tilde{Q}_I$).Then since the PMA is full, α is such that

$$\sum_{j\in\tilde{Q}_I} q_j + \sum_{j\in M\setminus\tilde{Q}_I} (\overline{x}_{Ij} + \alpha u_j) = k_I.$$

If we treat paper 1 as unconstrained ($\tilde{Q}'_I := \tilde{Q}_I \setminus \{1\}$), then again to obtain a full assignment, α' is such that

$$k_I = \sum_{j\in\tilde{Q}'_I} q_j + \sum_{j\in M\setminus\tilde{Q}'_I} (\overline{x}_{Ij} + \alpha' u_j) = \sum_{j\in\tilde{Q}'_I} q_j + \sum_{j\in M\setminus\tilde{Q}_I} (\overline{x}_{Ij} + \alpha' u_j) + (\overline{x}_{I1} + \alpha' u_1).$$

[6] If both $ob_I > 0$ and $ub'_I > 0$ hold, we can break the move into two parts by choosing δ appropriately.

We therefore have that:

$$\begin{aligned}0 = k_I - k_I &= q_1 - (\overline{x}_{I1} + \alpha' u_1) + \sum_{j \in M \setminus \tilde{Q}_I} (\alpha u_j) - \sum_{j \in M \setminus \tilde{Q}_I} (\alpha' u_j) \\ &= q_1 - (\overline{x}_{I1} + \alpha' u_1) + (\alpha - \alpha') \sum_{j \in M \setminus \tilde{Q}_I} u_j \\ &= (\overline{x}_{I1} + \alpha u_1) - (\overline{x}_{I1} + \alpha' u_1) + (\alpha - \alpha') \sum_{j \in M \setminus \tilde{Q}_I} u_j \\ &= (\alpha - \alpha') \sum_{j \in (M \setminus \tilde{Q}_I) \cup \{1\}} u_j .\end{aligned}$$

Then either $\sum_{j \in (M \setminus \tilde{Q}_I) \cup \{1\}} u_j = 0$, in which case α is meaningless (no excess papers to allocate), or $\alpha' = \alpha$, as required. □

D EMPIRICAL RESULTS

In order to test the effect of price-based bidding, we simulated PCMs that interact with a bidding system. The PCMs observe dynamic paper prices and bid in turn. To keep simulations as realistic as possible, we used bidding data from real conferences to generate PCMs' costs and behaviors.

The hypothesis we want to verify is that price-based bidding results in a better bidding matrix (with bids being more balanced across papers, and in particular, less underdemanded papers), which in turn leads to a better assignment, both from the point of view of reviewers (reviewer costs are lower) and papers (fewer papers are assigned to some PCM who did not bid for them – which suggests that papers are better reviewed).

D.1 Setup

Datasets. We use all 5 bidding datasets available on PrefLib [14, 15] (marked DS1-DS5). In addition, we used 3 random samples in varying proportions from the AAAI'17 bidding dataset (marked DSA1-DSA3). In all datasets, we use $r = 3$. Every bid in the input has up to three levels, that can be interpreted as "strong bid", "weak bid", and "no bid".

Private Costs. From each instance (original bid matrix) in Table 2 we derived a cost matrix and a quota matrix as follows. For papers with COI we set $q_{ij} = 0$ and otherwise $q_{ij} = 1$.[7] For the other papers, we generated costs in the ranges $[0, 1]$ for strong bid, $[1, 2]$ for a weak bid, and $[2, 8]$ for a no bid, so a stronger bid in the input file always indicates a higher preference. As these ranges are somewhat arbitrary, we also present metrics that are independent of the numerical costs.

Bidding scheme and PCM behavior. Recall that R is the bidding requirement for each PCM. We use integral bids since this is more realistic. In the fixed bidding scheme, we considered the following PCM behaviors. All use integral bids.

original The PCM bids exactly as in the original PrefLib file.

uniform The PCM bids on the R papers with lowest cost.

In the price bidding scheme, we let any PCM bid exactly once, in random order. Since as long as there are few bidders the demands are too low to induce an informative price, we use a "virtual bootstrap bid": every PCM starts with a virtual bid of $\frac{k}{m}$ on each paper, which entails an initial demand of exactly r for each paper. This virtual bid is replaced by her real bid when she acts. We update the prices every 5 bids.[8]

We also use the greedy price-sensitive behavior analyzed in Section 4.2:

greedy The PCM bids on papers in increasing $(C_{ij} - \beta \cdot p_j)$ order, until their cumulative price reaches or exceeds R. Unless specified otherwise, we use $\beta = 2$.

All behaviors only decide whether to bid positively or not. COI declarations are identical to what they are in the original PrefLib file. When the PCM chooses to bid on a paper, the strength of the bid (needed for the assignment algorithm) is the same as in the original PrefLib file. For the Uniform and Greedy behaviors, we also varied the bidding requirement R.

We emphasize that the bid strength was used only for calculating the final assignments. The demands and prices during the bidding process considered every positive bid as 1.

Evaluation. To evaluate the benefit of the proportional price scheme, we simulated different PCM behaviors using the same private costs generated from the datasets above. We then used standard assignment algorithms to match the papers. Following Garg et al. [8] as well as Lian et al. [12] and implementations from Aziz et al. [2] we used the following algorithms which have been proposed in the literature for discrete allocation. The Utilitarian and Rank Maximal algorithm have been used in real conference assignment (it is most likely that the Utilitarian algorithm is used by EasyChair, one of the largest conference management website; see footnote in Lian et al. [12]). We highlight that our PMA plays no role in the empirical evaluation and was used only for analysis purposes.

Utilitarian. The Utilitarian assignment algorithm takes a set of bids and returns the assignment that maximizes the sum of bids (utilities) of the papers assigned to each PCM.

Egalitarian. The Egalitarian assignment algorithm takes a set of bids and returns the assignment that maximizes the sum of bids (utility) of the least well-off PCM.

Rank Maximal. The Rank Maximal assignment algorithm *ignores* the value of the bids of the PCMs and only considers the ordinal rankings of the PCMs. It returns an assignment such that for each PCM i, the assignment X_i maximizes the lowest ranked paper received. This mechanism was studied and highly advocated for by Garg et al. [8].

Our results show that findings are similar across different assignment algorithms, we thus present the results mainly for the Utilitarian algorithm. Since assignments are always in $\{0, 1\}$, we use $X_i \subseteq M$ to denote the set of papers assigned to i. Recall also that $B_i = \{j \in M : f_{ij} > 0\}$.

Having the assignment $\boldsymbol{X}$, we measure how good it is using three metrics:

Social Cost The average cost to a PCM. Formally, $\frac{1}{n} \sum_{i \in N} \sum_{j \in M} C_{ij} x_{ij}$;

[7] We did not have the COI of AAAI'17.

[8] Varying the number of rounds between price updates had almost no effect on the results.

name	m	n	bids / PCM	strong bids/ PCM	$k = \frac{mr}{n}$	$< r$ bids	0 bids	$\geq$ 10 bids
DS1	176	146	8.9	5.6	3.6	28	6	49
DS2	52	24	14.3	8.5	6.1	4	0	14
DS3	54	31	10.4	5.2	5.2	18	3	15
DS4	442	161	17.6	5.1	8.2	62	8	91
DS5	613	201	21.1	6.3	9	125	30	155
DSA1	600	400	6.6	4	4.5	213	51	47
DSA2	1200	300	12.3	7.8	12	621	147	36
DSA3	2000	200	22.7	13.7	30	1264	362	16

Table 2: Anonymized bidding datasets used in our simulations.

Fraction of fulfilled bids The average (across PCMs) of $\frac{|X_i \cap B_i|}{|B_i|}$. A low number indicates the PCM failed to obtain many of her requested papers.

Fraction of assigned papers w/o bid The average (across PCMs) of $\frac{|X_i \setminus B_i|}{|X_i|}$. A high number indicates that the PCM was assigned many papers she did not bid for.

The first metric is the simplest but is sensitive to the way we generated to private costs. The latter two metrics (that are complementary to each other) can be thought of as recall and false discovery rate (complement of precision), respectively. The third one is particularly important, since papers with missing bids are assigned essentially at random, and therefore substantially decrease the review quality and will perhaps require manual re-assignment in practice. We repeated the process 5-20 times and calculated the average of each evaluation metric.

D.2 Results

Social cost. Table 3 compares the social cost obtained under the Original bidding behavior, with the social cost obtained under the proportional price scheme with the greedy behavior. It can be clearly seen that the latter substantially reduces the social cost, especially for the large datasets. In fact the only dataset where this is not true is the DS2, which is both tiny and has no problem of orphan papers in the first place (see Table 2).

We next turn to study what contributes to the improvement in the assignment.

Recall-precision tradeoff. A major factor that affects the assignment is the total number of bids. We therefore compared the Greedy behavior to the Uniform behavior, while varying the bidding requirement R in both. On Figure 2 we can see that, as expected, increasing the bidding requirement results in a lower fraction of assigned papers without bid (higher precision), and a lower fraction of fulfilled bids (lower recall).

More importantly, the graphs highlight the benefits of the proportional price scheme. First, the recall and precision under the greedy behavior are substantially better than under the uniform behavior. Second, this is obtained with much fewer bids and thus presumably less effort on behalf of the PCM. Finally, in the current (no price) bidding scheme, it is difficult to decide how to set the requirement R, or predict how PCMs will conform, whether in the price scheme, setting $R = k$ is useful as a reasonable anchor.

Sensitivity to behavior. Since we cannot control the behavior of the PCMs, it is important to check that the results are not too sensitive to small changes in the behavior. We therefore varied the parameter β in the greedy behavior (recall that higher β means the PCM will lean more towards ignoring low-price papers), as well as the fraction of PCMs who comply with the price-scheme bidding instructions (non-compliant PCMs follow the Original behavior).

We can see that the quality of assignment gradually improves as the sensitivity to prices (Figure 3, left) increase. This pattern means, perhaps counter-intuitively, that strictly following a sincere behavior (where $\beta = 0$) is not ideal from a social perspective, as it requires to bid on medium-ranked-but-highly-demanded papers that are likely to fit better some other PCM.

As for the compliance rate, as more PCMs switch to greedy bids the assignment improves in general (Figure 3, middle and right), but we can also see in the right figure that this only slightly hurts PCMs who choose to ignore the prices altogether, as can be seen from the very shallow increase in the cost for original PCMs.

E ADDITIONAL EMPIRICAL RESULTS

In this appendix we further show that our empirical findings are consistent across datasets and assignment algorithms.

Figures 4 and 5 show the recall and precision under both schemes for different bidding requirements. Results are similar to what we see in the main text.

Dataset	DS1	DS2	DS3	DS4	DS5	DSA1	DSA2	DSA3
Utilitarian algorithm								
Original	4.9	6.5	8.8	14.8	15.5	9	27	81.3
Greedy	4.6±0.11	7.5±0.22	8.2±0.16	11±0.1	11.8±0.09	5.6±0.04	15.4±0.09	42.9±0.09
Egalitarian algorithm								
Original	5.3	8.2	9.4	15.1	17	9.3	51.1	134.2
Greedy	4.7±0.15	8.3±0.35	8.8±0.25	11.7±0.62	12±0.21	6.5±0.9	29.7±0.31	73.2±0.64
Rank-Maximal algorithm								
Original	5	6.5	8.7	15	15.4	9	26.7	79.6
Greedy	4.6±0.08	7.5±0.21	8.1±0.16	11.1±0.12	11.7±0.05	5.7±0.27	15.3±0.16	42.8±0.11

Table 3: A comparison of the social cost under original bids and the proportional price scheme, with all three assignment algorithms. We compare the Original bids and the Proportional price scheme (with the Greedy behavior and $R = k$). We add a confidence interval of 2 standard deviations due to the random bidding order of the greedy behavior.

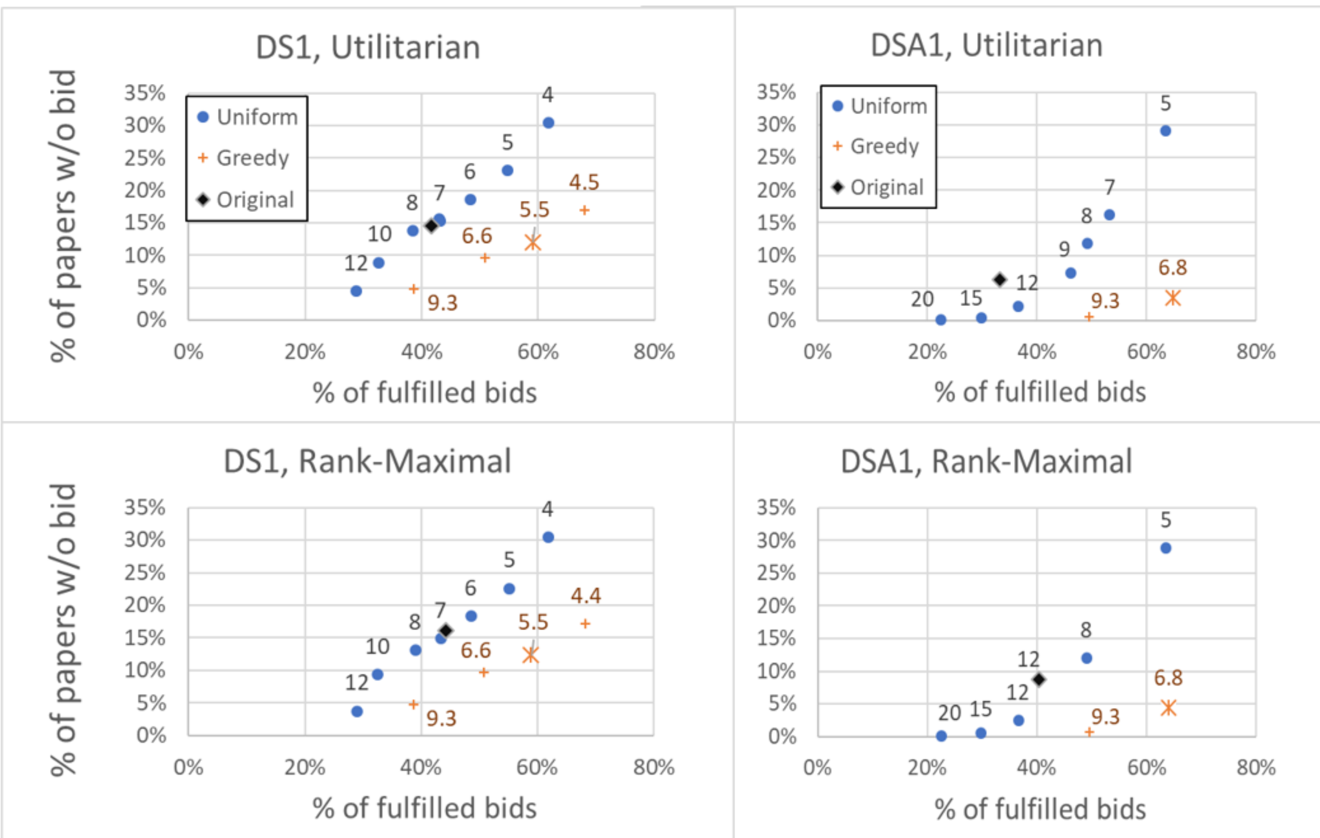

Figure 2: Assignment vs. bids. The top figures show the quality of the assignment obtained via the Utilitarian algorithm, under the three behaviors we consider. The number above each bullet is the average number of positive bids per PCM. The highlighted 'Greedy' bullet marks the outcome for $R = k$. The bottom figures show the assignments with the Rank-Maximal algorithm.

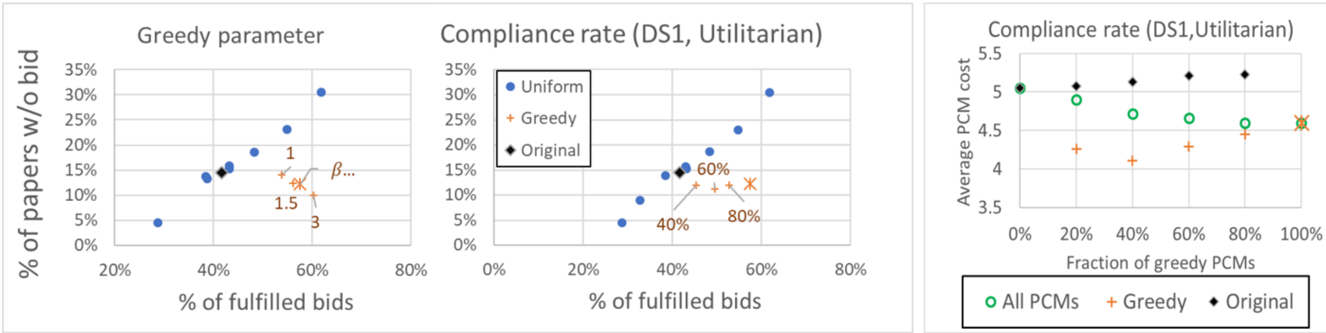

Figure 3: On the left, the effect of varying the parameter β in the greedy behavior (see labels on bullets). On the two right figures, we see the effect of increasing the compliance rate. The bullet labels in the middle figure (and the X-axis in the right figure) show the fraction of PCMs who follow the greedy behavior.

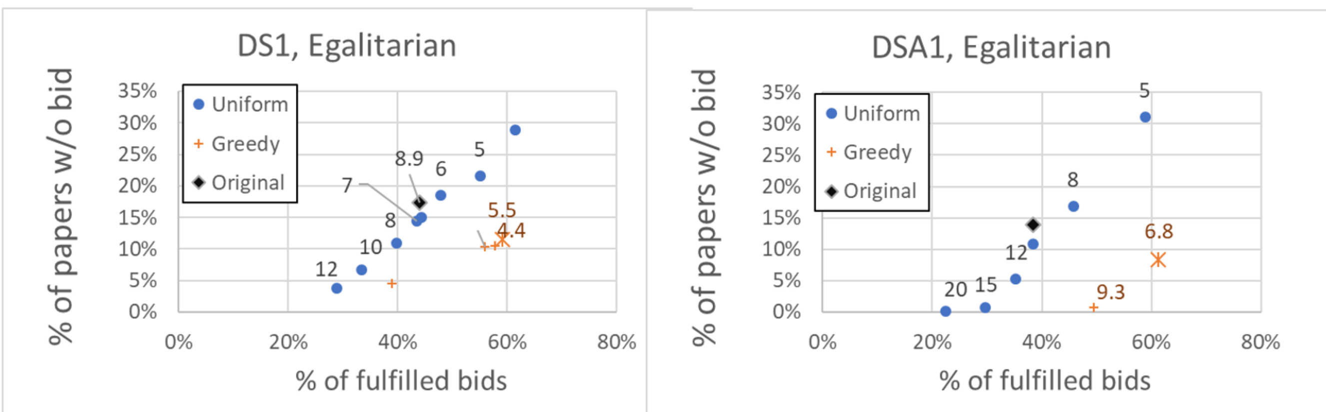

Figure 4: The tradeoff curves of the Egalitarian assignment for DS1 and DSA1. The number above each bullet is the average number of bids per PCM.

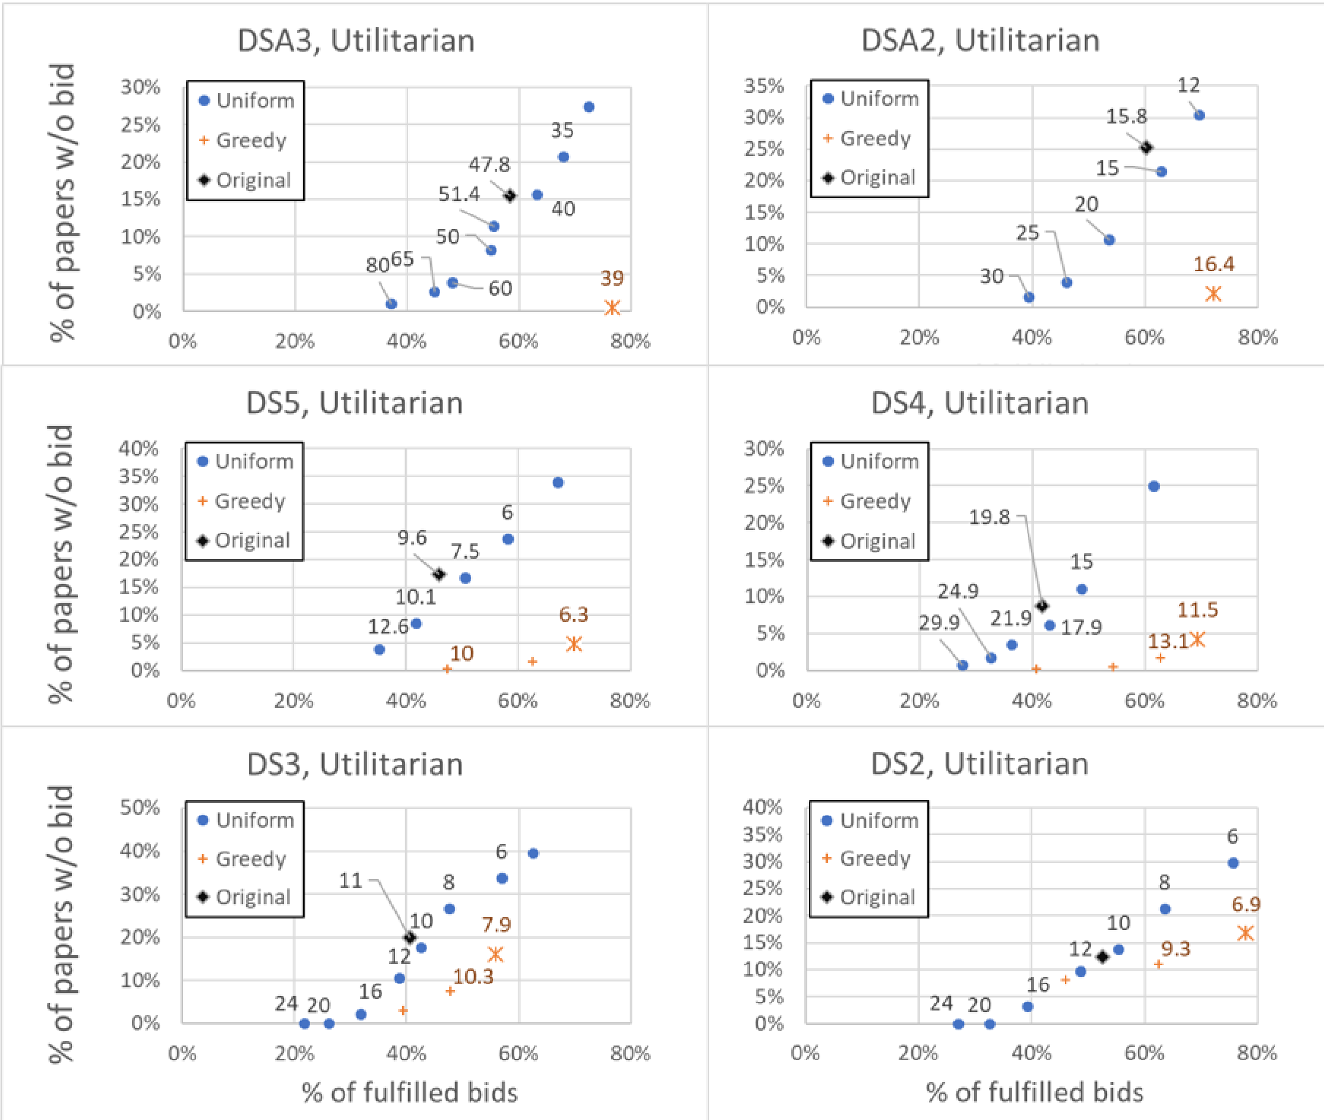

Figure 5: The tradeoff curves of the Utilitarian assignment for all remaining datasets (DS1 and DSA1 are shown in the main text). The number above each bullet is the average number of bids per PCM.