A Market-Inspired Bidding Scheme for Peer Review Paper Assignment

Reshef Meir,[1] Jérôme Lang,[2] Julien Lesca,[2] Nicholas Mattei,[3] Natan Kaminsky[1]
[1]Technion—Israel Institute of Technology
[2]Université Paris Dauphine
[3] Tulane University

Abstract

We propose a market-inspired bidding scheme for the assignment of paper reviews in large academic conferences. We provide an analysis of the incentives of reviewers during the bidding phase, when reviewers have both private costs and some information about the demand for each paper; and their goal is 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) 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. This game-theoretic analysis is based on a simple, prototypical assignment algorithm. We show via extensive simulations on bidding data from real conferences, that our bidding scheme would substantially improve both the bid distribution and the resulting assignment.

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 fields.

While in journals the assignment of papers to reviewers is typically handled manually by the editors, peer-reviewed conferences often use automated assignment algorithms 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. Though the frequency of these 'orphan' papers varies, they are a recurrent problem in the bidding processes. Fiez, Shah, and Ratliff (2019) discuss this issue and highlight the importance of studying and improving the bidding process.

Our ultimate goal is *to improve the assignment, to the benefit of all involved parties: reviewers, program chairs, and authors.* We argue that only designing new assignment algorithms will not help achieve this goal (or help marginally): what we need before all is to incentivize reviewers to bid in a more cooperative way. Thus we do not propose a new assignment algorithm but a new *bidding scheme* that improves over current bidding schemes. Crucially, the improvements we suggest are easy to implement: they only require revealing 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.

Evaluating our bidding scheme comes with a difficulty: each conference comes with its own assignment algorithm, and we do not want a separate analysis for each of these. However, we argue that what counts is not the precise assignment algorithm used, but a simple abstraction of it, representing what a reviewer believes about the possible assignments, based on what she can observe. For this we define a simple algorithm (*proportional mock assignment*), based on an adaptation of the *trading post mechanism* (Shapley and Shubik 1977), which we use to analyze bidders incentives.

We next 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 real-world conferences.

The Paper Assignment Process

A large CS conference such as AAAI, IJCAI, ICML, or NeurIPS has some 1,000 – 10,000 submissions and 1,000 – 3,000 reviewers. The assignment process typically proceeds through the following steps:

- The program chair recruits n program committee members (PCMs) to serve as reviewers.
- 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 interest (COI). Typically PCMs are asked (but not enforced) to bid positively on a minimum number R of papers, R being much higher than the actual number of expected reviews (for senior PCMs at AAAI-21, $R = 75$).
- The final bids are fed as input to an *assignment algorithm*, together with additional constraints such as a minimum

Conference	m	k	$< r$ bids	0 bids	$\overline{\text{bid}}$ / paper
IJCAI 18	3470	~ 5	140	5	29.7
AAAI 17	2414	~ 5.5	47	18	13.66
KSEM 06	235	10.7	18	5	6
KR 08	234	7.5	-	-	9.7
TARK 17	91	13.5	32	8	3.4

Table 1: Recent conferences' statistics.

and maximum number of papers per PCM. Each paper is assigned to a fixed number r of reviewers, typically 2 – 4.

On average, each PCM should get $k := \frac{mr}{n}$ papers for review. A conference comes with an assignment algorithm: either an off-the-shelf algorithm (there are many in use) or an algorithm written specifically for the conference (Garg et al. 2010).[1] In other cases, the actual assignment algorithm can be a black-box; see discussion by Lian et al. (2018).

Unfortunately, it is common that 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, see Table 1 .[2] For instance, at AAAI-17 each paper got > 13 bids on average, yet 18 papers got no bid at all.

Drawbacks of commonly used bidding-assignment protocols include: (1) program chairs spend a lot of time and effort coping with underdemanded papers; (2) the assignment of these papers is somewhat arbitrary, and results in additional work for PCMs and lower quality reviews; (3) PCMs spend time and effort bidding on papers they are unlikely to get, either because they are overdemanded, or because they already unknowingly bid on enough papers.

Paper Goal and Contribution

Our goal is to improve paper bidding, which in turn will improve paper assignment. We present a simple protocol called the *trading post bidding scheme*: PCMs are each assigned some initial budget R they are expected to exhaust, and the "price" of every paper changes dynamically, 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 trading post 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 define a "mock assignment" based on a modification of the *trading post mechanism* (Shapley and Shubik 1977). This is an *abstraction* that captures the key properties of assignment algorithms and the *perceived probability*, for a PCM, to be assigned a paper.
- We prove that in the game induced by the trading post bidding scheme and the mock assignment, there is an incentive for PCMs to make sincere bids that match the requested bidding amount R.
- We show via extensive simulations that if bidders are sensitive to price (less likely to bid on cheap, low probability papers) then the trading post bidding scheme obtains socially better outcomes *after a single bidding round*: bids are more balanced across papers, and PCMs get more desired papers.

A secondary claim that we make is that the trading post bidding scheme 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 proofs and simulation results to the appendix of the full version at https://tinyurl.com/y8z4y3z4.

Related Work

Algorithms. A first stream of papers consists of *information retrieval approaches* to the assignment problem.

From these, we only mention a recommendation-based approach that aims at easing paper bidding, namely the *Toronto Paper Matching System* (TPMS) (Charlin, Zemel, and Boutilier 2011): 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 trading post bidding scheme is compatible with TPMS or similar systems as well as advances in assignment and matching algorithms, and the PCM can sort/filter papers according to their fit, their price, or some combination thereof.

A second stream deals with *assignment algorithms*. Some of these are based on bids (Lian et al. 2018; Garg et al. 2010; Goldsmith and Sloan 2007); they are highly relevant to us, and we go back to them when describing our evaluation.

Bidding Behavior Very few papers focus on the behavior of bidders within the assignment system. The only empirical work we know of is (Rodriguez, Bollen, and de Sompel 2007), 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, and more. They find that the bidding behaviour is only weakly related to the subject of the submission, and is influenced by plenty of other (unidentified, conjectured) factors. They repeat the observation on unbalanced 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.*

Closer to our work, Fiez et al. (2019) assume a position-based click model from the literature, and determine an optimal order to present papers to PCMs based on their preferences and paper demands. We see this approach as complementary to ours, as PCMs are likely to be affected by both order of presentation (as in that paper) and explicit information on the demand (as in our paper). Ultimately, both interventions have similar effects, as they lead to more balanced bids and thus a better assignment.

[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.

Game Theory, Mechanism Design and Social Choice Some papers focus on strategyproof peer review (Kurokawa et al. 2015; Xu et al. 2018; Aziz et al. 2019c). Especially, Xu et al. (2018) consider scenarios where reviewers (who are also authors) may bid strategically so as to influence the assignment. Here we assume that reviewers bid independently from their interest as authors, but are still self-interested and would like to minimize their effort during the bidding and reviewing phases.

Paper assignment can also be seen as a *chore assignment problem*, which is known to be very different from the assignment of goods, because chores are non-disposable. The main two mechanisms used are the egalitarian mechanism (leximin) and the competitive mechanism based on prices that derive from the input of the problem (Bogomolnaia et al. 2019; Aziz et al. 2019a; Brânzei and Sandomirskiy 2019; Freeman et al. 2020; Chaudhury et al. 2020).

Market approaches (*e.g.*, using scrip money) have been considered for various allocation problems such as course allocation (Budish 2011; Babaioff, Nisan, and Talgam-Cohen 2017; Othman, Sandholm, and Budish 2010).

All these mechanisms do not readily adapt to the constraints of paper assignment, nor do they easily combine with existing bidding processes, and indeed we are unaware of a theoretical or practical application to paper assignment.

We conclude, similarly to Fiez, Shah, and Ratliff (2019), 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*.

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 *fractional partial assignment* is a matrix $\boldsymbol{X} = (x_{ij})_{i\in N, j\in M}$ — where $x_{ij} \in [0,1]$ can be seen as the probability that j is assigned to i — and subject to constraints:

Capacity Constraints $\forall i \in N, \sum_{j\in M} x_{ij} \leq k_i$ (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.

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

We say that X_i is *valid* if it satisfies capacity and quota constraints; X_i is *full* if $\sum_{j\in M} x_{ij} = k_i$. We say that $\boldsymbol{X}$ is *valid* [resp., *full*] if X_i is valid [full] for all $i \in N$. An assignment $\boldsymbol{X}$ is *integral* if $x_{ij} \in \mathbb{Z}$ for all i, j. 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, 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 on paper j. We always require $f_{ij} \leq q_{ij} \leq 1$, 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. 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.

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* takes as input a bidding profile $\boldsymbol{B}$, and outputs a valid partial fractional assignment $\boldsymbol{X}$.

The Trading Post Bidding Scheme

As explained in the introduction, we want to present PCMs with information on the demand in the form of prices. Our starting point is the *trading post* (TP) *mechanism* (Shapley and Shubik 1977), which assigns to each item j a "price" $P_j := \frac{d_j}{r}$. It then allocates items based on prices, but we will get back to this later on.

Our suggestion is simple: PCMs should be presented with the "inverse price" $p_j := \min\{1, \frac{r}{d_j}\}$ (henceforth, the "iprice") that is being updated dynamically as bids change.[3] Let $tb_i := \sum_{j\in M} f_{ij} p_j$ be the *total bid* of i. Given a bidding requirement R, the instructions to the PCM are to bid until $tb_i \cong R$, rather than bidding on R papers. Just like today, this requirement is not enforced, and PCMs may, but are not required to, login later on and revise their bid. Unlike current systems, they may see different iprices each time (see the Discussion section). To avoid changing the presented iprices as the PCM bids, the demand d_j during i's bidding is always computed as if $f_{ij} = 1$ (see Example 1, Item 2).

As we will show, the TP bidding scheme will facilitate both the formal analysis and the behavior of the PCMs.

Example 1. *We have $m = 4$ papers (a, b, c, d) and $n = 6$ reviewers; $r = 2$, therefore 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 $\mathbf{d} = (2, 1, 1, 0)$*
- *Reviewer 3 logs in. She sees the iprices $p_a = \min\{1, \frac{r}{|D_a \cup \{3\}|}\} = \frac{2}{3}, p_b = \min\{1, \frac{r}{|D_b \cup \{3\}|}\} = \min\{1, \frac{2}{2}\} = 1, p_c = p_d = 1$ (since the PCM is counted in the demand). Suppose she bids on $\{a, b, d\}$.*
- *Now Reviewer 4 logs in and sees iprices 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 iprices of $p_a = \frac{2}{3}, p_b = \frac{1}{2}, p_c = 1, p_d = \frac{2}{3}$.*

[3] We call it "inverse price" because it is not a price that the PCM pays for reviewing a paper but rather a 'payment' that she receives, since papers are chores and not goods.

Proportional Mock Assignments

In order to analyze 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 either unknown, or 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, computing the assignment requires the bidder to know the full bid matrix, and this is not available to her.

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. In the remainder of this section and the next one, we consider the assignment from the point of view of a particular PCM I. While the exact incentive model is described below, the general idea follows expected cost minimization: our assumption will be that every allocated paper in X_I has a certain 'cost', and I bids in a way that minimizes the total cost. Since the assignment algorithm maps any possible bid of I to an allocation X_I, the expected cost of every bid is well defined.

Mock TP Assignment As a first attempt, we can consider the (modified) trading post assignment (TPA) itself. That is, define the assignment to be $\overline{x}_{Ij} := f_{Ij}p_j$ for all j.

We make a distinction between two cases: if $tb_I > k_I$ we say that I *overbids*; and if $tb_I < k_I$ we say that she *underbids*. If $tb_I = k_I$ 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:

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

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

The TPA guarantees that each overdemanded paper is fully allocated, as $\sum_{i\in N} f_{ij}p_j = p_j \sum_{i\in N} f_{ij} = \frac{r}{d_j} d_j = r$, but puts no constraint on the number of assigned papers to each PCM, which is inconsistent with actual assignments. We would therefore like to extend the TPA $\overline{X}_I = (\overline{x}_{Ij})_{j\in M}$ to a full and valid assignment in a reasonable way.

Definition 1 (PMA). *A valid and full assignment $X_I = (x_{Ij})_{j\in M}$ is a* proportional mock assignment for PCM I *w.r.t. input $B_I, \boldsymbol{q}_I = (q_{Ij})_{j\in M}, \boldsymbol{d} = (d_j)_{j\in M}$, if there is a constant $\alpha \leq 1$ such that:*

(OB) *If I is weakly overbidding, then $x_{Ij} = \alpha \cdot f_{Ij}p_j$ for all $j \in M$;*

(UB) *If I is weakly underbidding, then $x_{Ij} = \min\{q_{Ij}, f_{Ij}p_j + \alpha \cdot u_j\}$ for all $j \in M$.*

A few explanations are in order. If I is overbidding, then the assignment is simply proportional to bids weighted by paper iprices. If I is underbidding, such as assignment 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 iprice, 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.

The UB step can also be thought of as a "water filling" process where each paper j is a tube of width u_j (the rate at which it is being filled) and height $q_{Ij} - \overline{x}_{Ij}$ (the cap).

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 2 (Extendable assignment). *A partial assignment $\overline{X}_I$ 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.*

Extendability means that there is *some way* to partially allocate leftover papers, so that I obtains a full and valid assignment. A TPA is almost always extendable: A sufficient condition is that either I overbids, or $k_I \leq \sum_{j\in M} \min\{q_{Ij}, u_j\}$. Thus, **for the rest of the paper we will assume that extendability always holds** (for PCM I).[4]

Proposition 1. *Given bid B_I, quotas $\boldsymbol{q}_I$ and demands $\boldsymbol{d}$, there exists a PMA if and only if the TPA $\overline{X}_I$ is extendable. Moreover, the PMA is unique.*

The proof is constructive and relies on a simple algorithm that roughly proceeds as follows: it starts from the TPA $\overline{X}_I$; then, if I is overbidding α is a simple normalization factor; if I is underbidding, α is computed at most m iterations.

We demonstrate here how the algorithm works on a simple example. The full details of the algorithm and the example are in the full version.

Example 2 (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} and assignments are shown only for $i \in \{1, 2\}$, together with the totals:*

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

This induces the following TP assignment $\overline{X}$:

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

PCM 1 is overbidding, thus $\alpha_1 = \frac{r}{tb_1} = \frac{60}{77}$. PCM 2 is underbidding. If we independently compute the PMA for each

[4] We do note that the extendability for each PCM separately does not mean that we can extend the partial allocation into a full allocation for all PCMs. This will not matter for our purpose.

agent, we get:

x_{ij}	a	b	c	d	e	f	α_i
1	$^{44}/_{77}$	$^{30}/_{77}$	$^{24}/_{77}$	$^{60}/_{77}$	0	0	$^{60}/_{77}$
2	0	$^{1}/_{2}$	$^{2}/_{5}$	$^{2}/_{5}$	1	$^{7}/_{10}$	$^{2}/_{5}$

Properties of PMA Two crucial properties of the PMA are *monotonicity* and *continuity*.

Definition 3 (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, by decreasing her bid on paper j PCM I gets no more of paper j and no less of all other papers.

Proposition 2. *The PMA is monotone in bids.*

The proof outline is simple: we consider a sufficiently small change in the bid such that no paper goes from underbid to overbid or vice versa. Then, for papers with overbid monotonicity is immediate as increasing the bid means higher allocation both before and after normalization. For papers with underbid this is more tricky since a higher initial assignment also means less leftovers to assign at the second step. However due to proportionality we can show that the net change is positive.

Proposition 3. *The (unique) PMA is continuous in the bid.*

Here the main obstacle is to show that the allocation does not change abruptly when one of the constraints becomes binding.

Incentives

The bidding process can be thought of as a game, where every time a PCM logs in, she sees the current iprices (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.

In order to define this utility, we assign a cost C_{Ij} reflecting the inconvenience of reviewing paper j to PCM I: $c_I(B_I, \boldsymbol{d}^{-I}) := \sum_{j\in M} x_{Ij} C_{Ij}$, where $X_I = (x_{Ij})_{j\in M}$ is the unique PMA corresponding to bid B_I. We assume that costs are *generic*, in the sense that no two sets of papers have exactly the same cost.

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.*, $X_I = X'_I$.

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 iprices $(p_j)_{j\in M}$,[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. We denote the sincere exact bid of I (w.r.t. $\boldsymbol{d}^{-I}, \boldsymbol{q}_I, k_I$) by b^*_I. Our main theoretical result is the following.

Theorem 4. *Any best response is equivalent to b^*_I.*

Proof sketch. We first show that every best response is sincere: 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, the "insincerity" of B_I is $IN(B_I) := j'' - j' - f_{Ij'} + f_{Ij''}$. We show that if $IN(B_I) > 0$ we can construct a bid that is just as good and is strictly more sincere: the easy case is when I is weakly overbidding, as slightly increasing the bid on j' will improve both cost and sincerity. If I is underbidding, we need to both increase the bid on j' and *decrease* the bid on j'' in a subtle way so as to strictly decrease the cost. Then, we take the sincere best response $b_I \in \mathbb{R}^+$. If b_I is an overbid, we slightly decrease it and get rid of the least favorable papers. If b_I is an underbid, we slightly increase the demand on the last paper, and show that this does not affect the more favorable papers. □

Observation 5. *If B_I is exact, then $x_{Ij} = f_{Ij} p_j$. In particular, if I places a full bid on j ($f_{Ij} = 1$) then $x_{Ij} = p_j$: the iprice is exactly the probability that I gets j if she bids on it.*

This means that by meeting the (voluntary) bidding requirement R, the PCM is both being rational (by Thm. 4), and has a decent idea on which papers they will get.

Price-Sensitive Bids

Intuitively, a PCM facing a bidding requirement will tend to skip papers with a low iprice, as they don't help reaching the requirement and are also less likely to be obtained (Obs. 5). This is especially true if learning one own's private cost requires some effort (like reading the abstract). We thus consider two bidding behaviors that capture this effect.

Sincere Bidding + Costly Exploration Preferring high-iprice 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}. 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. Note that now the strategy of the PCM is more complex: first decide (perhaps sequentially) on which papers to explore, and then decide

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

what bid to submit. The last step is already solved in the previous section, as given her current state of information the PCM is best off submitting her sincere exact bid b^*_I.

We argue that if exploration is *myopic*, i.e. the PCM always explores as if this is the last possible exploration, then the utility from exploration is weakly increasing in the iprice p_j. The formal details are in the full version. Indirectly, this also means that ceteris paribus, the PCM is more likely to bid on papers with a higher iprice.

Greedy Price-Sensitive Bidding The "attractiveness" of each paper is some function that is increasing (linearly) in iprice and decreasing (linearly) in private cost. The bidder selects papers greedily by decreasing attractiveness, until reaching the bidding quota. This 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 (Bettman, Johnson, and Payne 1991).

In-Silico Experiments

In order to test the effect of the TP bidding scheme, we simulated PCMs who interact with a bidding system. The PCMs observe dynamic paper iprices 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 the trading post bidding results in a better bidding matrix, which in turn leads to a better assignment, from both points of view of the reviewers and of the system (*i.e.*, the chair).

Datasets We used all 5 bidding datasets (DP1-DP5) available on PrefLib (Mattei and Walsh 2013, 2017). In addition, we used random samples in varying proportions from another large AI conference (DA1-DA3). In all datasets, we use $r = 3$. Every bid in the input has up to three levels, interpretable as "strong bid", "weak bid", and "no bid".

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

We also used two datasets (DI1, DI2) sampled from the ICLR'18 dataset used in Fiez, Shah, and Ratliff (2019). This dataset has a score $u_{ij} \in [0, 1]$ for each (i, j). The dataset does not contain the true bids.

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

original The PCM bids exactly as in the original data.

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

In the trading post bidding scheme, we let a PCM bid exactly once, in random order. Since as long as there are few bidders

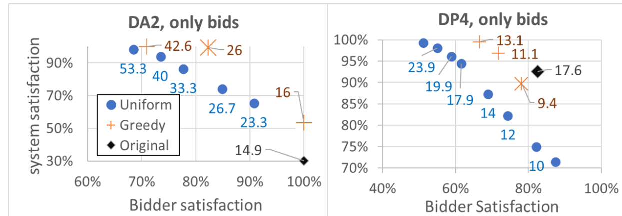

Figure 1: Tradeoff between system satisfaction (Y-axis) and bidder satisfaction (X-axis). In each figure, we show satisfaction under the three behaviors we consider, as we vary the bidding requirement, if available. The number above each bullet is the average number of positive bids per PCM. The highlighted 'Greedy' bullet marks the outcome for $R = k$.

the demands are too low to induce an informative iprice, 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 iprices every 5 bids.

We also use the greedy price-sensitive behavior:

greedy The PCM bids on papers in increasing order of $(C_{ij} - \beta \cdot p_j)$, until their cumulative iprice reaches or exceeds R. Unless specified otherwise, we use $\beta = 2$.[6]

All behaviors only decide whether to bid positively or not. COI declarations are as in the original data. 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 data. For the Uniform and Greedy behaviors, we also varied R.

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

Evaluation To evaluate the benefit of the trading post bidding scheme, we simulated different PCM behaviors using the same private costs generated from the datasets above. Since the greedy scheme has a random component (bidding order), we run it multiple times and take the average metrics.

We first evaluate paper-side gain and reviewer-side gain of the obtained bidding matrix, following Fiez, Shah, and Ratliff (2019). Their advantage is that they are independent of any particular assignment.

System satisfaction $\textsc{SysS}(\boldsymbol{B}) := \frac{1}{m} \sum_{j \in M} \frac{\min\{r, d_j\}}{r}$. Thus a maximal gain of 1 indicates all papers have sufficient bids, whereas a low score on this metric indicates there are many 'orphan papers' with insufficient bids. $\textsc{SysS}(\boldsymbol{B})$ is a rough upper bound on the ability of any assignment algorithm to satisfy all papers.

Bidder satisfaction Let $OPT_i := \frac{1}{R} \min\{C_i(S) : |S| = R\}$ be the best average paper cost that i could hope for. We set $\textsc{BidS}(\boldsymbol{B}) := \frac{1}{n} \sum_{i \in N} \frac{OPT_i}{\frac{1}{|B_i|} \sum_{j \in B_i} C_{ij}}$. Thus a low value indicates that reviewers get a far-from-optimal bunch of

[6] In the ICLR data we first transformed the cost to the right range by $C'_{ij} := 6 \cdot (1 - \frac{u_{ij}}{\max_{j'} u_{ij'}})^2$.

papers. Note that by definition, the original bids maximize BIDS if the total number of bids is fixed.

On Figure 1 we plot SYSS vs. BIDS and give the average number of positive bids per PCM. We neatly see the SYSS-BIDS tradeoff. We also see that the greedy scheme completely dominates the uniform bidding scheme. Moreover, it leads to the same tradeoff as the original bids, but with much fewer positive bids.

Evaluating Assignments Ultimately, we care about good assignments rather than good bids. Following Garg et al. (2010) and Lian et al. (2018), and implementations from Aziz et al. (2019b), 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 conferences. (It is most likely that the Utilitarian algorithm is used by EasyChair; see footnote in Lian et al. (2018).) Note that our PMA plays no role in the empirical evaluation and is used only for incentive analysis.

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 cardinal value of the bids. It returns an assignment minimizing the maximum, over all PCMs, of the rank of the lowest ranked paper received. It is studied and highly advocated for by Garg et al. (2010).

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 that $B_i = \{j \in M : f_{ij} > 0\}$.

Satisfaction Measures for Assignments After the assignment takes place, we return to measure the system and bidder satisfaction. $\text{SYSS}(\boldsymbol{X}, \boldsymbol{B}) := \frac{1}{rm} \sum_{i \in N} |X_i \cap B_i|$ measures the fraction of papers that were assigned to PCMs that bid on them (as all other papers may require intervention and possibly reassignment); and $\text{BIDS}(X_i, \boldsymbol{B}_i) := |X_i \cap B_i| / |B_i|$ measures, for each PCM, the fraction of realized bids. Note that both measures are based only on the bids and not on the private costs.

We also measure the social cost, which is independent of the bids: $\text{SC}(\boldsymbol{X}) := \frac{1}{n} \sum_{i \in N} \sum_{j \in M} C_{ij} x_{ij}$.

Results

Table 2 compares the social cost obtained under the Original bidding behavior, with the social cost obtained under the trading post bidding scheme with the greedy behavior. It can be clearly seen that the latter substantially reduces the social cost, especially for the large datasets. In the ICLR datasets where original bids are not available, we considered the minimum social cost of the uniform bid over all R we tried.

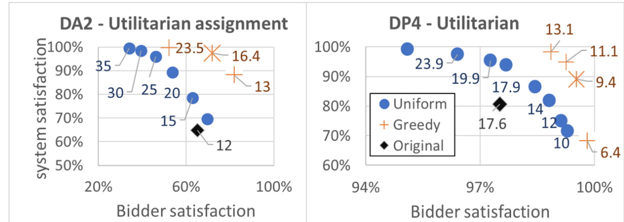

Figure 2: The Utilitarian assignment for datasets DP4 and DA2 under all three behaviors.

Recall-precision tradeoff Just like in the bid matrix, the tradeoff as we increase R clearly reflects in the assignment. Figure 2 shows the Utilitarian assignments. Increasing the bidding requirement results in a lower fraction of assigned papers without bid (higher precision, higher SYSS), and a lower fraction of fulfilled bids (lower recall, lower BIDS).

More importantly, the graphs highlight the benefits of the trading post bidding scheme. First, both BIDS and SYSS under the greedy behavior are substantially better than under the uniform or original behavior. Second, this is obtained with much fewer bids and thus presumably less effort on behalf of the PCM. Finally, in the current (no iprice) bidding scheme, the bidding requirement R is usually set to some arbitrary high number. In the new iprice scheme, we can recommend setting $R = k$ to guarantee sufficient bids, or a bit above for a safety margin.

In the full version we show that all of our datasets exhibit a similar pattern, under all three assignment algorithms. We also show that the satisfaction measures gracefully degrade if PCMs are less sensitive to the presented iprices, or if some PCMs ignore them altogether and stick to their original bids.

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 the PCM is likely to bid on moderately-desired papers that fit better many other PCMs. As for the compliance rate, as more PCMs switch to greedy bids the assignment improves at a linear rate, see Fig. 3 right.

Discussion

In this paper we proposed the *trading post bidding scheme*: reviewers observe dynamic "prices" that reflects papers' demands and update throughout the bidding process.

We showed via the introduction of a stylized assignment that changes linearly with the bids (the PMA), that bidders

Dataset	DP1	DP2	DP3	DP4	DP5	DA1	DA2	DA3	DI1	DI2
$m \times n$	176×146	52×24	54×31	442×161	613×201	600×400	1200×300	2000×200	600×400	900×300
Original	5.05	6.5	8.8	14.8	15.5	9	27	81.3	0.716	0.583
Greedy	4.6±0.04	7.5±0.22	8.2±0.16	11.2±0.06	11.8±0.09	5.5±0.05	15.6±0.06	42.9±0.02	0.684±0.001	0.554±0.001

Table 2: A comparison of the social cost under original bids and the trading post bidding scheme, under the Utilitarian algorithm. We compare the Original bids (For ICLR data, the minimum over all uniform bids) and the trading post bidding scheme (with the Greedy behavior and $R = k$). We add a confidence interval of 2 standard deviations due to random bidding order.

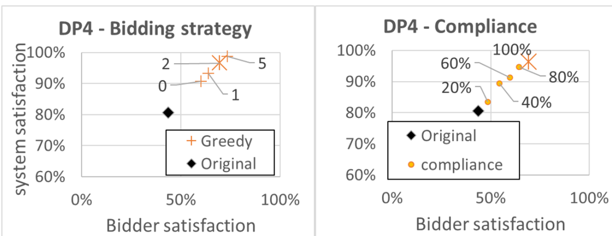

Figure 3: On the left, the effect of varying the parameter β in the greedy behavior (see labels on bullets). On the two figure, we see the effect of decreasing the compliance rate from 100% (rightmost bullet) to 0%.

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, as can be seen also from the results of Fiez, Shah, and Ratliff (2019). In contrast, any bidding that pushes PCMs to skip overdemanded papers will substantially improve the allocation under a broad range of conditions, and will reduce the exploration effort *during* bidding as a side effect. Thus everyone benefits: the PC chair has fewer orphan papers to deal with, reviewers are happier with their allocation of papers, and reviewers spend less time in the bidding process itself. This can be obtained either by reordering the papers as in (Fiez, Shah, and Ratliff 2019), or by placing a "price tag" on papers. These approaches can be naturally combined.

Our suggestions can be easily implemented in existing conference management platforms such as EasyChair, allowing reviewers to get a better lot of papers for less effort.

Temporal Considerations Since paper iprices change dynamically, one possible concern is that PCMs will choose strategically *when* to bid. We do not see this as a problem: PCMs are free to login and modify their bids any number of times, and in doing so they generally improve the assignment, as they can focus on the most relevant information. Thus there is no incentive to avoid early bidding. There is in fact one real reason and one perceived reason to bid early: placing an early bid on a popular paper will reduce its iprice and deter other potential reviewers, thereby increasing the chances of the early bidder to get the desired paper. Again this is not a problem in termas of fairness, as a reviewer that really wants the paper is free to bid on it regardless of its iprice, and the order of bids does not affect the assignment.

The perceived incentive for an early bid is to more easily reach the bidding requirement, as the iprices of popular paper will still be relatively high. However this is not a 'real' incentive since the requirement is only a recommendation anyways and is never enforced. Early PCMs who reached their budget requirement may turn out to be underbidding once more bidders entered their bids.. Crucially, this is also true in today's bidding scheme when PCMs have no information on the demand whatsoever. A possible improvement is to either remind early bidders that they can check and update their bids later on, or provide early bidders with a somewhat higher budget or project lower iprices. That said, more research is required to understand the temporal behavior of bidders and its effects on the outcome.

Multiple Bid Levels There are several ways of defining paper iprices in the presence of multiple bid levels. The technical properties of the PMA (with slight modifications) easily generalize. Bidders still have an incentive to be sincere: more precisely, they will place a maximal bid on some t most preferred papers, and no bid on papers below.

Privacy and Bias One may be concerned that information on the demand biases the judgment of PCMs is shown, and also that authors are unhappy about the popularity of their paper being public. We argue that these concerns are unjustified. First, there is no obvious connection between paper popularity and quality: PCMs bid on a paper either because of relevance to their specific interests, or because it seems to be an easy reject (Rodriguez, Bollen, and de Sompel 2007). 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. More importantly, revealing information on their demand does not harm unpopular papers: on the contrary, they are likely to get suitable reviewers in our trading post bidding scheme, whereas in the current system they may end up with reviewers who did not bid for them.

Next Steps We are currently designing lab experiments that will help us understand the bidding behavior of people with and without iprices. This should enable us to fine-tune the mechanism and understand how it can be best explained and implemented. We are also discussing using our scheme with the chairs of several medium-sized workshops.

Acknowledgements

This work was funded in part by the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute), and by the ANR project AGAPE ANR-18-CE23-0013.

We thank Laurent Gourves and Haris Aziz for valuable discussions on early versions of this paper. We thank Edith Elkind, Shaul Markovitch and Lirong Xia for sharing with us statistics of conferences they chaired.

References

Aziz, H.; Caragiannis, I.; Igarashi, A.; and Walsh, T. 2019a. Fair Allocation of Indivisible Goods and Chores. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI'19)*, 53–59.

Aziz, H.; Huang, X.; Mattei, N.; and Segal-Halevi, E. 2019b. The Constrained Round Robin Algorithm for Fair and Efficient Allocation. *arXiv preprint arXiv:1908.00161* .

Aziz, H.; Lev, O.; Mattei, N.; Rosenschein, J. S.; and Walsh, T. 2019c. Strategyproof peer selection using randomization, partitioning, and apportionment. *Artificial Intelligence* 275: 295–309.

Babaioff, M.; Nisan, N.; and Talgam-Cohen, I. 2017. Competitive equilibria with indivisible goods and generic budgets. *arXiv preprint arXiv:1703.08150* .

Bettman, J. R.; Johnson, E. J.; and Payne, J. W. 1991. Consumer Decision Making. In Robertson, T. S.; and Kassarjian, H. H., eds., *Handbook of Consumer Behavior*. Prentice-Hall.

Bogomolnaia, A.; Moulin, H.; Sandomirskiy, F.; and Yanovskaia, E. 2019. Dividing bads under additive utilities. *Soc. Choice Welf.* 52(3): 395–417.

Brânzei, S.; and Sandomirskiy, F. 2019. Algorithms for Competitive Division of Chores. *arXiv preprint arXiv:1907.01766* .

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

Charlin, L.; Zemel, R. S.; and Boutilier, C. 2011. A Framework for Optimizing Paper Matching. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence (UAI'11)*, 86–95.

Chaudhury, B. R.; Garg, J.; McGlaughlin, P.; and Mehta, R. 2020. Dividing Bads is Harder than Dividing Goods: On the Complexity of Fair and Efficient Division of Chores. *arXiv preprint arXiv:2008.00285* .

Fiez, T.; Shah, N.; and Ratliff, L. 2019. A SUPER* algorithm to optimize paper bidding in peer review. In *ICML workshop on Real-world Sequential Decision Making: Reinforcement Learning And Beyond*.

Freeman, R.; Sikdar, S.; Vaish, R.; and Xia, L. 2020. Equitable Allocations of Indivisible Chores. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'20)*, 384–392.

Garg, N.; Kavitha, T.; Kumar, A.; Mehlhorn, K.; and Mestre, J. 2010. Assigning Papers to Referees. *Algorithmica* 58(1): 119–136.

Goldsmith, J.; and Sloan, R. H. 2007. The AI Conference Paper Assignment Problem. In *Proceedings of the AAAI Workshop on Preference Handling for Artificial Intelligence*.

Kurokawa, D.; Lev, O.; Morgenstern, J.; and Procaccia, A. D. 2015. Impartial Peer Review. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI'15)*, 582–588.

Lian, J. W.; Mattei, N.; Noble, R.; and Walsh, T. 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)*, 1138–1145.

Mattei, N.; and Walsh, T. 2013. PrefLib: A Library for Preferences, HTTP://WWW.PREFLIB.ORG. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT'13)*.

Mattei, N.; and Walsh, T. 2017. A PREFLIB.ORG Retrospective: Lessons Learned and New Directions. In Endriss, U., ed., *Trends in Computational Social Choice*, chapter 15, 289–309. AI Access Foundation.

Othman, A.; Sandholm, T.; and Budish, E. 2010. Finding approximate competitive equilibria: Efficient and fair course allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'10)*, 873–880.

Rodriguez, M. A.; Bollen, J.; and de Sompel, H. V. 2007. Mapping the bid behavior of conference referees. *J. Informetrics* 1(1): 68–82.

Shapley, L.; and Shubik, M. 1977. Trade using one commodity as a means of payment. *Journal of Political Economy* 5(85): 937–968.

Xu, Y.; Zhao, H.; Shi, X.; and Shah, N. B. 2018. On Strategyproof Conference Peer Review. *arXiv preprint arXiv:1806.06266* .