

Into the Unknown: Assigning Reviewers to Papers with Uncertain Affinities

Cyrus Cousins^[0000–0002–1691–0282], Justin Payan^{*[0000–0001–7601–3500]},
and Yair Zick^[0000–0002–0635–6230]

University of Massachusetts Amherst, Amherst MA 01002, USA
{cbccousins, jpayan, yzick}@umass.edu

Abstract. A successful peer review process requires that qualified and interested reviewers are assigned to each paper. Most automated reviewer assignment approaches estimate a real-valued *affinity score* for each paper-reviewer pair that acts as a proxy for the quality of the match, and then assign reviewers to maximize the sum of affinity scores. Most affinity score estimation methods are inherently noisy: reviewers can only bid on a small number of papers, and textual similarity models and subject-area matching are inherently noisy estimators. Current paper assignment systems are not designed to rigorously handle noise in the peer-review matching market. In this work, we assume paper-reviewer affinity scores are located in or near a high-probability region called an *uncertainty set*. We maximize the worst-case sum of scores for a reviewer assignment over the uncertainty set. We demonstrate how to robustly maximize the sum of scores across various classes of uncertainty sets, avoiding potentially serious mistakes in assignment. Our general approach can be used to integrate a large variety of paper-reviewer affinity models into reviewer assignment, opening the door to a much more robust peer review process.

Keywords: Resource Allocation · Peer Review · Maximin Optimization

1 Introduction

Peer review is a fundamental institution for evaluating scientific knowledge. Over the 20th century, the scientific profession has grown significantly, and the institution of peer review has struggled with the increased scale. Modern computer science conferences receive thousands of submissions, matched to committees of similar size. Due to sheer scale, program chairs rely on automated reviewer assignment platforms, such as Microsoft CMT, OpenReview, and EasyChair, that utilize complex matching algorithms. Reviewing platforms generally implement a two-stage process to largely automate reviewer assignments [5, 18, 24]. First, the system estimates the “fit” between each paper-reviewer pair, called the *paper-reviewer affinity score*. Next, the system assigns reviewers via constrained optimization, maximizing a function of the computed affinity scores (usually the sum of scores, or *utilitarian welfare* [5]). Prior work identifies the reviewer assignment process as an important target for improving the overall quality of peer review in computer science [26, 31]. Efforts are underway to address shortcomings in reviewer assignment [14–16, 18, 21, 23, 25, 27, 30, 33], but none systematically addresses the fundamental issue of *uncertainty* in reviewer assignment.

Uncertainty in affinity score computation is a major source of error in assignment [18]. When we assign a reviewer to a paper, we are interested in ensuring the *quality of the future review*, which is fundamentally noisy. Because of this unpredictability, conferences typically construct affinity scores that reflect reviewer expertise and interest via four main sources of information. These sources include (a) *subject-area matching* (SAM) scores or keyword-based matching, where reviewer-provided areas of expertise are compared against keywords submitted by paper authors, (b) *textual similarity scores*, often implemented by the well-known Toronto Paper Matching System (TPMS) [5] or ACL scores [23], (c) *bidding*, where reviewers express their explicit ability and desire to review papers, and finally (d) *recommendations*, through which program committee members may suggest reviewers for papers. The overall affinity scores are typically computed as a linear combination of these four scores.¹ Recent conferences such as AAAI 2021 took a similar approach, linearly

* Corresponding Author. All authors contributed equally to this work.

¹ CMT implements their affinity scores this way, which can be seen from <https://cmt3.research.microsoft.com/docs/help/chair/auto-assign-reviewers.html>, as does OpenReview (source: personal correspondence).

combining TPMS scores, ACL scores, and SAM scores, and raising the sum to some power based on the reviewer bids [18].

Each of these common affinity score components can be missing or inaccurate. State-of-the-art document similarity measures disagree with expert judgments up to 43% of the time [34], and nearly 40% of TPMS scores were completely missing in AAAI 2021 [18].² Between 5% and 15% of papers in major AI conferences receive fewer than three positive bids, but there is evidence that many missing bids would be positive if collected [9, 21, 27]. Although no systematic study has been performed on keyword-based similarity scores, keyword matching accuracy depends on authors and reviewers using consistent terminology, and subtleties are invariably lost in the process. Even reviewers directly suggested by knowledgeable editors or the paper authors have been shown to perform surprisingly poorly on average, as measured by third-party annotators via the Review Quality Index [29, 36], showing that recommendations can be noisy as well.

To our knowledge, every reviewer assignment system still relies on affinity score estimates, but does not directly account for the fact that these scores are noisy estimates of assignment quality. Our work takes the first step towards addressing this fundamental gap. We investigate a generalized notion of affinity score, where organizers can implement affinity using any measure of fit between reviewers and papers. These measures may or may not be fully observable; for example, organizers may decide to estimate unknown bids as part of affinity computation. Our approach also enables even more advanced affinity measures, such as predictors of reviewer performance based on historical data.

1.1 Our Contributions

To properly account for uncertainty, we first construct a region called an *uncertainty set* which is close to the true affinities with high probability, then maximize the worst-case welfare over the uncertainty set. Uncertainty sets are a very general construct that allows conference organizers to introduce their own uncertainty models using available data and reasonable assumptions. Uncertainty sets generalize probability distributions — while it is possible to construct an uncertainty set from a probability distribution, non-Bayesian models will frequently *not* specify full probability distributions. In these cases, worst-case guarantees over an uncertainty set are quite natural. We call the problem of maximizing the worst-case welfare over an uncertainty set Reviewer Assignment under Uncertainty, or RAU.

We provide numerous examples of uncertainty sets throughout the paper, starting with axis-aligned, hyperrectangular uncertainty sets in Section 3.1 and spherical and ellipsoidal uncertainty sets in Section 3.2. Theorems 6 and 7 offer detailed end-to-end examples of how conference organizers can construct ellipsoidal uncertainty sets using bounds on the *square error* of an *affinity score estimator* from historic data or sampled bids (a frequentist approach that precludes optimizing for expected welfare). We also present a *calculus of uncertainty sets*, enabling construction of complex and highly informative uncertainty sets from simple components (Section 3.3). Our results are agnostic to the affinity model; organizers can define affinity scores arbitrarily, so long as they can be estimated for all paper-reviewer pairs, and sampled for some pairs.

We show that RAU is NP-hard over convex uncertainty sets (Theorem 10), and present an approximation algorithm called Robust Reviewer Assignment (RRA), which applies to any convex uncertainty set where worst-case welfare can be efficiently computed (Section 4). RRA applies *randomized rounding* methods to a *convex relaxation* of the discrete RAU problem, and we analyze both the *optimization error* due to convex programming methods and randomized rounding, and the *maximin error* due to operating with an uncertainty set, rather than known affinity scores (Section 4.2). We give bounds on the true welfare relative to the maximin welfare solution of RAU (Proposition 2), and explore the integrality gap of RAU (Proposition 13).

We empirically demonstrate the robustness of our approach relative to commonly-used baselines on publicly available data from five recent iterations of ICLR (Section 5.1). In addition, we explore synthetic settings where RRA avoids negative consequences faced by the most commonly used baseline (Section 5.2). We hope this work draws attention to the ad-hoc nature of affinity scores, spurring improvements to their computation and further study of their robustness. All proofs in this paper are contained in the appendix.

² Although the AAAI 2021 organizers do not explain why so many TPMS scores are missing, missing scores occur for several reasons, including reviewers opting out of the system or providing insufficient or empty publication records.

1.2 Related Work

Current automated peer review systems compute pointwise, ad-hoc affinity scores and maximize the total sum [5, 18, 25], the minimum value for papers [16, 33], or the minimum value for groups of papers [1].

Two recent studies start from our premise that commonly-used affinity scores may not be as accurate as they seem. Data is now available that directly compares elements of affinity scores to expert judgments of reviewer fit, and the authors of this dataset show that existing similarity score computation methods make many errors [34]. A recent study leverages randomness in assignment algorithms to directly judge assignment decisions, showing that higher weight should be placed on text similarity metrics over bids [28]. These results encourage further improvements to affinity score computation (especially text similarity), but also justify smarter utilization of the noisy sources of information that are available.

Other works use modern NLP techniques to improve document-based similarity scores [23], encourage reviewers to bid on underbid papers [9, 21, 27], or disincentivize strategic bidding behavior [13–15]. Although these approaches reduce uncertainty, they do not directly treat uncertainty in affinity scores.

Our robust optimization algorithm is based on an iterative supergradient-ascent approach; similar techniques have been applied to supervised learning with unknown labels [19] and fair learning with unknown group identities [7].

2 Reviewer Assignment under Uncertainty

Assume we have a set P of n papers submitted to a peer-review venue,³ and a set R of m reviewers. The key input to the reviewer assignment problem is a paper-reviewer *affinity score matrix* $S^* \in [0, 1]^{n \times m}$ (we will also refer to $[0, 1]^{n \times m}$ as the *unit hypercube*), where $S_{p,r}^* \in [0, 1]$ is the affinity of paper $p \in P$ to reviewer $r \in R$. This matrix S^* encodes the *true* affinities, or the value provided to the venue by assigning each reviewer to each paper. At this point, it is natural to ask, “How can one know the value a paper-reviewer assignment will provide ahead of time?” and “What do we mean by value provided?” These questions directly motivate our work; reviewer assignment is challenging because affinity depends on poorly-defined preferences over uncertain future outcomes.

Despite this fundamental challenge, all prior work assumes direct access to the true affinities [1, 5, 16, 18, 25, 33]. We relax this assumption, instead assuming only *partial knowledge* of S^* , as represented by an *uncertainty set* $\mathcal{S} \subseteq [0, 1]^{n \times m}$ that contains a point near S^* with high probability.

Definition 1 (((δ, γ) Uncertainty Set). Suppose $S^* \in [0, 1]^{n \times m}$ is the true affinity score matrix. A (δ, γ) uncertainty set obeys $\mathbb{P}(\inf_{S \in \mathcal{S}} \|S - S^*\|_1 > \gamma) < \delta$, i.e., it probably contains some S that is γ -close to S^* .

Once we compute a (δ, γ) uncertainty set \mathcal{S} , our goal is to find an *assignment matrix* A that assigns reviewers to papers with good worst-case guarantees on a score function defined by A and S^* , while satisfying some hard constraints. Assignments are deterministic, i.e. lie on vertices of the unit hypercube, or $\mathcal{A}_0 \doteq \{0, 1\}^{n \times m}$. We compute an assignment of reviewers to papers $A \in \mathcal{A}_0$, where $A_{p,r} = 1$ if and only if r is assigned to review the paper p . Hard constraints usually include the number of reviewers required per paper, upper bounds on reviewer loads, and conflicts of interest. Conflicts of interest consist of a set $\mathcal{C} \subseteq (P \times R)$, where $(p, r) \in \mathcal{C}$ implies that r cannot be assigned to p . $\mathcal{A}_{\text{CoI}} \subseteq \mathcal{A}_0$ denotes the set of all assignments respecting conflicts in \mathcal{C} . If $A \in \mathcal{A}_{\text{CoI}}$ and $(p, r) \in \mathcal{C}$, then $A_{p,r} = 0$.

Suppose that each paper p requires exactly k_p reviewers and each reviewer r must be assigned no more than u_r papers. Then

$$\mathcal{A}_P \doteq \left\{ A \in \mathcal{A}_0 \mid \forall p \in P: \sum_{r \in R} A_{p,r} = k_p \right\} \text{ \& } \mathcal{A}_R \doteq \left\{ A \in \mathcal{A}_0 \mid \forall r \in R: \sum_{p \in P} A_{p,r} \leq u_r \right\}$$

define the *paper coverage requirements* and *reviewer load bounds*, respectively. We occasionally refer to the *total review load* $K \doteq \sum_{p \in P} k_p$. Taken together, the hard constraints give us a set of permissible assignments $\mathcal{A} \doteq \mathcal{A}_{\text{CoI}} \cap \mathcal{A}_P \cap \mathcal{A}_R$.

³ A peer review venue is any entity which assigns reviewers to papers for the purposes of peer review. The prototypical venue is a peer-reviewed conference, but our approach applies to similar venues such as ACL Rolling Review [22].

We aim to compute assignments A that maximize a score-based objective function $W(A, S)$, while meeting all hard constraints. Throughout the paper we assume W is utilitarian social welfare $W(A, S) \doteq \frac{1}{n} \sum_{p \in P} \sum_{r \in R} A_{p,r} S_{p,r}$.

In summary, the problem of Reviewer Assignment under Uncertainty (RAU) takes as input a set of papers P , reviewers R , assignment constraints \mathcal{A} , and an uncertainty set \mathcal{S} . Although ideally we would compute $A^* \doteq \arg \max_{A \in \mathcal{A}} W(A, S^*)$, we cannot since we do not know S^* . Therefore, our goal is to find or approximate

$$A^{\text{MM}} \doteq \arg \max_{A \in \mathcal{A}} \inf_{S \in \mathcal{S}} W(A, S) . \quad (1)$$

Here A^{MM} maximizes welfare for adversarial (worst-case) affinity scores over the uncertainty set, which ensures robustness with high probability. Robust guarantees prevent catastrophic failures in exchange for some average case penalty.

We conclude this section with a simple observation relating the worst-case welfare over \mathcal{S} for any assignment A to the true welfare of A with S^* .

Proposition 2 (Relating True and Worst-Case Welfare over \mathcal{S}). *Suppose \mathcal{S} is a (δ, γ) uncertainty set with $\|S - S'\|_1 \leq L$ for all $S, S' \in \mathcal{S}$, and the true affinity score matrix is labeled S^* . Consider any assignment $A \in \mathcal{A}$. Then with probability at least $1 - \delta$,*

$$W(A, S^*) - \frac{L + \gamma}{n} \leq \inf_{S \in \mathcal{S}} W(A, S) \leq W(A, S^*) + \frac{\gamma}{n} .$$

Proposition 2 implies that if we aim for low \mathcal{L}_1 diameter uncertainty sets \mathcal{S} , we can approximately optimize true welfare using the robust objective (1). A similar bound holds when we compare $W(A^{\text{MM}}, S^*)$ to $W(A^*, S^*)$, i.e., to the welfare of the unknown, true optimal assignment A^* (Theorem 14). However, Proposition 2 states that the maximin objective itself is close to the true welfare under any assignment A , whereas Theorem 14 shows that the specific assignment A^{MM} has low regret against A^* .

3 Uncertainty Models for Affinity Scores

We start by analyzing simple uncertainty sets, namely the case where S is known, as well as hyperrectangular, spherical, and ellipsoidal uncertainty sets. We conclude the section by showing a compositionality rule that allows venue organizers to combine multiple uncertainty sets into a single RAU problem.

The case where S is known, $\mathcal{S} = \{S\}$, is quite straightforward. This corresponds to most current conference management system implementations [5, 18]. When $\mathcal{A} = \mathcal{A}_{\text{CoI}} \cap \mathcal{A}_P \cap \mathcal{A}_R$, the binary integer program $\arg \max_{A \in \mathcal{A}} W(A, S)$ is known to be polynomial-time solvable, as it is a linear program with totally unimodular constraints. Intuitively, these constraints introduce cuts to the assignment hypercube that never produce new vertices, thus all vertices of the resulting polytope occur on the binary integer lattice. Efficient implementations of this linear program can be found in multiple publicly available sources [16, 25, 33, 35]. Let us now consider hyperrectangular, spherical, and ellipsoidal uncertainty sets.

3.1 Hyperrectangular Uncertainty Sets from Confidence Intervals

Many simple and intuitive models for an uncertainty set take the form of axis-aligned hyperrectangles. A naïve uncertainty set might estimate confidence intervals for each $S_{p,r}$ independently and use a union bound to give a high-probability region for the affinity scores. Venue organizers might also make assumptions about intervals bounding affinity scores with certainty, taking the intersection of multiple such interval bounds. For example, they might start with the global constraints of the unit hypercube. Lower and upper bounds can then be given for pairs based on whether they receive certain bids, whether the program committee recommends the assignment, or whether a threshold on document similarity score is met. This model is ad-hoc and simple, but may be suitable in practice. Furthermore, if we assume that with probability at least $1 - \delta$, only a small constant fraction γ of these bounds can be violated, we can establish a (δ, γ) confidence interval under more realistic assumptions.

If we take all the lower bounds on paper-reviewer scores, we obtain a lower bound affinity score matrix \underline{S} . Similarly, taking all the maximal possible values for paper-reviewer scores yields an upper bound affinity score matrix \bar{S} . Our uncertainty set is thus $\mathcal{S}_\square \doteq \{X \in \mathbb{R}^{n \times m} \mid \underline{S}_{p,r} \leq X_{p,r} \leq \bar{S}_{p,r} \ \forall p, r\}$. Our first result is that RAU can be solved in polynomial time for axis-aligned, hyperrectangular uncertainty sets.

Theorem 3 (RAU under Hyperrectangular Uncertainty). *When the uncertainty set is an axis-aligned hyperrectangular region \mathcal{S}_\square , then*

$$\arg \max_{A \in \mathcal{A}} \inf_{S \in \mathcal{S}_\square} W(A, S) = \arg \max_{A \in \mathcal{A}} W(A, \underline{S}) .$$

Thus RAU maximizes W for \underline{S} , which requires polynomial time via LP reduction.

Axis-aligned, hyperrectangular uncertainty sets correspond to the case where uncertainty is bounded independently across paper-reviewer scores, hence their relative simplicity. Although hyperrectangular uncertainty sets are easy to work with, they are unnecessarily pessimistic, since it is very unlikely that all affinities take extreme values at once (i.e., \underline{S} is actually a very unlikely outcome). We can improve our estimates using uncertainty set models that account for the low probability of many simultaneous extreme values.

3.2 Ellipsoidal Uncertainty Sets with \mathcal{L}_2 Error Guarantees

Many standard models directly bound the \mathcal{L}_1 or \mathcal{L}_2 error of their predictions, which implies uncertainty sets that are more optimistic than hyperrectangular \mathcal{S} (and hence have tighter guarantees for Proposition 2).

We first analyze the case of symmetric uncertainty sets with \mathcal{L}_2 error guarantees. For example, we might solicit bids uniformly at random and then predict unsampled bids using collaborative filtering with \mathcal{L}_2 error guarantees [4, 8, 17]. \mathcal{S} could then be constructed as a linear combination of values known with certainty (document-based similarity scores and keyword-based matching scores) and the real and estimated bids, yielding a spherical uncertainty set \mathcal{S} .

We can consider a spherical $(\delta, 0)$ uncertainty set to consist of a point estimate S^0 and a radius ε limiting the \mathcal{L}_2 error from the point estimate S^0 . Formally, we aim to solve the problem $\arg \max_{A \in \mathcal{A}} \inf_{S \in B_\varepsilon(S^0)} W(A, S)$, where $B_\varepsilon(S^0) \doteq \{X \in \mathbb{R}^{n \times m} \mid \|X - S^0\|_F \leq \varepsilon\}$ denotes the ε Frobenius-norm ball around S^0 .

Theorem 4 (RAU under Spherical Uncertainty). *When \mathcal{S} is a sphere,*

$$\arg \max_{A \in \mathcal{A}} \inf_{S \in B_\varepsilon(S^0)} W(A, S) = \arg \max_{A \in \mathcal{A}} W(A, S^0) ,$$

which can be computed in polynomial time. Furthermore, for all $A \in \mathcal{A}$,

$$|W(A, S^0) - W(A, S^*)| \leq \frac{\varepsilon \sqrt{K}}{n} .$$

One way of looking at Theorem 4 is that A^{MM} over spherical uncertainty sets provides no additional robustness guarantees over $\arg \max_{A \in \mathcal{A}} W(A, S^0)$. Thus, it seems that in order to obtain meaningful robustness guarantees, we require both a limit to the total amount of variation in affinity scores (Theorem 3) as well as asymmetry between the noise on affinity scores (Theorem 4). We can explain these results intuitively; to decide how best to assign reviewers, we need to be able to make tradeoffs between assigning pairs with potentially high (but also potentially low) affinity or assigning pairs that have an average amount of affinity with higher certainty. Those tradeoffs are only meaningful if uncertainty varies across paper-reviewer pairs, and if there is a limited total amount of uncertainty.

With that intuition in mind, we generalize to the case of ellipsoidal uncertainty sets. In a simple model, we might model affinity scores as multivariate Gaussians, which we explore in greater detail in the experiments of Section 5. In this case, we obtain a mean vector⁴ $\vec{\mu} \in [0, 1]^{nm}$ and a positive semi-definite covariance matrix

⁴ Technically, a sample from this multivariate Gaussian is a vector in \mathbb{R}^{nm} , and must be reshaped into a matrix $S \in \mathbb{R}^{n \times m}$. We will convert matrices into vectors in row-major order and ignore the distinction between matrices and vectors when convenient.

$\Sigma \in \mathbb{R}^{nm \times nm}$. Given a confidence level $1 - \delta$, we create an uncertainty set

$$\mathcal{S} \doteq \{S \in \mathbb{R}^{nm} \mid (S - \bar{\mu})^T \Sigma^{-1} (S - \bar{\mu}) \leq \chi_{nm}^2(1 - \delta)\} \quad (2)$$

where χ_k^2 is the inverse CDF of the χ^2 distribution with k degrees of freedom. If we assume no model error (which is not a safe assumption generally), then the true affinity scores are contained within \mathcal{S} with probability at least $1 - \delta$. We know also that $\chi_{nm}^2(1 - \delta) \leq nm + 2\sqrt{nm \ln \frac{1}{\delta}} + 2 \ln \frac{1}{\delta}$. We can see that the size of the uncertainty set depends only logarithmically on $\frac{1}{\delta}$, and thus we can trade off between δ and the \mathcal{L}_1 diameter of the uncertainty set.

While the Gaussian model employed in (2) makes a strong modeling assumption, we can use a validation set and a predictive model with provable tail bounds to obtain uncertainty sets that do not require any such assumptions. To accomplish this, we require a predictive model of affinity scores. Venue organizers can then obtain \mathcal{L}_2 error bounds using validation data, and this will yield an ellipsoid due to sampling bias (only certain paper-reviewer pairs will be observed for any given venue). In this setting, the uncertainty set is not derived as a confidence interval of a probability distribution, but rather directly comes from a tail bound on total generalization error. Optimizing in this setting *requires* using a robust approach like RAU, and cannot be done with average case analysis.

We now develop this predictive model in more technical detail. Suppose the true affinity score of some paper-reviewer pair is $f^*(p, r)$, and we have access to a predictive model $\hat{f}(p, r)$, perhaps learned on historical venues. In practice, \hat{f} predicts the affinity of reviewer r for paper p based on any information available prior to reviewer assignment, and the specific definition of affinity is left to the venue organizers. For example, a venue may decide that affinity is best measured via reviewer bids, and they may use historical data to train a predictor \hat{f} to predict missing bids from document-based similarity scores and keywords. Alternatively, venue organizers may decide that the ground truth affinity $f^*(p, r)$ should correspond to a meta-reviewer's judgment of review quality, and \hat{f} can then be trained on historical data to predict these judgments.

We will take $\hat{S}_{p,r} \doteq \hat{f}(p, r)$ and $S_{p,r}^* \doteq f^*(p, r)$. We assume that we can evaluate \hat{f} on all paper-reviewer pairs in the current venue, and potentially on pairs from historical venues as well. We may be able to sample $f^*(p, r)$ for some, but not all, pairs in the current venue and historical venues (the validation data). We then probabilistically bound a weighted average of the square error between \hat{f} and f^* in terms of an estimate of expected square error computed on the validation set. The details of the validation set vary by application, but the overall strategy will be to estimate the square error $\mathbb{E}[(S^* - \hat{S})^2]$, where S^* and \hat{S} are given by f^* and \hat{f} on a random paper-reviewer pair. We show two such approaches, the first inductive, using historic or auxiliary data to form the validation set, and the second transductive, assuming a small random sample of true affinity scores (e.g., bids) can be queried within the current peer-review venue. We will need to define the notion of sampling a sequence of random variables *conditionally independently without replacement*.

Definition 5. A sequence of variables x_1, x_2, \dots, x_t , where all $x_i \in \mathcal{X}$, is sampled conditionally independently without replacement from a distribution \mathcal{D} with support \mathcal{X} if the variables x_i are sampled in order from x_1 to x_t , and for any $i \in \{1, \dots, t\}$, $p_{\mathcal{D}}(x_i = x | x_1, \dots, x_{i-1}) = \frac{p_{\mathcal{D}}(x_i = x)}{\int_{\mathcal{X} \setminus \{x_1, \dots, x_{i-1}\}} p_{\mathcal{D}}(x_i = y) dy}$ for all $x \in \mathcal{X} \setminus \{x_1, \dots, x_{i-1}\}$ and $p_{\mathcal{D}}(x_i = x | x_1, \dots, x_{i-1}) = 0$ for $x \in \{x_1, \dots, x_{i-1}\}$.

Theorem 6 (Ellipsoidal Uncertainty Sets from Inductive Predictors). Let \mathcal{D}' be a probability distribution over paper-reviewer pairs, and let $\mathcal{D}^{\mathcal{P}}$ and $\mathcal{D}^{\mathcal{R}}$ be distributions over papers and reviewers, respectively. Assume that P and R were drawn conditionally independently without replacement from $\mathcal{D}^{\mathcal{P}}$ and $\mathcal{D}^{\mathcal{R}}$, respectively.⁵ Suppose we sample T paper-reviewer pairs $\{(p_i, r_i)\}_{i=1}^T$ conditionally independently without replacement from \mathcal{D}' , and these paper-reviewer pairs have true and estimated affinity scores $\{f^*(p_i, r_i)\}_{i=1}^T$ and $\{\hat{f}(p_i, r_i)\}_{i=1}^T$, respectively.

Let

$$\alpha(p, r) \doteq \frac{\mathbb{P}_{\mathbf{p} \sim \mathcal{D}^{\mathcal{P}}}(\mathbf{p} = p) \mathbb{P}_{\mathbf{r} \sim \mathcal{D}^{\mathcal{R}}}(\mathbf{r} = r)}{\mathbb{P}_{(\mathbf{p}, \mathbf{r}) \sim \mathcal{D}'}((\mathbf{p}, \mathbf{r}) = (p, r))} \quad \& \quad \alpha_{\min} \doteq \inf_{p \in \mathcal{P}, r \in \mathcal{R}} \alpha(p, r)$$

⁵ Abusing notation slightly, we index the papers $p \in P$ and reviewers $r \in R$ so that p can represent either a paper in \mathcal{P} or an integer between 1 and n .

denote (1) the probability ratio of sampling p from \mathcal{D}^P and r from \mathcal{D}^R to sampling (p, r) from \mathcal{D}' , and (2) the infimum probability ratio, respectively. Now construct the ellipsoid matrix $\Sigma \in \mathbb{R}^{nm \times nm}$ as the diagonal matrix such that $\Sigma_{pm+r, pm+r} = \alpha(p, r)$ for all $p \in P, r \in R$.

Then for any $\delta \in (0, 1)$, the ellipsoid

$$\mathcal{S} \doteq \left\{ S \in \mathbb{R}^{nm} \left| \frac{1}{nm} (S - \hat{S})^\top \Sigma^{-1} (S - \hat{S}) \leq \underbrace{\frac{1}{T} \sum_{i=1}^T (f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2}_{=\hat{\xi}} + \underbrace{\sqrt{\left(\frac{1}{T} + \frac{n+m}{nm\alpha_{\min}^2}\right) \frac{\ln \frac{1}{\delta}}{2}}}_{=\eta} \right. \right\}$$

is a $(\delta, 0)$ uncertainty set, where $\hat{\xi}$ denotes the empirical square error of our estimated scores, and η denotes the excess error bound due to sampling.

Departing from the standard reviewer assignment setup, Theorem 6 assumes that both the historic data and the current venue are random. In particular, historic paper-reviewer pairs are sampled from \mathcal{D}' (modeling the historic data generation process), and papers and reviewers for the current venue are sampled from \mathcal{D}^P and \mathcal{D}^R (modeling the processes by which papers are submitted and reviewers volunteer). We then construct $\alpha(p, r)$ to reweight square error on the current venue to match expected square error on \mathcal{D}' (i.e., we use importance sampling to calibrate expectations over \mathcal{D}' versus those over \mathcal{D}^P and \mathcal{D}^R). For example, \mathcal{D}' reflects all elements of historic data generation, most importantly the availability of historic data from multiple venues with different focuses. We might then use the (relatively stable) topic areas of papers and reviewers to model \mathcal{D}^P and \mathcal{D}^R , and thus $\alpha(p, r)$ reflects the ratio of the popularity of p and r 's topic areas in the current venue to historic venues.

We show a similar result in the transductive setting. Instead of constructing a predictive function from historical data, we generalize a small set of known affinities for the current venue to the unknown affinities for the same venue. Note that \mathcal{D}' now reflects the process by which we obtain samples for (p_i, r_i) pairs from the current venue, rather than from historical venues.

Theorem 7 (Ellipsoidal Uncertainty Sets from Transductive Predictors). Suppose we sample T paper-reviewer pairs $\{(p_i, r_i)\}_{i=1}^T$ conditionally independently without replacement from \mathcal{D}' , and these paper-reviewer pairs have true and estimated affinity scores $\{f^*(p_i, r_i)\}_{i=1}^T$ and $\{\hat{f}(p_i, r_i)\}_{i=1}^T$, respectively. Let

$$\alpha(p, r) \doteq \frac{(nm)^{-1}}{\mathbb{P}_{(\mathbf{p}, \mathbf{r}) \sim \mathcal{D}'}((\mathbf{p}, \mathbf{r}) = (p, r))}$$

denote the probability ratio between sampling (p, r) uniformly at random and sampling (p, r) from \mathcal{D}' , and construct the ellipsoid matrix $\Sigma \in \mathbb{R}^{nm \times nm}$ as the diagonal matrix such that $\Sigma_{pm+r, pm+r} = \alpha(p, r)$ for all $p \in P, r \in R$.

Then for any $\delta \in (0, 1)$, the ellipsoid

$$\mathcal{S} \doteq \left\{ S \in \mathbb{R}^{nm} \left| \frac{1}{nm} (S - \hat{S})^\top \Sigma^{-1} (S - \hat{S}) \leq \underbrace{\frac{1}{T} \sum_{i=1}^T (f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2}_{=\hat{\xi}} + \underbrace{\sqrt{\frac{\ln \frac{1}{\delta}}{2T}}}_{=\eta} \right. \right\}$$

is a $(\delta, 0)$ uncertainty set, where $\hat{\xi}$ denotes the empirical square error of our estimated scores, and η denotes the excess error bound due to sampling.

The transductive result can be straightforwardly applied to many different contexts in which venue organizers can solicit samples of f^* on (p_i, r_i) pairs from the current venue, rather than historical data. This information must be obtained prior to assigning the majority of reviewers. For example, organizers could define f^* as a reviewer's hypothetical bid and Theorem 7 then requires soliciting a small number of bids to estimate the error of \hat{f} . Similarly, f^* could correspond to meta-reviewer judgments of review quality, accomplished by opting for a two-stage reviewing process, in which the reviews and feedback generated in the first stage are used to estimate the error of \hat{f} . These definitions of f^* are costly to sample, but organizers can still efficiently target sophisticated affinity models by solving RAU over the uncertainty sets of Theorem 7.

We naturally ask the question, "How many samples are sufficient to obtain a sharp confidence bound?" Observe that, by Proposition 2, the gap between adversarial and true welfare is $\frac{L}{n}$, where L denotes

the \mathcal{L}_1 diameter of \mathcal{S} . For the ellipsoidal uncertainty set of Theorem 7, $\frac{L}{n} \leq 2m\sqrt{\alpha_{\max}(\hat{\xi} + \eta)}$, where $\alpha_{\max} \doteq \sup_{p \in \mathcal{P}, r \in \mathcal{R}} \alpha(p, r)$. Furthermore, the empirical square error $\hat{\xi}$ converges to some ξ as T increases, thus we need only select $T \in \Omega\left(\frac{\log \frac{1}{\delta}}{\xi^2}\right)$ samples to ensure that the uncertainty set is constant-factor optimal, at which point the welfare gap is $\mathbf{O}(m\sqrt{\alpha_{\max}\xi})$, which is also optimal to within constant factors. Notably, the sufficient sample size T is *independent of the venue size* (i.e., n and m), thus the added burden of soliciting these extra bids is negligible. We also see that the fundamental limitation of this method is the average square error ξ , which depends on the predictor, the venue, and the sampling distribution \mathcal{D}' . It is thus paramount to use predictors for which this quantity will be small. Fortunately, this is often the case, as many predictive models are explicitly trained to minimize \mathcal{L}_2 error on some task, which motivates the choice of our ellipsoidal uncertainty sets. Note that while this argument pertains to Theorem 7, one can argue similarly for the necessary size of the validation set to ensure $\eta = \mathbf{O}(\xi)$ in Theorem 6.

Finally, we note that it is possible to extend either result under less favorable (more realistic) assumptions about the sampling process using the γ parameter (\mathcal{L}_1 error) of our uncertainty set construction. In particular, either result produces a $(\delta, T\gamma)$ uncertainty set if \hat{s} , s^* , and the associated (p, r) pairs are subject to *adversarial corruption* of $T\gamma$ of the validation set samples drawn from \mathcal{D}' , which has immediate applications in privacy, adversarial robustness, and various notions of strategy-proofness. Furthermore, to model more complicated and potentially not fully understood distribution shift, we obtain via Bennett’s inequality [2] a $\left(\delta + \delta', T\gamma + \frac{1}{3} \ln \frac{1}{\delta'} + \sqrt{2T\gamma(1-\gamma) \ln \frac{1}{\delta'}}\right)$ uncertainty set if the validation set is instead drawn from some \mathcal{D}'' such that $\text{TVD}(\mathcal{D}', \mathcal{D}'') \leq \gamma$.

3.3 Compositional Rules

We may often have more complicated uncertainty sets than the simple geometries described in the previous sections. For example, we can intersect the constraints of the unit hypercube with an ellipsoidal uncertainty set as described in Section 3.2. This produces a *truncated ellipsoid*, a common construction that we will see again in Section 5.

Lemma 8 (Uncertainty Set Intersection). *Suppose each \mathcal{S}_i for $i \in [k]$ is a $(\vec{\delta}_i, 0)$ uncertainty set. Then $\mathcal{S}_{\cap} \doteq \bigcap_{i=1}^k \mathcal{S}_i$ is a $(\|\vec{\delta}\|_1, 0)$ uncertainty set.*

We can also use Lemma 9 to convert (δ, γ) uncertainty sets to larger $(\delta, 0)$ uncertainty sets.

Lemma 9 (\mathcal{L}_1 Error Terms). *If \mathcal{S} is a (δ, γ) uncertainty set, then for any $\eta \in [0, \gamma]$, it holds that the Minkowski sum*

$$\mathcal{S}' \doteq \mathcal{S} + \{\vec{s} \in \mathbb{R}^{nm} \mid \|\vec{s}\|_1 \leq \eta\} = \{\vec{x} + \vec{s} \mid \vec{x} \in \mathcal{S}, \|\vec{s}\|_1 \leq \eta\}$$

is a $(\delta, \gamma - \eta)$ uncertainty set.

We can apply Lemmas 8 and 9 sequentially to intersect arbitrary (δ, γ) uncertainty sets. We first expand them via Lemma 9 to obtain larger $(\vec{\delta}_i, 0)$ uncertainty sets, and we then apply Lemma 8 to obtain their intersection.

We can also apply these results to the uncertainty sets previously described in Section 3. For example, the intersection of multiple axis-aligned, hyperrectangular constraints produces an axis-aligned hyperrectangle. This may occur when structural constraints defined by the venue (e.g., hard upper and lower bounds defined based on topic overlap) intersect with per-pair error bounds. Similarly, we might consider cases with two intersecting ellipsoidal error bounds derived from two different estimators using Theorems 6 and 7. This intersection is uninteresting if the two ellipsoids have the same centroid and one is strictly smaller than the other, but if these ellipsoids have different centroids (as when the estimators have different biases) their intersection can be quite beneficial.

Algorithm 1 Robust Reviewer Assignment (RRA)

Require: Error tolerance ε , supergradient norm bound λ , uncertainty set \mathcal{S} , constrained allocation space $\tilde{\mathcal{A}}$, total review load $K = \sum_{p \in P} k_p$ (i.e., the total number of assignments required)

- 1: Initialize $S^{(0)} \in \mathcal{S}$ arbitrarily
- 2: $A^{(0)} \leftarrow \arg \max_{A \in \tilde{\mathcal{A}}} W(A, S^{(0)})$ // Initialize $A^{(0)}$ to optimize $S^{(0)}$ (via LP reduction)
- 3: $\hat{A} \leftarrow A^{(0)}; \hat{w} \leftarrow -\infty$ // Maintain best allocation \hat{A} and adversarial welfare \hat{w}
- 4: $T \leftarrow \lceil 2K(\frac{\lambda}{\varepsilon})^2 \rceil; \alpha \leftarrow \frac{\varepsilon}{\lambda T}$ // Compute sufficient *step count* T and *step size* α
- 5: **for** $t \in \{1, 2, \dots, T\}$ **do**
- 6: $S^{(t)} \leftarrow \arg \inf_{S \in \mathcal{S}} W(A^{(t-1)}, S)$ // Adversary selects $S^{(t)}$ from \mathcal{S}
- 7: **if** $W(A^{(t-1)}, S^{(t)}) > \hat{w}$ // Update \hat{A} if adversarial welfare beats previous best
- 8: $\hat{A} \leftarrow A^{(t-1)}; \hat{w} \leftarrow W(A^{(t-1)}, S^{(t)})$
- 9: $A^{(t)} \leftarrow A^{(t-1)} + \alpha \nabla_{A^{(t-1)}} W(A^{(t-1)}, S^{(t)})$ // Update allocation with a supergradient step
- 10: $A^{(t)} \leftarrow \arg \min_{A \in \tilde{\mathcal{A}}} \|A - A^{(t)}\|_2$ // \mathcal{L}_2 project onto feasible allocation set $\tilde{\mathcal{A}}$
- 11: **return** $\text{ROUND}(\hat{A})$ // Sample integral assignment

4 Robust Reviewer Assignment (RRA)

We now present a general purpose algorithm for approximately solving the RAU problem over convex uncertainty sets, as long as the adversarial (worst-case) welfare can be computed in polynomial time. We first show in Theorem 10 that RAU is NP-hard in general for convex uncertainty regions of this type.

Theorem 10 (Hardness of RAU). *RAU is NP-hard over a convex uncertainty set \mathcal{S} , even for \mathcal{S} with a polynomial-time adversary. In particular, RAU remains NP-hard even when \mathcal{S} is restricted to bounded polytopes formed by intersections of polynomially many halfspaces.*

4.1 Robust Reviewer Assignment

Due to this hardness result, we outline an approach to approximately solve RAU efficiently for convex uncertainty sets with polynomial-time adversaries. We start by allowing *fractional* (rather than binary) assignments. We then apply *supergradient ascent* (analogous to subgradient descent) to approximate $\arg \max_{A \in \tilde{\mathcal{A}}} \inf_{S \in \mathcal{S}} W(A, S)$, where $\tilde{\mathcal{A}}$ is the convex closure of the feasible set of discrete allocations \mathcal{A} from Section 2. When the supergradient ascent algorithm terminates, we randomly round the assignment to a binary assignment.

In particular, we present Algorithm 1, termed RRA. RRA applies an iterative adversarial optimization strategy to the objective. In each iteration t , we take an *adversary step*, which identifies the pessimal $S^{(t)}$ given assignment $A^{(t-1)}$. We then take a *gradient ascent step* from $A^{(t-1)}$ to $A^{(t)}$ assuming the score matrix remains fixed at $S^{(t)}$, followed by a *projection step*, which ensures $A^{(t)}$ remains feasible (i.e., does not violate any constraints on assignments). The gradient ascent step is valid since the gradient $\nabla_A W(A, \arg \inf_{S \in \mathcal{S}} W(A, S))$ is an element of the supergradient $\nabla_A \inf_{S \in \mathcal{S}} W(A, S)$.

The number of iterations required to prove convergence depends on the *gradient norm bound* λ , which is the smallest term such that $\|\nabla_A W(A, S)\|_2 \leq \lambda$ for all A and S . We approximate the maximin optimal continuous matrix \tilde{A}^{MM} within an error of ε in number of iterations polynomial in λ , $\frac{1}{\varepsilon}$, and the total review load K . The time complexity of RRA also depends on the *adversarial minimization* and *projection* steps, but so long as these take polynomial time, then so too does RRA. We state the convergence results in Proposition 11. The proof applies standard convergence results for subgradient descent [32].

Proposition 11 (Supergradient Ascent Efficiency). *Let λ denote an upper bound on the \mathcal{L}_2 norm of the supergradient elements $\nabla_A W(A, S)$ used in RRA. The supergradient ascent component of RRA converges to within ε of the maximin optimal continuous assignment \tilde{A}^{MM} in $\lceil 2K(\frac{\lambda}{\varepsilon})^2 \rceil$ iterations. RRA runs in time $\mathcal{O}(2KC(\frac{\lambda}{\varepsilon})^2)$, where C is the time cost of one adversary and projection step.*

Although the bound on the number of iterations can be quite large, it proves the convex relaxation of RAU is solvable in polynomial time, as long as the adversary and projection steps can be solved in polynomial

time and λ is bounded. In addition, the required number of iterations until convergence will typically be much smaller in practice.

The complexity result improves in the case of (truncated) ellipsoidal uncertainty sets. The *adversarial minimization* step requires polynomial time under truncated ellipsoidal uncertainty sets, as it is a *linear objective* under *convex quadratic constraints* (and box constraints), which is a *second-order conic program*, and the *projection step* always requires polynomial time, as it is a *convex quadratic* objective under *linear constraints* (i.e., the assignment constraints $\tilde{\mathcal{A}}$). The bound λ can be difficult to compute in the general case, but we show λ is typically well-bounded in the case of truncated ellipsoidal uncertainty sets.

Corollary 12 (Supergradient Ascent Efficiency under Ellipsoidal Uncertainty). *For a truncated ellipsoidal uncertainty set, the supergradient ascent component of RRA converges to within ε of the maximin optimal continuous assignment \hat{A}^{MM} in $\mathbf{O}\left(\frac{2K_m}{n\varepsilon^2}\right)$ iterations.*

Finally, we can round using the extended Birkhoff von Neumann decomposition sampling algorithm [3, 10, 15]. This sampling algorithm generates an integral sample \hat{A}' from the distribution defined by the continuous assignment matrix \hat{A} . The sample \hat{A}' still satisfies the constraints of $\tilde{\mathcal{A}}$, and $\mathbb{E}_{\hat{A}'}[\hat{A}] = \hat{A}$. The time complexity of this sampling algorithm is $\mathbf{O}(mn(m+n))$, which is typically negligible compared to the complexity of supergradient ascent.

4.2 Maximin and Integrality Gaps

Aside from the optimization error of RRA (Algorithm 1), there are two more error sources: *maximin error* for working under uncertainty, and *rounding error*.

The integrality gap of RAU can be quite large; similarly, the \mathcal{L}_1 difference between a rounded assignment A' and a continuous assignment \tilde{A} may be quite large as well. Surprisingly, we show in this section that although the integrality gap is large, with high probability this does not translate to a large amount of suboptimality in the *true* W of the rounded solution A' . Intuitively, whenever the maximin optimal continuous solution \tilde{A}^{MM} has a high \mathcal{L}_1 distance from any valid binary integer assignment, the decisions made during rounding cancel out on average, and have relatively little impact on the true welfare of the assignment.

Proposition 13 (\mathcal{L}_1 Distance to Integral Solution). *Suppose an unrounded assignment \tilde{A} and a randomized rounding A' of \tilde{A} such that $\mathbb{E}[A'] = \tilde{A}$. Then the expected \mathcal{L}_1 deviation of the assignment due to rounding obeys*

$$\mathbb{E}_{A'}[\|A' - \tilde{A}\|_1] = 2(\|\tilde{A}\|_1 - \|\tilde{A}\|_2^2) \leq nm - 2\|\frac{1}{2} - \tilde{A}\|_1.$$

Although the assignments may need to be rounded quite significantly, Theorem 14 shows that the rounded assignment produced by RRA has near-optimal true welfare (in expectation). Theorem 14 also allows for ε error in the discrete/continuous maximin assignments. For this, we define $A^\varepsilon \in \mathcal{A}$, representing any ε -optimal discrete solution to RAU and $\tilde{A}^\varepsilon \in \tilde{\mathcal{A}}$, any ε -optimal continuous solution to RAU. A^ε and \tilde{A}^ε are formally defined by the properties

$$\max_{A \in \mathcal{A}} \inf_{S \in \mathcal{S}} W(A, S) - \inf_{S \in \mathcal{S}} W(A^\varepsilon, S) \leq \varepsilon \quad \& \quad \max_{A \in \tilde{\mathcal{A}}} \inf_{S \in \mathcal{S}} W(A, S) - \inf_{S \in \mathcal{S}} W(\tilde{A}^\varepsilon, S) \leq \varepsilon.$$

Theorem 14 (Maximin and Integrality Gaps in Welfare). *Suppose \mathcal{S} is a (δ, γ) uncertainty set with \mathcal{L}_1 diameter L . Let A^ε denote an ε -optimal discrete RAU solution, and \tilde{A}^ε denote an ε -optimal continuous RAU solution. Let A' denote the random variable that arises from applying the randomized rounding procedure ROUND to \tilde{A}^ε , and assume that ROUND preserves expectation, i.e., $\mathbb{E}_{A'}[A'] = \tilde{A}^\varepsilon$. Suppose also that the true affinity scores are S^* , and denote the optimal solution $A^* \doteq \arg \max_{A \in \mathcal{A}} W(A, S^*)$. The following then hold.*

1. *Maximin Gap:* $\mathbb{P}(W(A^*, S^*) - W(A^\varepsilon, S^*) > \varepsilon + \frac{2\gamma+L}{n}) < \delta$.
2. *Expected Regret of RRA:* $\mathbb{P}(W(A^*, S^*) - \mathbb{E}_{A'}[W(A', S^*)] > \varepsilon + \frac{2\gamma+L}{n}) < \delta$.
3. *Probabilistic Regret of RRA:* $\mathbb{P}(W(A^*, S^*) - W(A', S^*) > \frac{\varepsilon+(2\gamma+L)/n}{\delta'}) < \delta' + \delta$.

Table 1: Adversarial and average welfare (mean \pm standard deviation) for the naïve LP, FairFlow, PeerReview4All, FairSequence, and RRA methods on five ICLR conferences. Welfare is scaled by 100 for ease of comparison. Adversarial welfare is consistently highest (**bold**) using RRA, except for 2018, which is within one standard deviation.

Year	m	n	Adversarial USW \cdot 100					Average USW \cdot 100				
			LP	FF	PR4A	FS	RRA	LP	FF	PR4A	FS	RRA
2018	1657	546	17 \pm 3	7 \pm 3	17 \pm 3	16 \pm 3	16 \pm 3	179 \pm 2	134 \pm 12	177 \pm 2	177 \pm 2	160 \pm 4
2019	2620	851	22 \pm 2	12 \pm 2	22 \pm 2	22 \pm 2	27 \pm 3	184 \pm 1	139 \pm 9	184 \pm 1	183 \pm 1	161 \pm 3
2020	4123	1327	17 \pm 2	11 \pm 2	18 \pm 2	17 \pm 2	23 \pm 2	187 \pm 1	158 \pm 8	187 \pm 1	186 \pm 1	166 \pm 5
2021	4662	1557	23 \pm 2	18 \pm 2	23 \pm 2	23 \pm 2	33 \pm 3	192 \pm 1	177 \pm 2	192 \pm 1	191 \pm 1	174 \pm 6
2022	5023	1576	28 \pm 2	23 \pm 2	28 \pm 2	27 \pm 2	38 \pm 2	191 \pm 1	177 \pm 1	190 \pm 1	190 \pm 1	172 \pm 3

Since A' is the result of RRA, Theorem 14 directly bounds the expected and probabilistic regret of RRA. Note that the distribution for the probabilistic regret is over the randomness of the rounding procedure, while all bounds in Theorem 14 are probabilistic with respect to δ of the uncertainty set \mathcal{S} .

5 Experiments

We create uncertainty sets using five years of ICLR data, by constructing an asymmetric multivariate Gaussian and taking a confidence interval (as outlined by the example in (2)). We compare the adversarial and average-case performance of our approach to the adversarial and average-case performance of four commonly used baselines. We then demonstrate the importance of optimizing for the adversarial case. Code and data are publicly available on Github.⁶

5.1 Baseline Comparison

We examine the ability of our approach to robustly maximize utilitarian welfare in the case of *truncated* ellipsoidal uncertainty sets \mathcal{S} . This experiment is meant to mimic real-world conference scenarios, and thus we do not assume access to a ground truth affinity matrix. Consequently, we can only compare adversarial and average-case welfare for RRA vs. our baselines, but not true welfare.

We use the OpenReview API to collect all papers submitted (both accepted and rejected) to five recent iterations of ICLR (2018–2022). Following recent work, we use the pool of authors for each year as the reviewer pool, since we do not have access to the true reviewer identities for these conferences. The number of reviewers and papers for each conference year is shown in Table 1.

For each author in each year, we collect the multiset of keywords from papers the author submitted to ICLR in the current or previous years. We then follow a procedure similar to that of AAAI 2021 [18] to convert keywords into a mean vector $\vec{\mu} \in \mathbb{R}^{nm}$, and we also construct a covariance matrix $\Sigma \in \mathbb{R}_{\geq 0}^{nm \times nm}$ for paper-reviewer affinity scores. Vector \vec{p} is set so \vec{p}_i is an indicator for keyword i on paper p . Vector \vec{r} is initially set so \vec{r}_i is the number of times the keyword i appears on a paper written by that reviewer in this or previous years’ conferences. We then modify the values (but not the ordering) of \vec{r} such that the minimum non-zero value is 0.2, the maximum value is 1, and the remaining non-zero values are evenly-spaced between 0.2 and 1. Let $\vec{\lambda} \in \mathbb{R}^V$ be such that $\vec{\lambda}_i = (\frac{1}{2})^{i-1}$. *Sorted* represents the function that sorts values of a vector in decreasing order, $M_{\vec{p}}$ and $M_{\vec{r}}$ denote the number of non-zero entries in \vec{p} and \vec{r} respectively, and $Z = \sum_{i=1}^{M_{\vec{p}}} (\frac{1}{2})^{i-1}$. We set $\vec{\mu}_{pr} = \frac{\vec{\lambda} \cdot \text{Sorted}(\vec{p} \circ \vec{r})}{Z}$, $\Sigma_{pm+1, pm+1} = (M_{\vec{p}} M_{\vec{r}})^{-2}$, and all off-diagonal entries of Σ are 0. This procedure was chosen to roughly mirror the procedure used by AAAI 2021. For each year of ICLR, we set the uncertainty set \mathcal{S} for robust optimization to be the 95% confidence interval for the distribution $\mathcal{N}(\vec{\mu}, \Sigma)$ (as in (2)), intersected with the unit hypercube (Lemma 8).

For each year of ICLR, we sample without replacement 60% of the reviewers and 60% of the papers 100 times, to produce more data for statistical robustness of our experiments. We assume that all papers require 3 reviews and all reviewers can review up to 6 papers. There are no conflicts of interest. We then run RRA

⁶ <https://github.com/justinpayan/RAU>

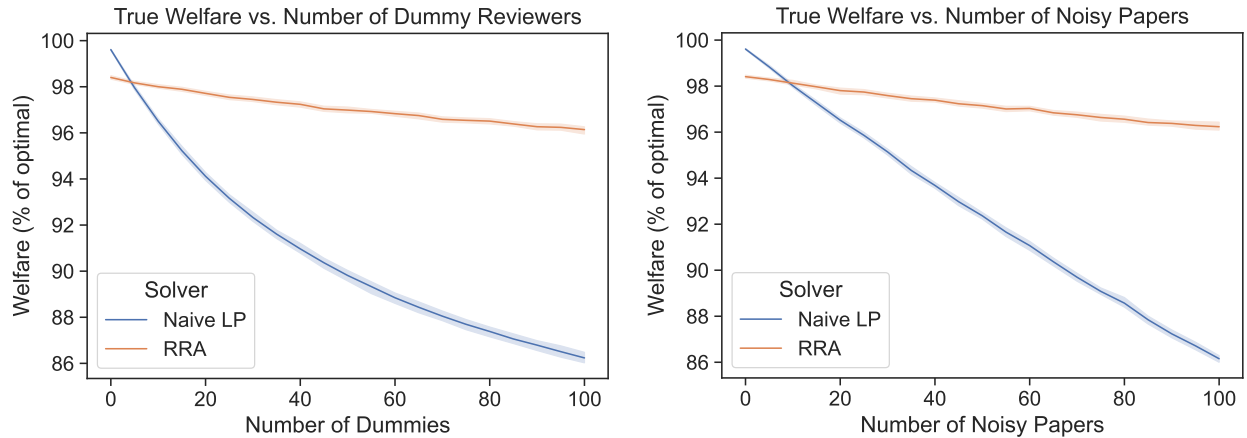


Fig. 1: Left: True welfare (as percentage of optimal, when S^* is known) of naïve LP approach vs. RRA on MIDL 2018 dataset with increasing number of “dummy” reviewers. Dummy reviewers have low true affinity for all papers, but higher noise in their estimated affinity scores. The true welfare (which is not known to the conference organizers) decreases sharply for the naïve LP formulation as we add more dummy reviewers, but RRA maintains high true welfare. Plot includes min, max, and average for 100 repetitions at each number of dummy reviewers. Right: True welfare on MIDL 2018 dataset with increasing number of noisy papers. Noisy papers have the 20th to 30th best reviewers’ affinities overestimated by 0.3. Once again, RRA is robust to this systematic overestimation, while LP degrades poorly.

using our calculated confidence region \mathcal{S} , and compare against the assignment given by $\arg \max_{A \in \mathcal{A}} W(A, \vec{\mu})$ (the “LP” solution); that is, the naïve solution that optimizes for the mean score. We also compare against three baselines commonly used for real conferences, FairFlow [16], PeerReview4All [33], and FairSequence [25]. The results of these experiments are shown in Table 1. All baselines lag far behind RRA in adversarial performance.

5.2 The Importance of Adversarial Analysis

To demonstrate the importance of optimizing for the adversarial case, we perform an experiment where we simulate the effect of adding many low-quality, high-variance reviewers to the dataset, and a similar experiment where we systematically overestimate some of the affinities for a subset of papers. These reviewers can often appear in modern conference reviewing. After being granted limited access to a recent, large AI conference ($m, n \approx 4000$), we manually inspected 100 random reviewers. 12 of these reviewers were PhD students, and 40 of them submitted 10 or fewer bids. PhD students with few papers will tend to have higher variance in true expertise relative to document-based similarity scores derived from their prior work, while reviewers with fewer bids will tend to have higher variance in true interest relative to their bids. Intuitively, the systematic overestimation of papers’ affinities can occur when paper authors submit a keyword that does not have the exact meaning they expected or is listed by reviewers in the wrong subcommunity. We believe that this occurs quite often, and leave a rigorous examination of this important question to future work.

We use the dataset of MIDL 2018, which is publicly available and has been used to validate many reviewer assignment algorithms in the past [16, 25, 33]. This dataset contains an $n \times m$ matrix of affinity scores that were used to assign reviewers to papers during the conference. We use this dataset for this experiment since it is small enough to run a large number of experiments with different settings ($n = 118$ papers and $m = 177$ reviewers). We assume the true affinity of the 177 original reviewers for each paper is equal to the affinity score present in the public dataset, but is *noisily estimated*. Thus, we assume that the conference organizers have access to estimated affinity scores which are equal to the original affinity scores plus normally distributed noise, $\mathcal{N}(0, .02)$. We also assume the conference organizers know that the estimation error is distributed according to $\mathcal{N}(0, .02)$ for each paper-reviewer pair. We also add a number of “dummy” reviewers to the dataset. For each dummy reviewer, we set the true affinity of that reviewer to be .1 for all papers, and sample

the estimated affinity from $\mathcal{N}(.1, .15)$. As with the original reviewers, we assume the conference organizers know the estimation error is distributed according to $\mathcal{N}(0, .15)$ for each paper-reviewer pair. This simple setup implies a multivariate Gaussian distribution over the true affinity scores, which are unknown to the conference organizers. We take a 95% confidence interval of this distribution and intersect it with the unit hypercube to define a truncated-ellipsoidal uncertainty set. We then assign reviewers using the naïve LP and RRA, and compare the *true*, but unknown, welfare for each approach.

For the setting with noisy papers, we take a subset of papers and identify the 20th to 30th (non-inclusive) ranked reviewers for each paper in decreasing order of affinity. For each paper, we add .3 to the estimated affinity for those 10 reviewers. We then assume that conference organizers estimate the standard deviation of these paper-reviewer pairs to be .15, while the remaining standard deviations are estimated at .02.

The results are shown in Figure 1. We report the welfare of naïve LP and RRA as a percentage of optimal, averaged over 100 runs per number of dummy reviewers/papers. We also report the minimum and maximum over all 100 runs for each setting. As we add more dummy reviewers or noisy papers, the true welfare of the naïve LP approach drops by up to roughly 15%, while RRA maintains high true welfare.

6 Extensions to RRA

Our approach is potentially much more general than the cases explored so far in this paper. This approach can incorporate modifications to the objective function, such as incorporation of additional soft constraints. We might also consider imposing penalties for unfair outcomes [16, 25, 33], or penalties to discourage manipulative bidding behavior [13–15]. An orthogonal extension is to directly modify the *welfare function*. However, it seems likely that the adversarial nature of our task precludes the use of more sophisticated welfare functions, such as the power-mean welfare, which would otherwise be amenable to maximization [6].

An orthogonal extension is to directly modify the *welfare function*, which can be used to tune the priorities of the assignment system, for instance more fairness (i.e., group-egalitarian [1], which is the *minimum average utility* of papers in any conference track or subject area) or to consider additional context, such as weighting the importance of conference track, journal track, and workshop track papers with a weighted USW function. We show in Table 1 that RRA suffers a reduction in average-case W in exchange for greatly improved worst-case W. It would be fairly straightforward to modify the objective to linearly interpolate between the worst-case and average-case W (when the expected value of S is known).

In general, fair welfare functions are concave, and thus the concave optimization methods of this work can be directly applied. However, we need to be able to compute and bound their gradients and modify RRA accordingly, and it may be difficult to solve the adversarial minimization problem for arbitrary welfare functions. Weighted utilitarian social welfare functions are a straightforward extension that cause essentially no difficulties, and group-egalitarian welfare can be adversarially optimized by separately minimizing for each group, and then taking the score matrix with the smallest group-average (a strategy adopted in the context of fair machine learning under adversarial uncertainty [7]). However, apart from such convenient special cases, it seems likely that the adversarial nature of our task precludes the use of more sophisticated welfare functions, such as the power-mean welfare, which would otherwise be amenable to maximization [6].

7 Conclusion

Uncertainty is endemic to reviewer assignment. Bias in bids due to partial information, prediction error in text similarity scores, and keyword terminology misalignment all result in inaccurate estimates of reviewers’ abilities. Furthermore, fundamental uncertainty about the quality of future reviews can be quantified, but never fully resolved. We treat uncertainty as a first-class citizen, formulating the problem of assigning reviewers under uncertainty as RAU and addressing it under broad conditions with RRA.

We define and explore the concept of uncertainty sets for reviewer assignment; our theory and examples demonstrate the flexibility of uncertainty sets to model real-world uncertainty-aware reviewer assignment workflows. Some special cases of uncertainty sets (singleton, hyperrectangular, or spherical) reduce to problems

without uncertainty, which can all be solved via linear programming. The general RAU problem is NP-hard, and our RRA algorithm provides approximate solutions with error guarantees dependent on the uncertainty set.

Because RRA optimizes against an uncertainty-aware adversary, we avoid assigning paper-reviewer pairs with high uncertainty. This method of accounting for uncertainty yields solutions that are more robust than optimizing with pointwise affinity score estimates. We hope this paper serves as a call to further investigate affinity score computation, ensuring affinities correlate with downstream review quality and incorporate the inherent uncertainty present in peer review.

Acknowledgments Thanks to OpenReview, Vignesh Viswanathan, Elita Lobo, Chang Zeng, Nihar Shah, and the reviewers at GAIW, IJCAI, and SAGT for feedback. Cyrus Cousins acknowledges postdoctoral fellowship support from the Center for Data Science at UMass Amherst. This work was performed using high performance computing equipment obtained under a grant from the Collaborative R&D Fund managed by the Massachusetts Technology Collaborative.

References

1. Aziz, H., Micha, E., Shah, N.: Group fairness in peer review. In: Proc. of the 22nd AAMAS (2023)
2. Bennett, G.: Probability inequalities for the sum of independent random variables. *Journal of the American Statistical Association* **57**(297), 33–45 (1962)
3. Budish, E., Che, Y.K., Kojima, F., Milgrom, P.: Implementing random assignments: A generalization of the Birkhoff-von Neumann theorem. In: Cowles Summer Conference (2009)
4. Cai, T.T., Zhou, W.X.: Matrix completion via max-norm constrained optimization. *Electronic Journal of Statistics* **10**(1), 1493–1525 (2016)
5. Charlin, L., Zemel, R.: The Toronto paper matching system: an automated paper-reviewer assignment system. In: Proc. of the 2013 ICML Workshop on Peer Reviewing and Publishing Models (2013)
6. Cousins, C.: An axiomatic theory of provably-fair welfare-centric machine learning. In: Proc. of the 35th NeurIPS. pp. 16610–16621 (2021)
7. Dong, E., Cousins, C.: Decentering imputation: Fair learning at the margins of demographics. In: Queer in AI Workshop @ ICML (2022)
8. Fang, E.X., Liu, H., Toh, K.C., Zhou, W.X.: Max-norm optimization for robust matrix recovery. *Mathematical Programming* **167**(1), 5–35 (2018)
9. Fiez, T., Shah, N., Ratliff, L.: A SUPER* algorithm to optimize paper bidding in peer review. In: Proc. of the 36th UAI (2020)
10. Gandhi, R., Khuller, S., Parthasarathy, S., Srinivasan, A.: Dependent rounding and its applications to approximation algorithms. *Journal of the ACM (JACM)* **53**(3), 324–360 (2006)
11. Garg, N., Kavitha, T., Kumar, A., Mehlhorn, K., Mestre, J.: Assigning papers to referees. *Algorithmica* **58**(1), 119–136 (2010)
12. Hoeffding, W.: Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* **58**(301), 13–30 (1963)
13. Jecmen, S., Shah, N.B., Fang, F., Conitzer, V.: Tradeoffs in preventing manipulation in paper bidding for reviewer assignment. In: Proc. of the 2022 ICLR Workshop on ML Evaluation Standards (2022)
14. Jecmen, S., Yoon, M., Conitzer, V., Shah, N.B., Fang, F.: A dataset on malicious paper bidding in peer review. In: Proc. of the 2023 WWW (2023)
15. Jecmen, S., Zhang, H., Liu, R., Shah, N., Conitzer, V., Fang, F.: Mitigating manipulation in peer review via randomized reviewer assignments. In: Proc. of the 34th NeurIPS (2020)
16. Kobren, A., Saha, B., McCallum, A.: Paper matching with local fairness constraints. In: Proc. of the 25th KDD (2019)
17. Lee, J.D., Recht, B., Srebro, N., Tropp, J., Salakhutdinov, R.R.: Practical large-scale optimization for max-norm regularization. Proc. of the 23rd NeurIPS (2010)
18. Leyton-Brown, K., Nandwani, Y., Zarkoob, H., Cameron, C., Newman, N., Raghu, D., et al.: Matching papers and reviewers at large conferences. arXiv preprint arXiv:2202.12273 (2022)
19. Mazzetto, A., Cousins, C., Sam, D., Bach, S.H., Upfal, E.: Adversarial multi class learning under weak supervision with performance guarantees. In: International Conference on Machine Learning. PMLR (2021)
20. McDiarmid, C.: On the method of bounded differences. *Surveys in combinatorics* **141**(1), 148–188 (1989)

21. Meir, R., Lang, J., Lesca, J., Mattei, N., Kaminsky, N.: A market-inspired bidding scheme for peer review paper assignment. In: Proc. of the 35th AAAI (2021)
22. Neubig, G., the ACL Reviewing Committee: ACL rolling review proposal (2020), aclweb.org/adminwiki/index.php?title=ACL_Rolling_Review_Proposal
23. Neubig, G., Wieting, J., McCarthy, A., Stent, A., Schluter, N., Cohn, T.: ACL reviewer matching code. github.com/acl-org/reviewer-paper-matching (2020)
24. OpenReview: How to do automatic assignments (2023), docs.openreview.net/how-to-guides/paper-matching-and-assignment/how-to-do-automatic-assignments
25. Payan, J., Zick, Y.: I will have order! Optimizing orders for fair reviewer assignment. In: Proc. of the 31st IJCAI (2022)
26. Rogers, A., Augenstein, I.: What can we do to improve peer review in NLP? In: Proc. of the 2020 EMNLP (2020)
27. Rozenzweig, I., Meir, R., Mattei, N., Amir, O.: Mitigating skewed bidding for conference paper assignment. In: Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (2023)
28. Saveski, M., Jecmen, S., Shah, N.B., Ugander, J.: Counterfactual evaluation of peer-review assignment policies. arXiv preprint [arXiv:2305.17339](https://arxiv.org/abs/2305.17339) (2023)
29. Schroter, S., Tite, L., Hutchings, A., Black, N.: Differences in review quality and recommendations for publication between peer reviewers suggested by authors or by editors. *Jama* **295**(3), 314–317 (2006)
30. Shah, N.B.: Principled methods to improve peer review. Tech. rep., Carnegie Mellon University (2019)
31. Shah, N.B.: Challenges, experiments, and computational solutions in peer review. *Communications of the ACM* **65**(6), 76–87 (2022)
32. Shor, N., Kiwiel, K.C., Ruszcayński, A.: Minimization methods for non-differentiable functions. Springer-Verlag (1985)
33. Stelmakh, I., Shah, N.B., Singh, A.: PeerReview4All: Fair and accurate reviewer assignment in peer review. In: Proc. of the 30th ALT (2019)
34. Stelmakh, I., Wieting, J., Neubig, G., Shah, N.B.: A gold standard dataset for the reviewer assignment problem. arXiv preprint [arXiv:2303.16750](https://arxiv.org/abs/2303.16750) (2023)
35. Taylor, C.J.: On the optimal assignment of conference papers to reviewers. Tech. rep., University of Pennsylvania (2008)
36. Van Rooyen, S., Black, N., Godlee, F.: Development of the review quality instrument (RQI) for assessing peer reviews of manuscripts. *Journal of clinical epidemiology* **52**(7), 625–629 (1999)

A Missing Proofs

We present proofs for all theoretical results stated in the main paper body without proof. Appendix A.1 proves the relation between true welfare and maximin welfare. Appendix A.2 proves the results enabling construction of uncertainty sets from estimators as well as by composing or extending other uncertainty sets. Finally, Appendix A.3 proves NP-hardness of RAU and contains constructive proofs for solving RAU exactly in special cases (hyperrectangles and spheres) as well as approximately in the case of convex uncertainty sets with polynomial-time adversaries.

A.1 Relation between True Welfare and Maximin Welfare

We prove the proposition that states that the minimum welfare of an assignment A over a (δ, γ) uncertainty set \mathcal{S} relates additively to the true welfare of A . The tightness of the approximation depends on the \mathcal{L}_1 diameter of the set \mathcal{S} and the additive \mathcal{L}_1 error γ allowed for the uncertainty set.

Proposition 2 (Relating True and Worst-Case Welfare over \mathcal{S}). *Suppose \mathcal{S} is a (δ, γ) uncertainty set with $\|S - S'\|_1 \leq L$ for all $S, S' \in \mathcal{S}$, and the true affinity score matrix is labeled S^* . Consider any assignment $A \in \mathcal{A}$. Then with probability at least $1 - \delta$,*

$$W(A, S^*) - \frac{L + \gamma}{n} \leq \inf_{S \in \mathcal{S}} W(A, S) \leq W(A, S^*) + \frac{\gamma}{n}.$$

Proof. For the right hand side, note that with probability at least $1 - \delta$ there exists some $S_\gamma^* \in \mathcal{S}$ with $\|S_\gamma^* - S^*\|_1 \leq \gamma$. Let $S' = \arg \min_{S \in \mathcal{S}} W(A, S)$. By definition, $W(A, S') \leq W(A, S_\gamma^*)$. If $W(A, S_\gamma^*) \leq W(A, S^*)$, we have the desired inequality. Otherwise, we can apply the fact that $W(A, S_\gamma^*) - W(A, S^*) = W(A, S_\gamma^* - S^*) \leq \frac{1}{n} \|S_\gamma^* - S^*\|_1 \leq \frac{1}{n} \gamma$, where the second-to-last inequality holds since every entry of A is in $\{0, 1\}$.

To derive the left hand side, we will aim to bound $W(A, S^*) - W(A, S')$ where $S' = \arg \min_{S \in \mathcal{S}} W(A, S)$. Again with probability at least $1 - \delta$ there exists some $S_\gamma^* \in \mathcal{S}$ with $\|S_\gamma^* - S^*\|_1 \leq \gamma$. So $W(A, S^*) - W(A, S_\gamma^*) = W(A, S^* - S_\gamma^*) \leq \frac{1}{n} \|S^* - S_\gamma^*\|_1 \leq \frac{\gamma}{n}$. Applying similar logic, we can bound $W(A, S_\gamma^*) - W(A, S') = W(A, S_\gamma^* - S') \leq \frac{1}{n} \|S_\gamma^* - S'\|_1 \leq \frac{L}{n}$. \square

A.2 Constructing Uncertainty Sets

We first prove our two results showing how ellipsoidal uncertainty sets can be constructed using inductive or transductive tail bounds for estimated affinity scores.

Theorem 6 (Ellipsoidal Uncertainty Sets from Inductive Predictors). *Let \mathcal{D}' be a probability distribution over paper-reviewer pairs, and let \mathcal{D}^P and \mathcal{D}^R be distributions over papers and reviewers, respectively. Assume that P and R were drawn conditionally independently without replacement from \mathcal{D}^P and \mathcal{D}^R , respectively.⁷ Suppose we sample T paper-reviewer pairs $\{(p_i, r_i)\}_{i=1}^T$ conditionally independently without replacement from \mathcal{D}' , and these paper-reviewer pairs have true and estimated affinity scores $\{f^*(p_i, r_i)\}_{i=1}^T$ and $\{\hat{f}(p_i, r_i)\}_{i=1}^T$, respectively.*

Let

$$\alpha(p, r) \doteq \frac{\mathbb{P}_{\mathbf{p} \sim \mathcal{D}^P}(\mathbf{p} = p) \mathbb{P}_{\mathbf{r} \sim \mathcal{D}^R}(\mathbf{r} = r)}{\mathbb{P}_{(\mathbf{p}, \mathbf{r}) \sim \mathcal{D}'}((\mathbf{p}, \mathbf{r}) = (p, r))} \quad \& \quad \alpha_{\min} \doteq \inf_{p \in P, r \in R} \alpha(p, r)$$

denote (1) the probability ratio of sampling p from \mathcal{D}^P and r from \mathcal{D}^R to sampling (p, r) from \mathcal{D}' , and (2) the infimum probability ratio, respectively. Now construct the ellipsoid matrix $\Sigma \in \mathbb{R}^{nm \times nm}$ as the diagonal matrix such that $\Sigma_{pm+r, pm+r} = \alpha(p, r)$ for all $p \in P, r \in R$.

⁷ Abusing notation slightly, we index the papers $p \in P$ and reviewers $r \in R$ so that p can represent either a paper in P or an integer between 1 and n .

Then for any $\delta \in (0, 1)$, the ellipsoid

$$\mathcal{S} \doteq \left\{ S \in \mathbb{R}^{nm} \left| \underbrace{\frac{1}{nm}(S - \hat{S})^\top \Sigma^{-1}(S - \hat{S}) \leq \frac{1}{T} \sum_{i=1}^T (f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2}_{=\hat{\xi}} + \underbrace{\sqrt{\left(\frac{1}{T} + \frac{n+m}{nm\alpha_{\min}^2}\right) \frac{\ln \frac{1}{\delta}}{2}}}_{=\eta} \right. \right\}$$

is a $(\delta, 0)$ uncertainty set, where $\hat{\xi}$ denotes the empirical square error of our estimated scores, and η denotes the excess error bound due to sampling.

Proof. We aim to show the ellipsoid \mathcal{S} is a $(\delta, 0)$ uncertainty set. In other words, it must hold that with probability $1 - \delta$, the true affinity scores S^* satisfy

$$\frac{1}{nm}(S^* - \hat{S})^\top \Sigma^{-1}(S^* - \hat{S}) \leq \underbrace{\frac{1}{T} \sum_{i=1}^T (f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2}_{=\hat{\xi}} + \underbrace{\sqrt{\left(\frac{1}{T} + \frac{n+m}{nm\alpha_{\min}^2}\right) \frac{\ln \frac{1}{\delta}}{2}}}_{=\eta}.$$

We will give a probabilistic bound on the difference

$$Z \doteq \frac{1}{nm}(S^* - \hat{S})^\top \Sigma^{-1}(S^* - \hat{S}) - \hat{\xi} = \frac{1}{nm} \sum_{p=1}^n \sum_{r=1}^m \frac{(S_{p,r}^* - \hat{S}_{p,r})^2}{\alpha(p, r)} - \frac{1}{T} \sum_{i=1}^T (f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2.$$

We first note that $\mathbb{E}[Z] = 0$. We assume that our papers $p \sim \mathcal{D}^{\mathcal{P}}$, our reviewers $r \sim \mathcal{D}^{\mathcal{R}}$, and our samples (which are reviewer-paper pairs) are drawn as $(p_i, r_i) \sim \mathcal{D}'$. For simplicity of notation, we will write $\{(p_i, r_i)\}_{i=1}^T \sim \mathcal{D}'$ to denote the set of all paper-reviewer pairs $(p_i, r_i) \sim \mathcal{D}'$. Also recall that $\alpha(p, r) = \frac{\mathbb{P}_{\mathbf{p} \sim \mathcal{D}^{\mathcal{P}}}(\mathbf{p}=p) \mathbb{P}_{\mathbf{r} \sim \mathcal{D}^{\mathcal{R}}}(\mathbf{r}=r)}{\mathbb{P}_{(\mathbf{p}, \mathbf{r}) \sim \mathcal{D}' }((\mathbf{p}, \mathbf{r})=(p, r))}$. We have that

$$\begin{aligned} \mathbb{E}_{\substack{P \sim \mathcal{D}^{\mathcal{P}}, R \sim \mathcal{D}^{\mathcal{R}} \\ \{(p_i, r_i)\}_{i=1}^T \sim \mathcal{D}'}} [Z] &= \mathbb{E}_{\substack{P \sim \mathcal{D}^{\mathcal{P}}, R \sim \mathcal{D}^{\mathcal{R}} \\ \{(p_i, r_i)\}_{i=1}^T \sim \mathcal{D}'}} \left[\frac{1}{nm} \sum_{p=1}^n \sum_{r=1}^m \frac{(S_{p,r}^* - \hat{S}_{p,r})^2}{\alpha(p, r)} - \frac{1}{T} \sum_{i=1}^T (f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2 \right] \\ &= \mathbb{E}_{P \sim \mathcal{D}^{\mathcal{P}}, R \sim \mathcal{D}^{\mathcal{R}}} \left[\frac{1}{nm} \sum_{p=1}^n \sum_{r=1}^m \frac{(S_{p,r}^* - \hat{S}_{p,r})^2}{\alpha(p, r)} \right] - \mathbb{E}_{\{(p_i, r_i)\}_{i=1}^T \sim \mathcal{D}'} \left[\frac{1}{T} \sum_{i=1}^T (f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2 \right]. \end{aligned}$$

Following the standard argument for importance sampling, it is also clear that

$$\begin{aligned} \mathbb{E}_{P \sim \mathcal{D}^{\mathcal{P}}, R \sim \mathcal{D}^{\mathcal{R}}} \left[\frac{1}{nm} \sum_{p=1}^n \sum_{r=1}^m \frac{(S_{p,r}^* - \hat{S}_{p,r})^2}{\alpha(p, r)} \right] &= \frac{1}{nm} \sum_{p=1}^n \sum_{r=1}^m \mathbb{E}_{P \sim \mathcal{D}^{\mathcal{P}}, R \sim \mathcal{D}^{\mathcal{R}}} \left[\frac{(S_{p,r}^* - \hat{S}_{p,r})^2}{\alpha(p, r)} \right] \\ &= \frac{1}{nm} \sum_{i=1}^{nm} \mathbb{E}_{(p_i, r_i) \sim \mathcal{D}'} \left[(f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2 \right] \\ &= \frac{1}{T} \sum_{i=1}^T \mathbb{E}_{(p_i, r_i) \sim \mathcal{D}'} \left[(f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2 \right] \\ &= \mathbb{E}_{\{(p_i, r_i)\}_{i=1}^T \sim \mathcal{D}'} \left[\frac{1}{T} \sum_{i=1}^T (f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2 \right]. \end{aligned}$$

In addition, Z is a function of a system of negatively-dependent (due to sampling without replacement) random variables comprising T auxiliary terms, m reviewers, and n papers. Modifying any of the T auxiliary terms can result in at most $\frac{1}{T}$ change in Z , each of the n papers has impact at most $\frac{1}{n\alpha_{\min}}$, and each of m reviewers has impact at most $\frac{1}{m\alpha_{\min}}$. The sum of square bounded differences is thus $\frac{1}{T} + \frac{1}{n\alpha_{\min}^2} + \frac{1}{m\alpha_{\min}^2}$.

Consequently, by McDiarmid's inequality [20], it holds

$$\mathbb{P} \left(Z \geq \sqrt{\left(\frac{1}{T} + \frac{1}{n\alpha_{\min}^2} + \frac{1}{m\alpha_{\min}^2} \right) \frac{\ln \frac{1}{\delta}}{2}} \right) \leq \delta .$$

□

Theorem 7 (Ellipsoidal Uncertainty Sets from Transductive Predictors). *Suppose we sample T paper-reviewer pairs $\{(p_i, r_i)\}_{i=1}^T$ conditionally independently without replacement from \mathcal{D}' , and these paper-reviewer pairs have true and estimated affinity scores $\{f^*(p_i, r_i)\}_{i=1}^T$ and $\{\hat{f}(p_i, r_i)\}_{i=1}^T$, respectively. Let*

$$\alpha(p, r) \doteq \frac{(nm)^{-1}}{\mathbb{P}_{(\mathbf{p}, \mathbf{r}) \sim \mathcal{D}'}((\mathbf{p}, \mathbf{r}) = (p, r))}$$

denote the probability ratio between sampling (p, r) uniformly at random and sampling (p, r) from \mathcal{D}' , and construct the ellipsoid matrix $\Sigma \in \mathbb{R}^{nm \times nm}$ as the diagonal matrix such that $\Sigma_{pm+r, pm+r} = \alpha(p, r)$ for all $p \in P, r \in R$.

Then for any $\delta \in (0, 1)$, the ellipsoid

$$\mathcal{S} \doteq \left\{ S \in \mathbb{R}^{nm} \left| \frac{1}{nm} (S - \hat{S})^\top \Sigma^{-1} (S - \hat{S}) \leq \underbrace{\frac{1}{T} \sum_{i=1}^T (f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2}_{=\hat{\xi}} + \underbrace{\sqrt{\frac{\ln \frac{1}{\delta}}{2T}}}_{=\eta} \right. \right\}$$

is a $(\delta, 0)$ uncertainty set, where $\hat{\xi}$ denotes the empirical square error of our estimated scores, and η denotes the excess error bound due to sampling.

Proof. Proof of result follows similarly to that of Theorem 6. Here, reviewers and papers are fixed. We will bound the deviation from the mean for the random variable

$$Z \doteq \frac{1}{nm} \sum_{p=1}^n \sum_{r=1}^m \frac{(S_{p,r} - \hat{S}_{p,r})^2}{\alpha(p, r)} - \frac{1}{T} \sum_{i=1}^T (f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2 .$$

We show that $\mathbb{E}[Z] = 0$, as

$$\begin{aligned} \mathbb{E}_{\{(p_i, r_i)\}_{i=1}^T \sim \mathcal{D}'}[Z] &= \mathbb{E}_{\{(p_i, r_i)\}_{i=1}^T \sim \mathcal{D}'} \left[\frac{1}{nm} \sum_{p=1}^n \sum_{r=1}^m \frac{(S_{p,r} - \hat{S}_{p,r})^2}{\alpha(p, r)} - \frac{1}{T} \sum_{i=1}^T (f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2 \right] \\ &= \frac{1}{nm} \sum_{p=1}^n \sum_{r=1}^m \frac{(S_{p,r} - \hat{S}_{p,r})^2}{\alpha(p, r)} - \frac{1}{T} \sum_{i=1}^T \mathbb{E}_{\{(p_i, r_i)\}_{i=1}^T \sim \mathcal{D}'} \left[(f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2 \right] \\ &= \frac{1}{nm} \sum_{p=1}^n \sum_{r=1}^m \frac{(S_{p,r} - \hat{S}_{p,r})^2}{\alpha(p, r)} - \frac{1}{nm} \sum_{i=1}^{nm} \mathbb{E}_{\{(p_i, r_i)\}_{i=1}^{nm} \sim \text{Unif}(P \times R)} \left[\frac{(f^*(p_i, r_i) - \hat{f}(p_i, r_i))^2}{\alpha(p_i, r_i)} \right] \\ &= \frac{1}{nm} \sum_{p=1}^n \sum_{r=1}^m \frac{(S_{p,r} - \hat{S}_{p,r})^2}{\alpha(p, r)} - \frac{1}{nm} \sum_{p=1}^n \sum_{r=1}^m \frac{(S_{p,r} - \hat{S}_{p,r})^2}{\alpha(p, r)} . \end{aligned}$$

To apply the final bound, we need only consider the T independent auxiliary terms, each of bounded difference $\frac{1}{T}$. In this case, we have the bounded difference term $\frac{1}{T}$ for each of T auxiliary terms. This yields, via the Hoeffding [12] or McDiarmid [20] inequalities, the final result

$$\mathbb{P} \left(Z \geq \sqrt{\frac{\ln \frac{1}{\delta}}{2T}} \right) \leq \delta .$$

□

Now that we have shown how to construct uncertainty sets from affinity score estimators, we finish by showing how to compose uncertainty sets. We proved Lemma 8 in Section 3.3, showing that uncertainty sets

can be intersected to form new uncertainty sets. Lemma 8 can be combined with Lemma 9 to create many types of uncertainty sets from simpler uncertainty sets.

Lemma 9 (\mathcal{L}_1 Error Terms). *If \mathcal{S} is a (δ, γ) uncertainty set, then for any $\eta \in [0, \gamma]$, it holds that the Minkowski sum*

$$\mathcal{S}' \doteq \mathcal{S} + \{\vec{s} \in \mathbb{R}^{nm} \mid \|\vec{s}\|_1 \leq \eta\} = \{\vec{x} + \vec{s} \mid \vec{x} \in \mathcal{S}, \|\vec{s}\|_1 \leq \eta\}$$

is a $(\delta, \gamma - \eta)$ uncertainty set.

Proof. Denote the true affinity scores as S^* . If $S^* \in \mathcal{S}'$, then we have $\|S - S^*\|_1 = 0 \leq \gamma - \eta$. We argue the case when $S^* \notin \mathcal{S}'$. By Definition 1, with probability at least $1 - \delta$, there is some $S \in \mathcal{S}$ such that $\|S - S^*\|_1 \leq \gamma$. We can assume that $\eta \leq \|S^* - S\|_1$, since if not, then $S^* \in \mathcal{S}'$. Consider $S' = S + \eta \frac{S^* - S}{\|S^* - S\|_1}$. To show $S' \in \mathcal{S}'$, we must show that $\left\| \eta \frac{S^* - S}{\|S^* - S\|_1} \right\|_1 \leq \eta$. This can be shown easily, as $\left\| \eta \frac{S^* - S}{\|S^* - S\|_1} \right\|_1 = \frac{\eta}{\|S^* - S\|_1} \|S^* - S\|_1 = \eta$. In addition,

$$\begin{aligned} \|S^* - S'\|_1 &= \left\| S^* - \left(S + \eta \frac{S^* - S}{\|S^* - S\|_1} \right) \right\|_1 \\ &= \left(1 - \frac{\eta}{\|S^* - S\|_1} \right) \|S^* - S\|_1 \\ &= \|S^* - S\|_1 - \eta \\ &\leq \gamma - \eta, \end{aligned}$$

where the final inequality holds with probability at least $1 - \delta$. Therefore, with probability at least $1 - \delta$, \mathcal{S}' contains a point S' within $\gamma - \eta$ \mathcal{L}_1 distance from S^* , and \mathcal{S}' is a $(\delta, \gamma - \eta)$ uncertainty set. \square

A.3 NP-hardness and Algorithms for Solving RAU

We first show that RAU is NP-hard, even for the case of bounded polytopes formed by intersections of polynomially many halfspaces.

Theorem 10 (Hardness of RAU). *RAU is NP-hard over a convex uncertainty set \mathcal{S} , even for \mathcal{S} with a polynomial-time adversary. In particular, RAU remains NP-hard even when \mathcal{S} is restricted to bounded polytopes formed by intersections of polynomially many halfspaces.*

Proof. We reduce from the problem of maximizing egalitarian welfare under the reviewer assignment constraints, which is known to be NP-hard [11, 33]. In the maximal egalitarian welfare problem for reviewer assignment, we have the same feasible set of assignments \mathcal{A} , as well as a fixed score matrix S . The goal is to find an assignment A maximizing the minimum total score of any paper, or equivalently $\min_{p \in P} \sum_{r \in R} A_{p,r} S_{p,r}$.

Given an instance of max egalitarian welfare (P, R, \mathcal{A}, S) , we now construct an instance of RAU over a convex uncertainty set \mathcal{S} . The set of n papers P , m reviewers R , and valid assignments \mathcal{A} remain the same. To construct an uncertainty set \mathcal{S} , first consider the set $\mathcal{S}' = \{S^{(1)}, \dots, S^{(n)}\}$, where $S^{(p)}$ is defined so that $S_{i,j}^{(p)} = S_{i,j}$ when $i = p$ and $S_{i,j}^{(p)} = 0$ when $i \neq p$. Let \mathcal{S} be the convex hull of \mathcal{S}' . In other words, \mathcal{S} is the set

$$\mathcal{S} \doteq \left\{ X \in [0, 1]^{n \times m} \mid X = \sum_{i=1}^n \alpha_i S^{(i)}, \sum_{i=1}^n \alpha_i = 1 \right\}.$$

Since \mathcal{S} is a convex set, $(P, R, \mathcal{A}, \mathcal{S})$ is a RAU instance over a convex uncertainty set. Furthermore, \mathcal{S} is a bounded polytope formed by the intersection of polynomially many halfspaces.

For any $A \in \mathcal{A}$, we have that

$$S_{\min} \doteq \arg \min_{S \in \mathcal{S}} nW(A, S) = \arg \min_{S \in \mathcal{S}} W(A, S).$$

We will show that $S_{\min} \in \mathcal{S}'$, completing the proof. Fix A and consider some element $X \in \mathcal{S}$, where $X = \sum_{i=1}^n \alpha_i S^{(i)}$ and (w.l.o.g.) $\alpha_1, \alpha_2 > 0$. Assume also w.l.o.g. that $\sum_{r \in R} A_{1,r} S_{1,r} \leq \sum_{r \in R} A_{2,r} S_{2,r}$.

Consider the alternative $X' \in \mathcal{S}$ with $X' = (\alpha_1 + \alpha_2)S^{(1)} + \sum_{i=3}^n \alpha_i S^{(i)}$. Then

$$\begin{aligned} nW(A, X) &= \alpha_1 \sum_{r \in R} A_{1,r} S_{1,r} + \alpha_2 \sum_{r \in R} A_{2,r} S_{2,r} + \sum_{i=3}^n \alpha_i \sum_{r \in R} A_{i,r} S_{i,r} \\ &\geq (\alpha_1 + \alpha_2) \sum_{r \in R} A_{1,r} S_{1,r} + \sum_{i=3}^n \alpha_i \sum_{r \in R} A_{i,r} S_{i,r} \\ &= nW(A, X') . \end{aligned}$$

Since we can always decrease $W(A, S)$ by restricting the support of the score matrices (i.e., setting some $\alpha_i \neq 0$ to be 0), the minimal W is reached for some $S \in \mathcal{S}'$. Therefore, since for any given A , we have that

$$\min_{S \in \mathcal{S}'} nW(A, S) = \min_{p \in P} \sum_{r \in R} A_{p,r} S_{p,r} ,$$

the maximizer for the RAU problem and the maximizer for the egalitarian welfare problem are equivalent. In other words, if we can efficiently compute a robust solution to RAU, we can solve the egalitarian welfare problem. Therefore RAU is NP-hard. \square

Although RAU is NP-hard for even fairly simple cases, there are certain special cases in which RAU reduces to a problem solvable in polynomial time. The first case is that of hyperrectangular uncertainty sets.

Theorem 3 (RAU under Hyperrectangular Uncertainty). *When the uncertainty set is an axis-aligned hyperrectangular region \mathcal{S}_{\square} , then*

$$\arg \max_{A \in \mathcal{A}} \inf_{S \in \mathcal{S}_{\square}} W(A, S) = \arg \max_{A \in \mathcal{A}} W(A, \underline{S}) .$$

Thus RAU maximizes W for \underline{S} , which requires polynomial time via LP reduction.

Proof. For any assignment $A \in \mathcal{A}$, $\inf_{S \in \mathcal{S}_{\square}} W(A, S) = W(A, \underline{S})$, i.e., the lowest welfare is achieved (possibly non-uniquely) if we assume that all scores are the worst they can be under the rectangular constraints. To see why, consider any matrix $X \in \mathcal{S}_{\square}$ with $X_{p,r} > \underline{S}_{p,r}$ for some p and r . Consider X' where $X'_{p,r} = \underline{S}_{p,r}$ and $X'_{i,j} = X_{i,j}$ for $i \neq p$ or $j \neq r$. Either $A_{p,r} = 0$ and hence $W(A, X) = W(A, X')$ or $A_{p,r} = 1$ and thus $W(A, X) > W(A, X')$.

We also observe that the problem $\arg \max_{A \in \mathcal{A}} W(A, \underline{S})$ is an instance of RAU with known S . This problem is solvable in polynomial time via linear programming, as discussed at the beginning of Section 3. \square

RAU is also solvable in polynomial time for spherical uncertainty sets, and the proof provides a bit more insight than the proof of Theorem 3.

Theorem 4 (RAU under Spherical Uncertainty). *When \mathcal{S} is a sphere,*

$$\arg \max_{A \in \mathcal{A}} \inf_{S \in B_{\varepsilon}(S^0)} W(A, S) = \arg \max_{A \in \mathcal{A}} W(A, S^0) ,$$

which can be computed in polynomial time. Furthermore, for all $A \in \mathcal{A}$,

$$|W(A, S^0) - W(A, S^*)| \leq \frac{\varepsilon \sqrt{K}}{n} .$$

Proof. Let $\mathbf{0}$ denote the $n \times m$ matrix of all 0 values. For any $A \in \mathcal{A}$, it holds that

$$\begin{aligned}
\inf_{S \in B_\varepsilon(S^0)} W(A, S) &= \min_{X \in B_\varepsilon(\mathbf{0})} W(A, S^0 + X) \\
&= \min_{X \in B_\varepsilon(\mathbf{0})} W(A, X) + W(A, S^0) && \text{LINEARITY OF } W \\
&= W\left(A, S^0 - A \frac{\varepsilon}{\|A\|_F}\right) && \text{SEE BELOW} \\
&= W(A, S^0) - W\left(A, A \frac{\varepsilon}{\|A\|_F}\right) && \text{LINEARITY OF } W \\
&= W(A, S^0) - \frac{1}{n} \|A\|_F \varepsilon \\
&= W(A, S^0) - \frac{1}{n} \sqrt{K} \varepsilon .
\end{aligned}$$

Because $W(A, S^0)$ does not depend on X , we just need to find $\arg \min_{X \in B_\varepsilon(\mathbf{0})} W(A, X)$. Because $\|A\|_F$ is fixed (and the argmin occurs at $\|X\|_F = \varepsilon$) the argmin must be $X = -A \frac{\varepsilon}{\|A\|_F}$.

Again, \sqrt{K} , ε , and n are all fixed. So we see that our function has the same argmax as $W(A, S^0)$. The problem $\arg \max_{A \in \mathcal{A}} W(A, S^0)$ is an instance of RAU with known S , which is solvable in polynomial time.

We can quantify the worst-case welfare loss due to spherical uncertainty with radius ε as

$$W(A, S^0) - W(A, S^*) \leq \frac{1}{n} \sqrt{K} \varepsilon .$$

A similar argument to the above shows that

$$\max_{S \in B_\varepsilon(S^0)} W(A, S) = W(A, S^0) + \frac{1}{n} \sqrt{K} \varepsilon ,$$

and thus we also have

$$W(A, S^*) - W(A, S^0) \leq \frac{1}{n} \sqrt{K} \varepsilon .$$

□

Finally, we will prove Proposition 11 from Section 4, which enables approximate solutions to RAU when the uncertainty set \mathcal{S} admits a polynomial-time minimizing adversary. We use the following result from [32].

Theorem 15. *Suppose that $w : \mathbb{R}^n \rightarrow \mathbb{R}$ is a convex function with minimizer w^* and optimal set $X^* = \{\vec{x} \in \mathbb{R}^n : w(\vec{x}) = w^*\}$. Let \vec{x}_0 denote the starting point for subgradient descent. Suppose we have an upper bound $\rho \geq \max_{\vec{x} \in X^*} \|\vec{x}_0 - \vec{x}\|_2$. Suppose we also have an upper bound λ for the Euclidean norm of the subgradient at each step of the subgradient descent algorithm. After T steps of subgradient descent with constant step size α , if \hat{w} represents the best function value found after T steps, then*

$$\hat{w} - w^* \leq \frac{\rho^2 + \lambda^2 \alpha^2 T}{2\alpha T} .$$

Proposition 11 (Supergradient Ascent Efficiency). *Let λ denote an upper bound on the \mathcal{L}_2 norm of the supergradient elements $\nabla_A W(A, S)$ used in RRA. The supergradient ascent component of RRA converges to within ε of the maximin optimal continuous assignment \tilde{A}^{MM} in $\lceil 2K(\frac{\lambda}{\varepsilon})^2 \rceil$ iterations. RRA runs in time $\mathcal{O}(2KC(\frac{\lambda}{\varepsilon})^2)$, where C is the time cost of one adversary and projection step.*

Proof. First, we show that $\rho \leq \sqrt{2K}$ (recall ρ is defined as the upper bound $\rho \geq \max_{\vec{x} \in X^*} \|\vec{x}_0 - \vec{x}\|_2$). Because the initial assignment $A^{(0)}$ and the optimal assignment A^* both lie in $\tilde{\mathcal{A}}$, they both must satisfy the constraint that for all papers p , $\sum_{r \in R} A_{p,r} = k_p$. This implies that even if there is no overlap in the two assignments, they will differ on at most $2K$ entries. Taking the Euclidean norm over the difference of the two assignments in the worst case gives the bound on ρ .

Our function is concave, and we are maximizing over the input space, so we can apply Theorem 15 to minimize its negative. From Theorem 15 we know that after T iterations of subgradient ascent with step size

α , the error $\varepsilon = \inf_{S \in \mathcal{S}} W(A^*, S) - \inf_{S \in \mathcal{S}} W(\hat{A}, S)$ satisfies

$$\varepsilon \leq \frac{\rho^2 + \lambda^2 \alpha^2 T}{2\alpha T}.$$

The right-hand side is minimized at $\alpha = \frac{\rho}{\lambda\sqrt{T}}$. If we substitute this value of α , we get

$$\varepsilon \leq \frac{\rho\lambda}{\sqrt{T}}.$$

Solving for T then yields

$$T \geq \left(\frac{\rho\lambda}{\varepsilon}\right)^2.$$

Since $\rho \geq \sqrt{2K}$, the larger number of iterations $T \geq 2K(\frac{\lambda}{\varepsilon})^2$ will suffice. Substituting this value of T into the equation given for α above yields $\alpha = \frac{\varepsilon}{\lambda^2}$. \square

We also prove the following corollary, which gives a more precise bound on the number of required iterations when \mathcal{S} is a (truncated) ellipsoid.

Corollary 12 (Supergradient Ascent Efficiency under Ellipsoidal Uncertainty). *For a truncated ellipsoidal uncertainty set, the supergradient ascent component of RRA converges to within ε of the maximin optimal continuous assignment \hat{A}^{MM} in $\mathbf{O}\left(\frac{2Km}{n\varepsilon^2}\right)$ iterations.*

Proof. Because the subgradient is always contained in \mathcal{S} , we can easily upper-bound the norm of it. We have for all $A \in \mathcal{A}$ and $S \in \mathcal{S}$ that

$$\|\nabla_A W(A, S)\|_2 \leq \max_{S \in \mathcal{S}} \frac{1}{n} \|S\|_2.$$

In turn, if the ellipsoid has center $\bar{\mu}_{\mathcal{S}}$ and radius q , then we can use the triangle inequality to show

$$\max_{S \in \mathcal{S}} \|S\|_2 \leq \|\bar{\mu}_{\mathcal{S}}\|_2 + q,$$

where q can be computed as the maximum eigenvalue of $\Sigma_{\mathcal{S}}$. Both $\|\bar{\mu}_{\mathcal{S}}\|_2$ and q are $O(\sqrt{nm})$, so we see that λ is $O(\sqrt{\frac{m}{n}})$. Applying Proposition 11 completes the proof. \square

Theorem 14 (Maximin and Integrality Gaps in Welfare). *Suppose \mathcal{S} is a (δ, γ) uncertainty set with \mathcal{L}_1 diameter L . Let A^ε denote an ε -optimal discrete RAU solution, and \tilde{A}^ε denote an ε -optimal continuous RAU solution. Let A' denote the random variable that arises from applying the randomized rounding procedure ROUND to \tilde{A}^ε , and assume that ROUND preserves expectation, i.e., $\mathbb{E}_{A'}[A'] = \tilde{A}^\varepsilon$. Suppose also that the true affinity scores are S^* , and denote the optimal solution $A^* \doteq \arg \max_{A \in \mathcal{A}} W(A, S^*)$. The following then hold.*

1. *Maximin Gap:* $\mathbb{P}(W(A^*, S^*) - W(A^\varepsilon, S^*) > \varepsilon + \frac{2\gamma+L}{n}) < \delta$.
2. *Expected Regret of RRA:* $\mathbb{P}(W(A^*, S^*) - \mathbb{E}_{A'}[W(A', S^*)] > \varepsilon + \frac{2\gamma+L}{n}) < \delta$.
3. *Probabilistic Regret of RRA:* $\mathbb{P}(W(A^*, S^*) - W(A', S^*) > \frac{\varepsilon + (2\gamma+L)/n}{\delta'}) < \delta' + \delta$.

Proof. We begin with the maximin gap bound. With probability at least $1 - \delta$, there is some $S_\gamma^* \in \mathcal{S}$ such that $\|S_\gamma^* - S^*\|_1 \leq \gamma$, and thus $W(A^*, S^*) \leq W(A^*, S_\gamma^*) + \frac{\gamma}{n}$ (by Proposition 2). Likewise, $W(A^\varepsilon, S_\gamma^*) \leq W(A^\varepsilon, S^*) + \frac{\gamma}{n}$. Together, these imply

$$\begin{aligned} W(A^*, S^*) - W(A^\varepsilon, S^*) &\leq W(A^*, S_\gamma^*) - W(A^\varepsilon, S_\gamma^*) + \frac{\gamma}{n} \\ &\leq W(A^*, S_\gamma^*) - W(A^\varepsilon, S_\gamma^*) + 2\frac{\gamma}{n}. \end{aligned}$$

Now, $W(A^*, S_\gamma^*) \leq \max_{A \in \mathcal{A}} \inf_{S \in \mathcal{S}} W(A, S) + \frac{L}{n}$, and by definition $W(A^\varepsilon, S_\gamma^*) \geq \inf_{S \in \mathcal{S}} W(A^\varepsilon, S)$, thus

$$W(A^*, S_\gamma^*) - W(A^\varepsilon, S_\gamma^*) + 2\frac{\gamma}{n} \leq \max_{A \in \mathcal{A}} \inf_{S \in \mathcal{S}} W(A, S) - \inf_{S \in \mathcal{S}} W(A^\varepsilon, S) + \frac{2\gamma+L}{n}.$$

The definition of A^ε implies $\max_{A \in \mathcal{A}} \inf_{S \in \mathcal{S}} W(A, S) - \inf_{S \in \mathcal{S}} W(A^\varepsilon, S) \leq \varepsilon$, yielding the desired result.

To obtain the expected regret of RRA, we first apply the two facts that $\mathbb{E}_{A'}[A'] = \tilde{A}^\varepsilon$ and W is a linear objective function. We now must bound $W(A^*, S^*) - W(\tilde{A}^\varepsilon, S^*)$, which we can do using the same proof we used to obtain the maximin gap above (replacing A^ε and \mathcal{A} with \tilde{A}^ε and $\tilde{\mathcal{A}}$).

Finally, the probabilistic worst-case regret bound follows from Markov's inequality. \square