Bayesian Inferences on Uncertain Ranks and Orderings: Application to Ranking Players and Lineups[*]

Andrés F. Barrientos[†], Deborshee Sen[‡], Garritt L. Page[§], and David B. Dunson[¶]

Abstract. It is common to be interested in rankings or order relationships among entities. In complex settings where one does not directly measure a univariate statistic upon which to base ranks, such inferences typically rely on statistical models having entity-specific parameters. These can be treated as random effects in hierarchical models characterizing variation among the entities. In this paper, we are particularly interested in the problem of ranking basketball players in terms of their contribution to team performance. Using data from the National Basketball Association (NBA) in the United States, we find that many players have similar latent ability levels, making any single estimated ranking highly misleading. The current literature fails to provide summaries of order relationships that adequately account for uncertainty. Motivated by this, we propose a Bayesian strategy for characterizing uncertainty in inferences on order relationships among players and lineups. Our approach adapts to scenarios in which uncertainty in ordering is high by producing more conservative results that improve interpretability. This is achieved through a reward function within a decision theoretic framework. We apply our approach to data from the 2009–2010 NBA season.

Keywords: Bayesian, ordering statements, ranking, decision theory, sports statistics.

1 Introduction

Making inference on orderings among parameters is of widespread importance. These may represent, for example, abilities of individuals or teams performing a given task (Cattelan et al., 2013), treatment effects in clinical trials (Rücker and Schwarzer, 2015), health and social indices of geographical regions (Marriott, 2017), consumer preferences in search engines (Park et al., 2015), and educational systems (Millot, 2015). Standard statistical techniques for inferring such orderings aim to estimate a ranking of parameters. These have classically included ranking the posterior means or using the posterior expected ranks (Laird and Louis, 1989; Lin et al., 2006); we refer the reader to

Henderson and Newton (2016) for a more nuanced technique. Other methods have been developed based on the type of data at hand. For example, when data originate directly from paired comparisons between competing units, the Bradley-Terry model (Bradley and Terry 1952) has become a popular option to estimate each unit's "ability" which are then used to produce a ranking (Agresti 2002; Hunter 2004; Caron and Doucet 2012). Recent work dedicated to Bradley-Terry type models has focused on generalizations of the base model. For example, Dittrich et al. (1998) allow unit-specific covariates to inform ability estimates, and Huang et al. (2008) use group information to estimate individual abilities. When preference data are available, Vitelli et al. (2018) recently developed a method based on the Mallows model that produces rankings. In this same setting, Caron et al. (2014) developed a method that produces rankings based on a Plackett-Luce model.

This article is motivated by our interest in providing an approach for producing ordering statements for entities in professional sports leagues, with a particular focus on the National Basketball Association (NBA) of the United States of America. In what follows, we use the term "ordering statement" to mean the collection of rankings of a subset of objects of interest. We are interested in comparing players not on the basis of their individual statistics but in terms of their contribution to team performance. Our primary dataset is from the 2009–2010 NBA season, and our goal is to provide ordering statements associated with individuals' and groups of individuals' (lineups') abilities. This is a particularly challenging case study because many players have similar abilities and, based on the available data, the abilities need to be estimated from measurements taken at a group level, making them imprecise. Ranking statements typically summarize orderings in parameters and it can be challenging to assess and summarize uncertainty in such statements. This is especially true in scenarios where parameter values are similar or the amount of data is limited relative to the number of parameters (as in our case study). In such settings, it is particularly important to not over-interpret a single estimated ranking but instead to carefully account for uncertainty in rank-based decision making and inferences.

Despite the rich statistical literature on ranking problems, few contributions focus on characterizing uncertainty. In the Bayesian framework, methods for measuring uncertainty include point-wise credible intervals for ranks (Rodríguez and Moser, 2015) and identifying rankings with high posterior probability (Soliman and Ilyas, 2009; Vitelli et al., 2018; see also Jewett et al., 2019 for some visual tools). In the frequentist framework, some theoretical contributions focus on studying asymptotic conditions under which point-wise confidence intervals of population ranks have the claimed coverage probability (Xie et al., 2009), and under which ranking estimates converge to the truth as the number of parameters and sample size increase to infinity (Hall and Miller, 2010). Other contributions aim to define confidence regions for rankings using multiple confidence intervals or hypothesis tests for the parameters while controlling the family-wise error rate (Wright et al., 2018; Klein et al., 2020; Mohamad et al., 2021). However, controlling the family-wise error rate can be challenging and often leads to regions that are so wide as to be practically meaningless. Loss functions have also been used in this regard, with the choice of the loss depending on the ranking problem at hand; examples include ranking by using a weighted loss that improves estimates of extrema (Wright

et al., 2003), finding the correct rank for a particular parameter, finding a parameter corresponding to a particular rank, or correctly ordering a pair of parameters (all in Jewett et al., 2019), identifying the top or bottom rankings (Lin et al., 2006; Henderson and Newton, 2016), optimizing the area under a receiver operating characteristic curve (Rendle et al., 2009), and learning to rank (Chen et al., 2009).

In this paper, we adopt a Bayesian framework and propose an approach for producing ordering statements of parameters along with their associated uncertainty. These ordering statements are not designed to make definitive conclusions about all parameters being compared; instead, they maintain accuracy by remaining silent when uncertainty associated with the ranking is high. These ordering statements thus represent uncertainty more realistically as compared to usual rankings. We achieve this by forming two particular types of ordering statements, which we term local and global statements. Local statements provide orderings relative to individual parameters, and global statements combine local statements. In our framework, combining local statements to form global statements leads to a potentially massive number of global statements. To address this, we propose a decision-theoretic approach to pick a global statement that balances the number of elements included in the statement and its uncertainty. The proposed approach can be parallelised, which makes it attractive from a computational perspective. We also show that the optimal global statement that maximizes the expected utility function converges to the true ranking. We highlight the fact that our approach is based solely on posterior distribution probabilities. Thus, it can be employed regardless of the statistical model used to fit the data. For this reason, we focus primarily on developing and presenting our approach that produces local and global ordering statements and ancillarily on the model employed to fit the NBA data.

The rest of the paper is organized as follows. Section 2 discusses the challenges of ranking parameters under different scenarios; this is motivated concretely by the NBA application. Because our approach is applicable regardless of the statistical model utilized, Section 3 first provides formal definitions of local and global statements, introduces the reward function used to find optimal statements, and presents an asymptotic result showing the convergence of the optimal statements to the true ranking. Section 4.1 describes the statistical model used in the NBA application, and the rest of Section 4 is dedicated to conducting a simulation study to assess the performance of our approach and compare to the cumulative probability consensus ranking described in Vitelli et al. (2018). Section 5 presents results for our analysis of the NBA data. Computer codes used to carry out the analysis in Section 5 can be found in our R-package at https://github.com/anfebar/anfebar.github.io/tree/master/Software/BOARS. Finally, we end with some concluding remarks in Section 6.

2 Motivation through basketball application

Since our focus is on NBA performance data, we begin by briefly providing some context. In basketball, a competition between two teams is based on five players within a lineup. A basketball game is a timed collection of competitions between lineups, which we refer to as "encounters". The lineup compositions for both teams change frequently during

the course of a game based on substitutions made by team coaches. It is common for teams to record a number of metrics during encounters to evaluate lineup performance. Among these are the total number of points scored and the total number of points conceded. In this work, we use the difference of these two totals as a way to measure the ability of a lineup while taking into account the quality of opponent. The available information we have is thus at the lineup level rather than the player level.

It is often stated that the NBA is a "player's league", meaning that a team's success is highly dependent on its top players' talent level. Being able to discriminate between player abilities and attract high-performing players is thus crucial for general managers/decision makers. In addition, basketball is a sport that requires precise interactions between players, and thus being able to identify collections of players that play effectively together is a coach's key responsibility. This makes it important to be able to order players and lineups. Examples of ordering statements that could prove useful to coaches and general managers are a ranking of all players or lineups across the NBA, rankings of subsets of players or lineups (within a team or between teams), pairwise comparisons between players and/or lineups, lists of players or lineups that are at the top or bottom of a team or a subset of teams, et cetera. To obtain such statements, a common approach is to first fit a model that produces a posterior distribution of player/lineup abilities, and then compute a ranking based on either a summary measure of the posterior distribution or by means of loss functions. Examples of summary measures include posterior means, the so-called r-values (Henderson and Newton, 2016), and posterior expected ranks (Laird and Louis, 1989; Lin et al., 2006), while loss functions are explored in Jewett et al. (2019).

If a posterior summary measure is used to construct an ordering of players/lineups, it is desirable that this be accompanied by a measure of uncertainty. Since an ordering statement defines a region of the parameter space, a natural measure of uncertainty is the amount of posterior probability that such a region accumulates. Ideally, one would like the posterior probability of an ordering statement to be large and also for it to be precise. For example, it would be undesirable to report a wide credible interval for the position at which a player is ranked even though the posterior probability of the interval is high. Figure 1 illustrates the challenges of producing precise ordering statements with small uncertainty in our application. This displays a caterpillar plot of each players' average per game difference between points scored and allowed for each encounter in which the player participates. This could be thought of as a type of adjusted plus/minus (Sill, 2010) in that it takes into account lineup information. It is clear that there is a massive overlap in the average difference in points scored and allowed across players. Reasons for this are: *(i)* player abilities are inferred using measurements collected at an aggregate, that is, lineup level, *(ii)* players have very similar abilities, which is not unexpected for a top-level competition, *(iii)* the number of observations for a player varies widely from around twenty to two thousand, and *(iv)* encounters are often competitive, resulting in differences in points scored close to zero. This overlap makes it so that any ordering statements (as those listed in the previous paragraph) that contain a moderate to large subset of players (and thus lineups) would likely have negligible posterior probability. For this reason, it is imperative to attach a measure of uncertainty to these types of
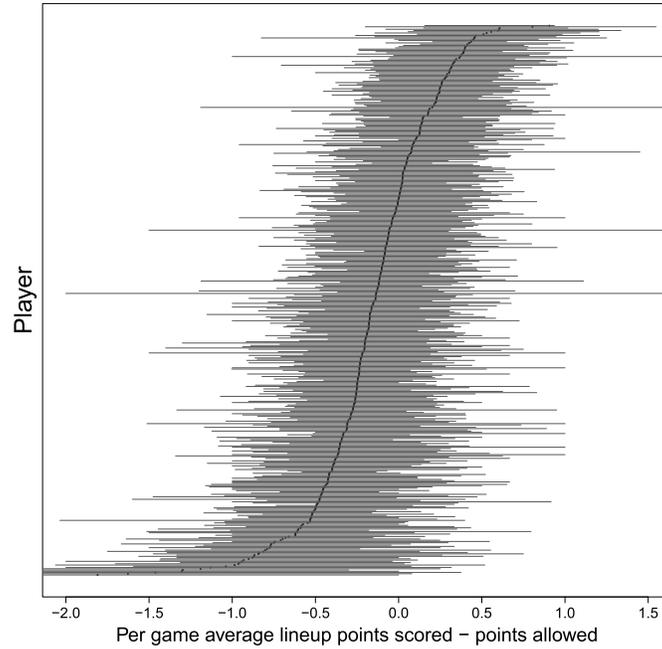
Figure 1: Per game average of points scored minus points allowed of each encounter in which a player participated. Each bar represents the interquartile range of the per game average between points scored and allowed of the all encounters in which a player participated.

ordering statements so that coaches and general managers are aware that they are making decisions based on statements that may be very uncertain.

In scenarios like our NBA application where any ranking/ordering of a large number of players/lineups is likely to have massive uncertainty, an alternative option is to focus on making statements for a small subset of players/lineups. This can include, for example, pairwise comparisons. Doing so would produce ordering statements that are likely to be less uncertain. However, if a coach/general manager makes decisions based on a collection of statements that consist of a small subset of players/lineups, then in order to correctly quantify uncertainty, they will need to compute the joint posterior probability of each of the "smaller" statements occurring simultaneously. Unfortunately, this joint probability is again likely to be low. In Section 5.2, we provide some examples of how typical approaches lead to ordering statements which have posterior probability that are essentially zero.

In this article, we propose an approach that attempts to find ordering statements which accumulate a large joint posterior probability. Our approach focuses on exploring the support of the posterior distribution and identifying a region (or regions) in the parameter space for which a non-negligible concentration of posterior probability

occurs, and in turn connecting the region to a set of ordering statements. The posterior probability of the identified regions is equal to the joint probability of the corresponding ordering statements, which corresponds to our measure of uncertainty. In addition, not only do we attempt to find an ordering statement on which a non-negligible amount of posterior probability is concentrated, but, if possible, we find one that involves a reasonable number of players/lineups as well. We achieve this by forming local and global statements, where local statements provide orderings relative to individual parameters, and global statements combine local statements; concrete definitions of these statements are provided in Section 3. When comparing the abilities of two players/lineups within a local statement, say ξ_1 and ξ_2, the global statement should refrain from saying whether player 1 is better or worse than player 2 (that is, $\xi_1 > \xi_2$ or $\xi_1 < \xi_2$, respectively) if the posterior probability of the event $\xi_1 > \xi_2$ is not reasonably large or small, respectively. Thus, local statements can end up being empty. When the global statement only involves a few local statements comprised of a few players/lineups, we call it a sparse global statement. If sparsity needs to be reduced, there are two options: *(i)* work with a smaller posterior probability, or *(ii)* introduce a mechanism that allows coaches or managers to introduce some error and accept that some players or lineups may be misordered in some of the ordering statements. Thus, for each ordering statement, we need to provide not only the posterior probability, but also a proper quantification of such error. Our proposed approach to achieve this is both novel and complementary to the existing literature of alternate methods to estimate ordering statements.

Decision makers will be able to employ the ordering statements our approach provides with confidence that uncertainty is explicitly and properly calibrated. However, it may be the case that comparisons between two or more players/lineups of interest is not available in the global statement, implying that uncertainty associated with the comparisons is large. If this is the case, then the coach/general manager's options are twofold: *(i)* make a decision based on, for example, the posterior mean of the abilities, r-values, posterior expected ranks in the face of high uncertainty (which essentially boils down to a "guess"), or *(ii)* collect a different set of data that is able to better discriminate between desired players or lineups and run our procedure based on these data. The next section presents a detailed description of our proposed approach.

Before describing our approach, we briefly mention that analyzing player contributions to lineup effectiveness and overall team strength has appeared in the basketball analytics literature. A small sampling of these are Kalman and Bosch (2020), who discover effective lineups based on new position definitions, Page et al. (2007), use box-score variables to highlight player skills that are conducive to winning games, and Deshpande and Jensen (2016), who use situational contributions to estimate player quality and produce a player ranking. That being said, we are not aware of any work in sports analytics that addresses uncertainty quantification associated with rankings or orderings.

3   Bayesian ordering statements

This section contains the main conceptual contribution of this article. We begin by providing a formal description of the setup we are working under in Section 3.1. We then

introduce notions of local and global statements in Section 3.2. These are formed using loss parameters that users are free to pick according to their needs; this is discussed further in Section 3.2. We provide a decision-theoretic approach to tune the loss parameters in Section 3.3; this leads to optimal global statements. We show in Section 3.4 that the optimal global statement converges to the true ranking asymptotically as the sample size increases.

3.1 Setup

Formally, our goal is to compare L unknown parameters $\xi_1, \ldots, \xi_L \in \Xi$; these represent player abilities in our application, and will also later be used to construct lineup abilities. We use the shorthand notation $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_L) \in \Xi^L$, and we let $\mathcal{B}(\Xi^L)$ denote the Borel sigma-algebra on Ξ^L. We assume that Ξ^L allows orderings among parameters, that is, for every $(\xi_1, \ldots, \xi_L) \in \Xi^L$ and every pair (l, l'), $1 \leq l \neq l' \leq L$, one of $\xi_{l'} > \xi_l$ or $\xi_l = \xi_{l'}$ or $\xi_l > \xi_{l'}$ is true.

We adopt a Bayesian approach in this article and assume that there is a statistical model with parameters $\boldsymbol{\xi}$, and we let P be the posterior distribution of $\boldsymbol{\xi}$ given observations, where we omit its explicit dependence on observations to keep the notation simple; P is defined on the space $(\Xi^L, \mathcal{B}(\Xi^L))$. It is reasonable to assume that no two players have exactly the same ability. In other words, either $\xi_{l'} > \xi_l$ or $\xi_l > \xi_{l'}$ for every $l \neq l'$ and $\boldsymbol{\xi} \in \Xi^L$, and we shall assume this in the remainder of the article.

Throughout this article, we shall be interested in making ordering statements among individual parameters, which we formally define as follows.

Definition 1. *An ordering statement is an event in $\mathcal{B}(\Xi^L)$ that establishes order relations among (a subset of) individual parameters. No order relations are assumed among the remaining parameters.*

A trivial ordering statement is the empty statement, that is, entire space Ξ^L. While this has posterior probability one, it provides no information among the relative orderings of any individual parameters. The simplest meaningful ordering statements are those which order only two individual parameters, for example, $\{\boldsymbol{\xi} \in \Xi^L : \xi_1 < \xi_2, \xi_3, \ldots, \xi_L \in \Xi\}$. For simplicity, we shall just write this as $(\xi_1 < \xi_2)$ and not make explicit that $\xi_3, \ldots, \xi_L \in \Xi$. An example of a strict ordering statement is $\xi_1 < \cdots < \xi_L$, which provides an ordering among all the individual parameters; this typically has negligible posterior probability. The above examples are of relatively simple ordering statements, and we shall construct more sophisticated ordering statements in the sequel. In the remainder of this section, we fix the posterior distribution P and construct statements based on it.

3.2 From local to global statements

We begin by considering pairwise comparisons. We define events $E_{l,l'} = (\xi_l > \xi_{l'})$, $1 \leq l \neq l' \leq L$ and call these elementary statements because they compare pairs of parameters. These will be used as building blocks in the sequel. In particular, $E_{l,l'}$

denotes the event that the lth parameter is ordered higher than the l'th parameter. We then choose a lower bound $(1 - \alpha)$ for the posterior probability of the elementary statements for $\alpha \in [0, 1]$, we define sets

$$\overline{\mathcal{A}}_{l,\alpha} = \{l' \in \{1, \ldots, L\} : \mathsf{P}(E_{l',l}) > 1 - \alpha\},$$
$$\underline{\mathcal{A}}_{l,\alpha} = \{l' \in \{1, \ldots, L\} : \mathsf{P}(E_{l,l'}) > 1 - \alpha\}. \tag{3.1}$$

The set $\overline{\mathcal{A}}_{l,\alpha}$ consists of all parameters that are ordered higher than the lth parameter with posterior probability at least $(1 - \alpha)$, and the set $\underline{\mathcal{A}}_{l,\alpha}$ consists of all parameters that are ordered lower than the lth parameter with posterior probability at least $(1-\alpha)$; recall that we have fixed the posterior distribution before constructing statements. In other words, the sets $\underline{\mathcal{A}}_{l,\alpha}$ and $\overline{\mathcal{A}}_{l,\alpha}$ are sets of players whose posterior probability of having an ability different from that of player l is high. We order the lth parameter only with respect to parameters in $\underline{\mathcal{A}}_{l,\alpha}$ and $\overline{\mathcal{A}}_{l,\alpha}$, since ordering with respect to other parameters has a high degree of uncertainty.

We now define the following local statements which involve multiple elementary comparisons. We use the notation $|\mathcal{A}|$ to denote the cardinality of a finite set $\mathcal{A}$.

1. **Statement** $A_{l,\alpha}$: parameter l is ordered higher than all parameters in $\underline{\mathcal{A}}_{l,\alpha}$ and ordered lower than all parameters in $\overline{\mathcal{A}}_{l,\alpha}$, that is,

$$A_{l,\alpha} = (\bigcap_{l' \in \underline{\mathcal{A}}_{l,\alpha}} E_{l,l'}) \bigcap (\bigcap_{l' \in \overline{\mathcal{A}}_{l,\alpha}} E_{l',l}).$$

   In other words, $A_{l,\alpha}$ is the event where all the elementary statements associated with $\underline{\mathcal{A}}_{l,\alpha} \cup \overline{\mathcal{A}}_{l,\alpha}$ simultaneously hold.

2. **Statement** $A_{l,\alpha,t}$: at least $(1 - t) \times 100$ percent of the elementary statements in $A_{l,\alpha}$ are α-correct for $t \in [0, 1]$, where α-correct refers to an elementary statement $E_{l,l'}$ or $E_{l',l}$ such that $\mathsf{P}(E_{l,l'}) > (1 - \alpha)$ or $\mathsf{P}(E_{l',l}) > (1 - \alpha)$, respectively; we will refer to t as the local error. Formally,

$$A_{l,\alpha,t} = \bigcup_{(\underline{\mathcal{A}},\overline{\mathcal{A}}) \in \mathscr{A}_{l,\alpha,t}} (\bigcap_{l' \in \underline{\mathcal{A}}} E_{l,l'}) \bigcap (\bigcap_{l' \in \overline{\mathcal{A}}} E_{l',l}), \tag{3.2}$$

   where

$$\mathscr{A}_{l,\alpha,t} = \left\{ (\underline{\mathcal{A}}, \overline{\mathcal{A}}) : \underline{\mathcal{A}} \subseteq \underline{\mathcal{A}}_{l,\alpha}, \overline{\mathcal{A}} \subseteq \overline{\mathcal{A}}_{l,\alpha}, |\underline{\mathcal{A}}| + |\overline{\mathcal{A}}| \geq (1 - t) \times |\underline{\mathcal{A}}_{l,\alpha} \cup \overline{\mathcal{A}}_{l,\alpha}| \right\}.$$

   Notice that $\underline{\mathcal{A}} \cup \overline{\mathcal{A}}$ with $(\underline{\mathcal{A}}, \overline{\mathcal{A}}) \in \mathscr{A}_{l,\alpha,t}$ is a collection of indices with cardinality greater than or equal to $(1-t) \times |\underline{\mathcal{A}}_{l,\alpha} \cup \overline{\mathcal{A}}_{l,\alpha}|$ and for which either $\mathsf{P}(E_{l',l}) > 1 - \alpha$ or $\mathsf{P}(E_{l',l}) > 1 - \alpha$.

$A_{l,\alpha}$ is the subset of Ξ^L where $\xi_{l'} < \xi_l$ for all $l' \in \underline{\mathcal{A}}_{l,\alpha}$ and $\xi_l < \xi_{l'}$ for all $l' \in \overline{\mathcal{A}}_{l,\alpha}$; where we note that the sets $\underline{\mathcal{A}}_{l,\alpha}$ and $\overline{\mathcal{A}}_{l,\alpha}$ were chosen based on the posterior distribution P. This is relaxed in the event $A_{l,\alpha,t}$, which does not require all the elementary statements in $A_{l,\alpha}$ to hold. The posterior probability of at least one of the elementary statements in $A_{l,\alpha}$ being α-incorrect is a non-decreasing function of the number of such

statements in $A_{l,\alpha}$, and therefore $\mathsf{P}(A_{l,\alpha})$ is a non-increasing function of the number of elementary statements in $A_{l,\alpha}$. Since our final aim is to define statements having high posterior probability, statement $A_{l,\alpha,t}$ considers the same elementary statements as $A_{l,\alpha}$ but also allowing for a local error determined by t.

To create statements involving multiple local statements for a fixed t and α, we consider only those players whose corresponding $A_{l,\alpha,t}$ have posterior probability lower-bounded by $(1 - \gamma)$ for $\gamma \in [0,1]$,

$$\mathcal{G}_{\alpha,t,\gamma} = \{l \in \{1, \ldots, L\} \, : \, \mathsf{P}\left(A_{l,\alpha,t}\right) \geq 1 - \gamma\}. \tag{3.3}$$

We now define the following global statements.

1. **Statement $G_{\alpha,t,\gamma}$**: statements $A_{l,\alpha,t}$ simultaneously hold for all $l \in \mathcal{G}_{\alpha,t,\gamma}$, that is, $G_{\alpha,t,\gamma} = \bigcap_{l \in \mathcal{G}_{\alpha,t,\gamma}} A_{l,\alpha,t}$. In other words, $G_{\alpha,t,\gamma}$ is the event where all the local statements associated with $\mathcal{G}_{\alpha,t,\gamma}$ simultaneously hold.

2. **Statement $G_{\alpha,t,\gamma,q}$**: at least $(1-q) \times 100$ percent of the local statements comprising $A_{l,\alpha,t}$ are γ-correct for $q \in [0,1]$, where γ-correct refers to a statement $A_{l,\alpha,t}$ such that $\mathsf{P}(A_{l,\alpha,t}) \geq 1 - \gamma$; we will refer to q as the global error. Formally,

$$G_{t,\alpha,\gamma,q} = \bigcup_{\mathcal{G} \in \mathscr{G}_{\alpha,t,\gamma,q}} \bigcap_{l \in \mathcal{G}} A_{l,\alpha,t}, \tag{3.4}$$

   where

$$\mathscr{G}_{\alpha,t,\gamma,q} = \{\mathcal{G} \, : \, \mathcal{G} \subseteq \mathcal{G}_{\alpha,t,\gamma}, \, |\mathcal{G}| \geq (1 - q) \times |\mathcal{G}_{\alpha,t,\gamma}|\}.$$

   Notice that $\mathcal{G} \subseteq \mathcal{G}_{\alpha,t,\gamma}$ is a collection of indices l' with cardinality greater than or equal to $(1 - q) \times |\mathcal{G}_{\alpha,t,\gamma}|$ and for which $\mathsf{P}\left(A_{l,\alpha,t}\right) \geq 1 - \gamma$.

Recall that we choose the set $\mathcal{G}_{\alpha,t,\gamma}$ based on the posterior probability of events $A_{l,\alpha,t}$. Having chosen $\mathcal{G}_{\alpha,t,\gamma}$, $G_{\alpha,t,\gamma}$ is then the subset of Ξ^L where the events $A_{l,t,\alpha}$ happen for all $l \in \mathcal{G}_{\alpha,t,\gamma}$. This is relaxed in the global statement $G_{\alpha,t,\gamma,q}$, which requires only a proportion of the events $A_{l,t,\alpha}$ for $l \in \mathcal{G}_{\alpha,t,\gamma}$ to happen.

The global statement $G_{\alpha,t,\gamma,q}$ allows a global error q that plays a role similar to that of t at the local level. These global statements are of principal interest as they can be employed to produce rankings. Section 1 of the supplementary material (Barrientos et al., 2022) details an algorithm that is able to find local and global statements along with the corresponding posterior probabilities using samples drawn from P.

One of our goals is to produce global statements that balance sparsity and the desire for statements having high posterior probability, which can be achieved through specific choices of (α, t, γ, q). The following inequalities provide intuition about the relation between (α, t, γ, q) and the posterior probability of global statements and local statements.

1. $\mathsf{P}(G_{\alpha,t,\gamma,q}) \geq \mathsf{P}(G_{\alpha,t,\gamma,q'})$ for $q \geq q'$, that is, global statements with smaller global error have smaller posterior probability.

2. $\mathsf{P}(G_{\alpha,t,\gamma,q=0}) = \mathsf{P}(G_{\alpha,t,\gamma}) \geq 1 - \gamma|\mathcal{G}_{\alpha,t,\gamma}|$, that is, when the global error is equal to zero, the posterior probability of the global statement increases as γ or the number of local statements in $G_{\alpha,t,\gamma}$ decreases.

3. $\mathsf{P}(A_{l,\alpha,t}) \geq \mathsf{P}(G_{\alpha,t,\gamma})$ if $l \in \mathcal{G}_{\alpha,t,\gamma}$, that is, when the global error is zero, the posterior probability of the global statement is upper bounded by the posterior probability of any local statement in $G_{\alpha,t,\gamma}$.

4. $\mathsf{P}(A_{l,\alpha,t=0}) = \mathsf{P}(A_{l,\alpha}) \geq 1 - \alpha|\underline{\mathcal{A}}_{l,\alpha} \cup \overline{\mathcal{A}}_{l,\alpha}|$, that is, when the local error is equal to zero, the posterior probability of the local statement increases as α or the number of elementary statements in $A_{l,\alpha}$ decreases.

Moreover, for fixed q and t, the number of elementary statements in $\mathcal{A}_{l,\alpha,t}$ and local statements in $\mathcal{G}_{\alpha,t,\gamma,q}$ decreases as either α or γ decreases. We can thus obtain less sparse statements by setting α and γ at large values, which in turn increases the posterior probability of the global statement. The roles of q and t are less obvious. For example, if we increase q and t, it is not clear how this will impact the posterior probability of $G_{\alpha,t,\gamma,q}$. On one hand, the posterior probability of the local statements will increase leading to less sparse global statements which may decrease its posterior probability; while on the other hand, an increase of q will also allow more errors which may result in an increase in the posterior probability. To deal with the effect of variations of (α, t, γ, q) on the posterior probability and on sparsity of $\mathcal{G}_{\alpha,t,\gamma,q}$, we introduce a decision theoretic framework in the next Section 3.3 to select values of (α, t, γ, q), and illustrate their impact on the ordering statements in Section 5.

3.3   Decision theoretic framework

In the previous section, we defined a hierarchy of ranking statements (elementary, local, and global). Our goal is to identify global statements having high posterior probability while simultaneously including as many local statements as possible. To this end, we formulate this as an optimization problem within a decision theoretic framework. We use a reward function and assume that the space of actions corresponds to the set of all possible global statements, that is, each action is equivalent to a global statement. We can in principle construct global statements by putting together any arbitrary combination of elementary and local statements without using a posterior distribution. Since the space of actions is finite, we should then compute the reward function for each of these statements and choose the one that maximizes the reward. However, the space of actions is massive, and it turns out to be infeasible to evaluate the reward function for all actions. For computational efficiency, we discard in advance actions that have low posterior probability. This is why we consider global statements constructed using the posterior distribution, that is, statements of the form $G_{\alpha,t,\gamma,q}$ with non-negligible posterior probability. We remark that since we only consider global statements constructed using the posterior distribution (in particular, using the data), our strategy is not strict decision theoretic approach, but rather an approximate one.

When optimizing the reward function, locally, we seek statements that maximize the number of elementary comparisons, which is given by $|\underline{\mathcal{A}}_{l,\alpha} \cup \overline{\mathcal{A}}_{l,\alpha}|$ for the lth

parameter. However, each local statement $A_{l,\alpha,t}$ allows an error that depends on t, and we seek to simultaneously minimize this. Globally, we seek to maximize the number of local statements $|\mathcal{G}_{\alpha,t,\gamma}|$ while minimizing the error q. Finally, we want the posterior probability of the final global statement $G_{\alpha,t,\gamma,q}$ to be high. With this in mind and defining $a = (\alpha, t, \gamma, q) \in [0,1]^4$ to be an action (that in turn leads to a statement $G_{\alpha,t,\gamma,q}$), we couch this in a decision theoretic framework and seek to design a reward function satisfying the following Condition 3.1.

Condition 3.1. *The reward function assessing $G_{\alpha,t,\gamma,q}$ satisfies the following.*

1. *It is a non-decreasing function of $|\underline{\mathcal{A}}_{l,\alpha} \cup \overline{\mathcal{A}}_{l,\alpha}|$ for $l \in \mathcal{G}_{\alpha,t,\gamma}$, and of $|\mathcal{G}_{\alpha,t,\gamma}|$.*

2. *It is a non-increasing function of t and q.*

3. *It is a non-decreasing function of $\mathsf{P}(G_{\alpha,t,\gamma,q})$.*

An example of such a reward function is

$$R(a) = R(\alpha, t, \gamma, q) = C(G_{\alpha,t,\gamma,q}) \times \mathsf{P}(G_{\alpha,t,\gamma,q}), \tag{3.5}$$

where

$$C(G_{\alpha,t,\gamma,q}) = \lfloor (1-q) \times |\mathcal{G}_{\alpha,t,\gamma}| \rfloor \times \left\{ \sum_{l \in \mathcal{G}_{\alpha,t,\gamma}} \lfloor (1-t) \times |\underline{\mathcal{A}}_{l,\alpha} \cup \overline{\mathcal{A}}_{l,\alpha}| \rfloor \right\} \tag{3.6}$$

satisfies the first two parts of Condition 3.1. In this case, the factor $C(G_{\alpha,t,\gamma,q})$ equals the $(1-q) \times 100$ percent of the number of local statements within $G_{\alpha,t,\gamma,q}$ multiplied by $(1-t) \times 100$ percent of the total number of elementary statements within the local statements comprising $G_{\alpha,t,\gamma,q}$. Thus, adding (or removing) local statements in $G_{\alpha,t,\gamma,q}$ will raise (or lower) $C(G_{\alpha,t,\gamma,q})$ by $(1-q) \times 100$ percent of the number of added (or removed) local statements multiplied by $(1-t) \times 100$ percent of the total number of elementary statements within the added local statements. If the number elementary statements increases (or decreases) within a local statement already in $G_{\alpha,t,\gamma,q}$, the factor $C(G_{\alpha,t,\gamma,q})$ will change by $(1-t) \times 100$ percent of the number of added (or removed) elementary statements. It is worth noting that the reward function in equation (3.5) changes in direct proportion to $\mathsf{P}(G_{\alpha,t,\gamma,q})$ and $C(G_{\alpha,t,\gamma,q})$. If either $C(G_{\alpha,t,\gamma,q})$ or $\mathsf{P}(G_{\alpha,t,\gamma,q})$ change in $\pm 1\%$, the reward function will also change in $\pm 1\%$. Changes in reward function values are primarily governed by the trade-off between $C(G_{\alpha,t,\gamma,q})$ and $\mathsf{P}(G_{\alpha,t,\gamma,q})$ because less sparse global statements (i.e., larger $C(G_{\alpha,t,\gamma,q})$) leads to smaller posterior probabilities (i.e., smaller $\mathsf{P}(G_{\alpha,t,\gamma,q})$). Depending on the application, the dynamics of this trade-off will differ.

With an abuse of notation, the function C is interpreted as a measure of the quality of a global statement G comprised of local statements as described in Section 3.2; this does not necessarily have to be constructed using the parameters (α, t, γ, q). Restricting ourselves to the set of global statements that are indeed constructed using parameters (α, t, γ, q), the reward function (3.5) corresponds to the posterior mean of a utility

function of the form $u\{(\xi_1, \ldots, \xi_L), a\} = C(G_{\alpha,t,\gamma,q}) \times \mathbb{1}(G_{\alpha,t,\gamma,q})$. Doing this allows us to reduce the space of all possible global statements to a more manageable size and constrain the optimization to statements having high posterior probability.

An efficient algorithm to approximate the optimizer of the reward function $\widehat{a} = (\widehat{\alpha}, \widehat{t}, \widehat{\gamma}, \widehat{q}) = \mathrm{argmax}_{a \in [0,1]^4} R(a)$ is provided in Section 2 of the supplementary material. Section 3.4 provides asymptotic guarantees showing that, under this strategy, the optimal global statement $G_{\widehat{\alpha}, \widehat{t}, \widehat{\gamma}, \widehat{q}}$ converges to the "true" ranking. Since users may also be interested in fixing local and global errors at desired values or making statements that ensure a minimum posterior probability, the algorithm in Section 2 of the supplementary material also enables users to approximate optimal global statements over any subset of (α, t, γ, q) such that $\mathsf{P}(G_{\alpha,t,\gamma,q})$ is above a given value.

The computational cost of finding a global statement given M samples drawn from the posterior distribution is $\mathcal{O}(ML^2)$; this is because we have to consider all possible pairwise comparisons to form the sets $E_{l,l'}$, $\overline{\mathcal{A}}_{l,\alpha}$ and $\underline{\mathcal{A}}_{l,\alpha}$. However, this can be trivially parallelised since they can be computed independently of each other for each parameter. Having written code in R and C++, even without parallelisation, it takes only a few minutes to compute global statements for up to 4×10^3 parameters and 2×10^3 posterior samples. Precise results for computational times required to approximate the optimal global statements are provided in Section 4.3 for a simulation study, and in Sections 5.1 and 5.2 for the NBA application.

3.4   Large data limit

We study the asymptotic behavior of the proposed procedure when there is a "true" underlying $\boldsymbol{\xi}^\star = (\xi_1^\star, \ldots, \xi_L^\star)$ that generates the data and the probability measure P concentrates around this truth as the amount of data increases. This is the case when P is a consistent posterior distribution. Writing the measure P as a function of N as $\mathsf{P}^{(N)}$, where N is the number of observations, we assume that the following consistency condition holds:

$$\mathsf{P}^{(N)}\{\boldsymbol{\xi} \in B_\epsilon(\boldsymbol{\xi}_\star)\} \to 1 \text{ in } \mathbb{P}_{\boldsymbol{\xi}^\star}\text{-probability for any } \epsilon > 0, \tag{3.7}$$

where $B_\epsilon(\boldsymbol{\xi}_\star)$ denotes a ball around $\boldsymbol{\xi}^\star$ of radius $\epsilon > 0$, $B_\epsilon(\boldsymbol{\xi}_\star) = \{\boldsymbol{\xi} \in \Xi^L : \mathrm{dist}(\boldsymbol{\xi}, \boldsymbol{\xi}_\star) \leq \epsilon\}$, $\mathrm{dist}(\cdot, \cdot)$ denotes a distance measure on Ξ^L, and $\mathbb{P}_{\boldsymbol{\xi}^\star}$ denotes the data-generating mechanism under parameter value $\xi^\star$. In the following Proposition 2, we show that our method is guaranteed to produce a traditional ranking with high posterior probability as $N \to \infty$ while assuming that, without loss of generality, $\xi_1^\star \prec \cdots \prec \xi_L^\star$.

Proposition 2. *Let $\boldsymbol{\xi}^\star$ be such that $\xi_1^\star \prec \cdots \prec \xi_L^\star$. If the distribution $\mathsf{P}^{(N)}$ concentrates at $\boldsymbol{\xi}^\star$ in $\mathbb{P}_{\boldsymbol{\xi}^\star}$-probability as $N \to \infty$ (that is, equation (3.7) holds), the optimal action $(\widehat{\alpha}^{(N)}, \widehat{t}^{(N)}, \widehat{\gamma}^{(N)}, \widehat{q}^{(N)})$ converges to $(\widehat{\alpha}^{(\infty)}, \widehat{t}^{(\infty)}, \widehat{\gamma}^{(\infty)}, \widehat{q}^{(\infty)}) \equiv (0, 0, 0, 0)$ in $\mathbb{P}_{\boldsymbol{\xi}^\star}$-probability, with $G_{0,0,0,0}$ stating that $\xi_1 \prec \cdots \prec \xi_L$.*

Proof of Proposition 2. We can find $\epsilon > 0$ such that $\xi_1 \prec \cdots \prec \xi_L$ for all $\boldsymbol{\xi} \in B_\epsilon(\boldsymbol{\xi}^\star)$, for example, by choosing $\epsilon = \min_{l \in \{1, \ldots, L\}} |\xi_{l+1} - \xi_l|/3$ if $\xi_l \in \mathbb{R}$ for all l. By equation (3.7),

for any $\delta > 0$, we can choose N large enough such that

$$\mathsf{P}^{(N)}\left\{B_\epsilon(\boldsymbol{\xi}^\star)\right\} > 1 - \delta, \tag{3.8}$$

which implies $\mathsf{P}^{(N)}(\xi_1 \prec \cdots \prec \xi_L) > 1 - \delta$. Choosing

$$\delta < \min\{\alpha, \gamma\} \tag{3.9}$$

implies that $\mathsf{P}^{(N)}(E_{l,l'}) > 1 - \delta$ for $l > l'$ and $\mathsf{P}^{(N)}(E_{l',l}) > 1 - \delta$ for $l < l'$. Since $\delta < \alpha$, from equation (3.2) we have that $A_{l,\alpha,t} \supseteq B_\epsilon(\boldsymbol{\xi}^\star)$ for all $l \geq 0$, and thus $\mathsf{P}^{(N)}(A_{l,\alpha,t}) \geq \mathsf{P}^{(N)}\{B_\epsilon(\boldsymbol{\xi}^\star)\} \geq 1 - \delta$ for all $l \geq 0$. Since $\delta < \gamma$, this implies that $\mathcal{G}_{\alpha,t,\gamma} = \{1, \ldots, L\}$ for any $\gamma > 0$. This in turn implies that $G_{\alpha,t,\gamma,q} \supseteq B_\epsilon(\boldsymbol{\xi}^\star)$ for all $q \geq 0$. Thus the reward function (3.5) is maximized for $t = q = 0$. Since the reward function (3.5) is a decreasing function of α and γ, and since $\mathsf{P}(G_{\alpha,t,\gamma,q}) \geq 1 - \delta$, the reward function is maximized for $a = (\widehat{\alpha}^{(\infty)}, \widehat{t}^{(\infty)}, \widehat{\gamma}^{(\infty)}, \widehat{q}^{(\infty)})$. The result now follows from equations (3.9) and (3.8). □

4 Simulation study

We perform a simulation study to investigate the global statements defined in Section 3. The simulation study is based on a subset of the NBA lineup data so that synthetic datasets mimic the inherent complexities of the NBA lineup data. Section 4.1 introduces the approach adopted to model the data (which is used later in Section 5) and generate synthetic datasets. Section 4.2 describes simulation of the synthetic datasets and factors considered in the simulation study, and Section 4.3 presents the obtained results.

4.1 Statistical model

The model used to estimate player abilities based on the NBA data is detailed; this will be used in the sequel to illustrate our approach. We emphasize that the framework of Section 3 is applicable regardless of the model posed and the dataset considered.

Let z_i and z_i' denote the points scored by the two lineups that competed in the ith encounter, where $i = 1, \ldots, N$, and N denotes the number of encounters. Each lineup is comprised of five players and, as before, we assume that $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_L) \in \mathbb{R}^L$ denotes the abilities of the L total players. Let $\mathcal{L}_i \subset \{1, \ldots, L\}$ and $\mathcal{L}_i' \subset \{1, \ldots, L\}$ such that $\mathcal{L}_i \cap \mathcal{L}_i' = \emptyset$ identify the five players of the two lineups in the ith encounter. The ability of lineup $\mathcal{L}_i$ is defined as

$$\tau_{\mathcal{L}_i} = \sum_{k \in \mathcal{L}_i} \xi_k. \tag{4.1}$$

By defining the ability of lineup $\mathcal{L}_i$ as the sum of the abilities of individuals in the lineup, we induce correlation among lineups that share players and achieve a more parsimonious setup. Following Huang et al. (2008), we model the difference in points scored as

$$(z_i - z_i') \mid \boldsymbol{\xi}, \sigma^2 \overset{\text{ind}}{\sim} \text{Normal}(\tau_{\mathcal{L}_i} - \tau_{\mathcal{L}_i'}, \sigma^2), \quad i = 1, \ldots, N. \tag{4.2}$$

We estimate $\boldsymbol{\xi}$ and σ^2 under a Bayesian approach, and choices for the prior distribution of $\boldsymbol{\xi}$ and σ^2 are discussed in Section 4.2. Notice that (4.1) and (4.2) correspond to an

extension of the Bradley-Terry model (Bradley and Terry, 1952) for continuous scores and that includes group information (i.e., lineups). If our goal were to create ordering statements for lineups only, a potential modeling strategy would be to have lineup-specific parameters. However, this would result in a less parsimonious model, which would introduce more uncertainty, and prediction for unobserved lineups would not be available.

We note that the model given by equations (4.1) and (4.2) is not identifiable. Different conditions can be imposed for identifiability, such as requiring $L^{-1} \sum_{l=1}^{L} \xi_l = 0$, which is interpreted as an average player having ability zero, or requiring $\xi_{\widetilde{l}} = 0$ for a fixed $\widetilde{l} \in \{1, \ldots, L\}$, in which case all other abilities are interpreted as relative to the ability of the $\widetilde{l}$th player. For this paper, we impose the latter constraint with $\widetilde{l}$ being the player with the highest frequency of appearance.

We acknowledge that the structure of the NBA data is particularly complex and model (4.2) does not account for some of the structure present in the data. For example, point differences are discrete random variables, player abilities are expected to be correlated with playing times, some players might have similar abilities (i.e., the modeling strategy could account for ties), information beyond lineup composition (e.g., homecourt) could influence point differences, and the definition of lineup abilities should include interaction terms between players. Nonetheless, this simple model is expected to be able to capture important orderings among players' and lineups' abilities.

4.2   Simulation setting

We conduct a simulation study to examine the performance of the proposed method under four different conditions: (i) the prior distributions for $\boldsymbol{\xi}$ and σ^2, (ii) the sample size (which is given by the number of encounters), (iii) the true abilities to be estimated, and (iv) the number of parameters to be ordered (which is given by the number of players). In what follows, we provide specific details about how synthetic datasets are generated; as a high level description, our approach is to first fit model (4.2) to a subset of the NBA data and then employ the posterior distribution of $\boldsymbol{\xi}$ and σ^2 as the seed to generate true player abilities.

To complete the model specification described in the previous section, we consider the following priors for $\boldsymbol{\xi}$ and σ^2.

Prior 1: $\xi_l \mid \mu \overset{\text{iid}}{\sim} \text{Laplace}(\mu, 3), \quad l = 1, \ldots, L,$

Prior 2: $\xi_l \mid \mu \overset{\text{iid}}{\sim} \text{Laplace}(\mu, \sigma \lambda^{-1/2}), \quad l = 1, \ldots, L, \quad \lambda \sim \text{Gamma}(1, 1),$

Prior 3: $\xi_l \mid \mu \overset{\text{iid}}{\sim} \text{Normal}(\mu, \sigma^2 \lambda), \quad l = 1, \ldots, L, \quad \lambda \sim \text{Half-Cauchy}(0, 1),$

with $\pi(\sigma^2) \propto \sigma^{-2}$ and $\mu \sim \text{Normal}(0, 3^2)$. Prior 1 represents an informative scenario that assumes knowledge of the approximate variability of ξ_ls. Priors 2 and 3 are less informative and induce shrinkage towards the mean μ, with prior 3 producing more shrinkage than prior 2. Prior 2 is motivated by Park and Casella (2008) and Hans (2009, 2010), and prior 3 by Gelman (2006) and Polson and Scott (2012). The prior

variance for μ was set to 9 due to the fact that most encounters between lineups are competitive and so few result in point differences of double digits.

Since the idea is to simulate data that mimic the complexities of the NBA lineup data, we consider two subsets of data with five and ten randomly selected teams, respectively. The subset with five teams has 78 players and 844 encounters, while the one with ten teams has 167 players and 3604 encounters. For both subsets, we sample from the posterior distribution of the abilities of players based on the three priors described above. We use Stan (Stan Development Team, 2018) to collect 10^4 samples from each posterior distribution, from which we discard the initial 20% as burn-in and thin every 8 after that. We use the no-U-turn sampler (NUTS) (Hoffman and Gelman, 2014), which is an adaptive version of Hamiltonian Monte Carlo (Duane et al., 1987). From the 10^4 samples, we randomly select fifty samples denoted by $(\tilde{\boldsymbol{\xi}}_{(j,k,d)}, \tilde{\sigma}^2_{(j,k,d)})$, with $j = 1, \ldots, 50$, indexing the samples, $k = 1, 2, 3$ indexing the prior used for data generation, and $d = 5, 10$ being the number of teams. The use of priors 1, 2, and 3 for data generation offers different levels of similarity among true abilities, with prior 1 and 3 leading to the least and most similar abilities, respectively.

We generate synthetic datasets as follows. For each value of (j, k, d), we assume that $\tilde{\boldsymbol{\xi}}_{(j,k,d)}$ and $\tilde{\sigma}^2_{(j,k,d)}$ are the true abilities and variance, respectively. Given the shrinkage properties of the considered priors, the abilities in $\tilde{\boldsymbol{\xi}}_{(j,3,d)}$ are more similar to each other than those in $\tilde{\boldsymbol{\xi}}_{(j,2,d)}$, and $\tilde{\boldsymbol{\xi}}_{(j,1,d)}$ represents the case where abilities are the least similar. Using model (4.2), we simulate differences in points scored assuming that the lineups in each encounter are the same as observed in the NBA data. The synthetic datasets thus are comprised of 844 (for $d = 5$) and 3604 (for $d = 10$) encounters. Moreover, for each (j, k, d), we consider two sets of experiments indexed by $s = 1, 2$. In the first set ($s = 1$), we consider the encounters as they are, while in the second set ($s = 2$), we generate two independent realizations of the differences in points scored z_i for each encounter. Thus for the second set of experiments, the synthetic datasets are comprised of 1688 (for $d = 5$) and 7208 (for $d = 10$) encounters. In total, we are generating 600 ($50 \times 3 \times 2 \times 2$) synthetic datasets, where each of them are associated with a particular combination of (j, k, d, s). We use $\mathrm{D}^{(j,k,d,s)}$ to denote the synthetic dataset corresponding to combination (j, k, d, s).

For each synthetic dataset we sample from the posterior distributions of the abilities of players and lineups using the three priors described above. We use the algorithm in Section 2 of the supplementary material and restrict the search to statements with posterior probability greater than 0.9 to obtain the optimal global statements. Let $G^{(j,k,d,s,m)}_{\widehat{a}}$ denote the optimal global statement – note that $\widehat{a} = (\widehat{\alpha}, \widehat{t}, \widehat{\gamma}, \widehat{q})$ is also a function of (j, k, d, s, m) – obtained from model (4.2), prior m, and dataset $\mathrm{D}^{(j,k,d,s)}$. To evaluate the performance of our procedure at each combination $\boldsymbol{v} = (k, d, s, m)$, we compute the following measures: (i) posterior probability of the global statements; (ii) frequentist coverage, that is, fraction of global statements that are consistent with the true abilities; (iii) the number of local statements and average number of players (or lineups) within local statements for each global statement; and (iv) local and global errors.

We compare our proposed approach with the cumulative probability consensus ranking (CPCR) methodology proposed by Vitelli et al. (2018). Because the information

provided by our statements and CPCR cannot be compared directly (our approach produces a joint statement while the CPCR produces a type of marginal statement), we define a type of joint CPCR. This permits us to compare the two approaches which we do based on measures equivalent to (i) and (ii), i.e., posterior probability of a type of joint CPCR; (ii) fraction of joint CPCRs that are consistent with the true abilities. Vitelli et al. (2018) define CPCRs as follows: first select the item which has the maximum a posteriori marginal probability of being ranked 1st – we denote this item with $\xi_{(1)}$; then the item which has the maximum a posteriori marginal posterior probability of being ranked 1st or 2nd among the remaining ones – denoted by $\xi_{(1:2)}$, etc. Thus, we define joint CPCRs for players as $\bigcap_{l=L_0}^{L}[\xi_{(1:l)}$ ranked 1st, 2nd, …, or lth], $L_0 = 1, \ldots, L$, and for lineups as $\bigcap_{l=L_0}^{T}[\tau_{(1:l)}$ ranked 1st, 2nd, …, or lth], $L_0 = 1, \ldots, T$, where T is the number of observed lineups.

4.3 Results

We apply our approach using the reward function in equation (3.5) to find global statements for lineups and players. For lineups, we draw a sample from the posterior distribution of player abilities, which simultaneously leads to a sample from the posterior distribution of lineup abilities $\tau_{\mathcal{L}_i}$. The optimization described in Section 3.3 is not performed over all 5-tuples, but only over the observed lineups. Given the similarities between the patterns observed for players and lineups, we report the results for lineups in Figure 2, while results for players are deferred to Figure 1 of the online Supplementary Material. Local and global errors for statements associated with players and lineups are both reported in the online Supplementary Material (see Figure 2). Moreover, statements associated with lineups are based on the lineups from the Cleveland Cavaliers and San Antonio Spurs.

The first observation from Figure 2 is that the proposed procedure is, as expected, able to find global statements having high posterior probability (above 0.9). The number of teams $d \in \{5, 10\}$ does not appear to have a direct impact on the posterior probabilities or frequentist coverage as no clear patterns emerge. However, as expected, relative differences between true player abilities (reflected by k), coupled with the prior distribution specification m, impacts the frequentist coverage as well as the number of local statements and average number of lineups within local statements.

Generally speaking, the average number of lineups within local statements are smallest for prior 3 (that is, $m = 3$). Since making statements with a large number of lineups is desirable, this is unsatisfactory if the frequentist coverage does not remain high. Notice, however, that as the relative similarity among player and lineup abilities increases (that is, k increases), the coverage associated with priors 1 and 2 decreases quite dramatically even with sparse global statements (that is, low number of local statements and low average number of players or lineups within local statements); this too is undesirable. Thus, the balancing act alluded to earlier (correctness of statement versus the number of items the statement contains) is impacted by the prior distribution and similarity of true parameter values. Indeed, when the true abilities are most similar ($k = 3$) and the model accurately captures this ($m = 3$), we observe global statements that are sparser,
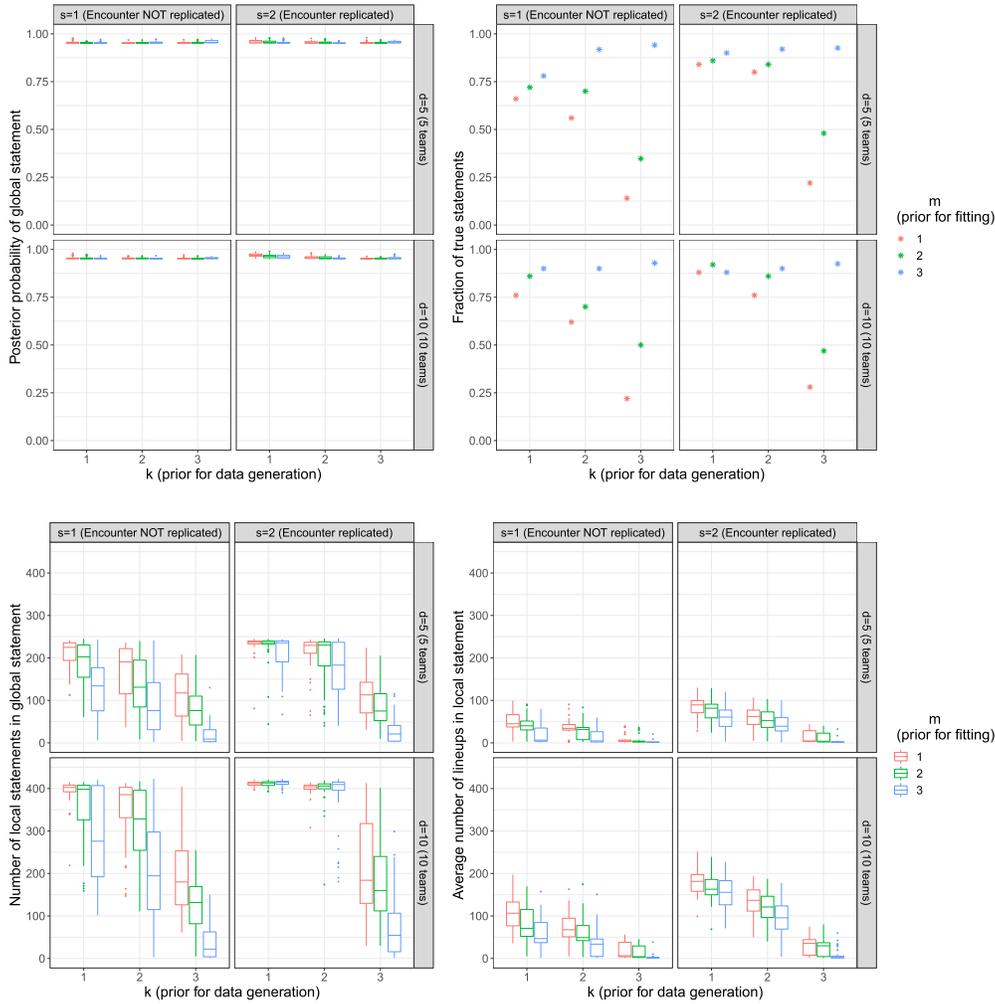
Figure 2: Simulation results for lineups assessing the effect of prior distribution used for fitting (captured by m), prior distribution used for data generation (captured by k), number of parameters (captured by d), and doubling the sample size (captured by s) with $\boldsymbol{v} = (k, d, s, m)$. The top-left plot displays the posterior probability of the global statements; the top-right plot shows the frequentist coverage; the bottom-left plot presents the number of local statements within global statements; and the bottom-right plot shows the distribution of the average number of lineups within local statements across global statements. These plots summarize results over 50 synthetic datasets.

indicating that the method refrains from collecting many local statements when forming the global statement, which is the right decision in such cases. Finally, increasing the number of encounters (moving from $s = 1$ to $s = 2$) has the desired effect of producing
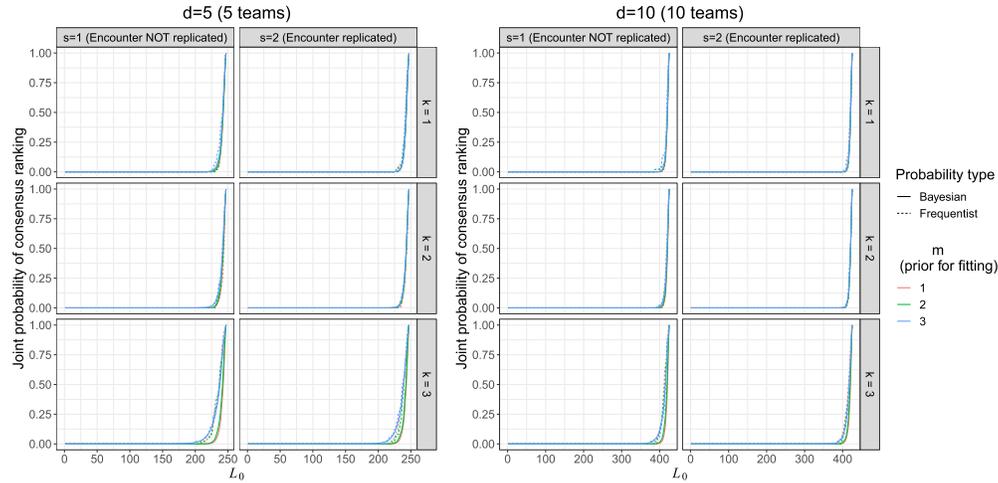
Figure 3: Simulation results for joint CPCRs for lineups assessing the effect of prior distribution used for fitting (captured by m), prior distribution used for data generation (captured by k), number of parameters (captured by d), and doubling the sample size (captured by s) with $\boldsymbol{v} = (k, d, s, m)$. Each plot displays the average posterior probability and frequentist coverage of the joint CPCRs (i.e., $\bigcap_{l=L_0}^{T}[\tau_{(1:l)}$ ranked 1, 2, ..., or $l]$). These plots summarize results over 50 synthetic datasets.

statements that are less sparse and have higher posterior probability and smaller local and global errors (see Figure 2 of the Supplementary Material).

We now analyze the results obtained under the CPCR methodology. As originally intended by Vitelli et al. (2018), the CPCR approach identifies, for each $l = 1, \ldots, T$, events of the form $[\tau_{(1:l)}$ ranked 1st, 2nd, ..., or lth] that concentrates the highest posterior probability (see Figure 3 in the supplementary material). Both the average posterior probability and the frequentist coverage for each CPCR are similar to each other across all combinations of simulation factors. The CPCRs representing the lineups with highest ability have probabilities less than or around 0.3, while CPCRs representing lineups that have the lowest ability have probabilities close to one (by design CPCRs for the worst lineups are more imprecise). As mentioned, global statements and CPCR are not directly comparable because the former provides joint statements and the latter provides a type of marginal statement. For this reason, we consider joint CPCRs defined as $\bigcap_{l=L_0}^{T}[\tau_{(1:l)}$ ranked 1st, 2nd, ..., or lth], $L_0 = 1, \ldots, T$. Figure 3 shows the effect that L_0 has on the probabilities for joint CPCRs. The larger L_0 (less lineups in the joint statement), the higher the probability. Notice how quickly these probabilities decay to zero. In fact, if the goal is to report a joint statement that concentrates a probability close to 0.9, we must select $L_0 = 243$ for $d = 5$ (5 teams = 247 lineups) and $L_0 = 422$ for $d = 10$ (10 teams = 425 lineups). For the case $L_0 = 243$ with $d = 5$, the most precise statement in $\bigcap_{l=243}^{247}[\tau_{(1:l)}$ ranked 1, 2, ..., or $l]$ states that lineup $\tau_{(1:243)}$ is ranked 1st, 2nd, ..., or 243th out of 247 lineups, which does not provide much information. No-
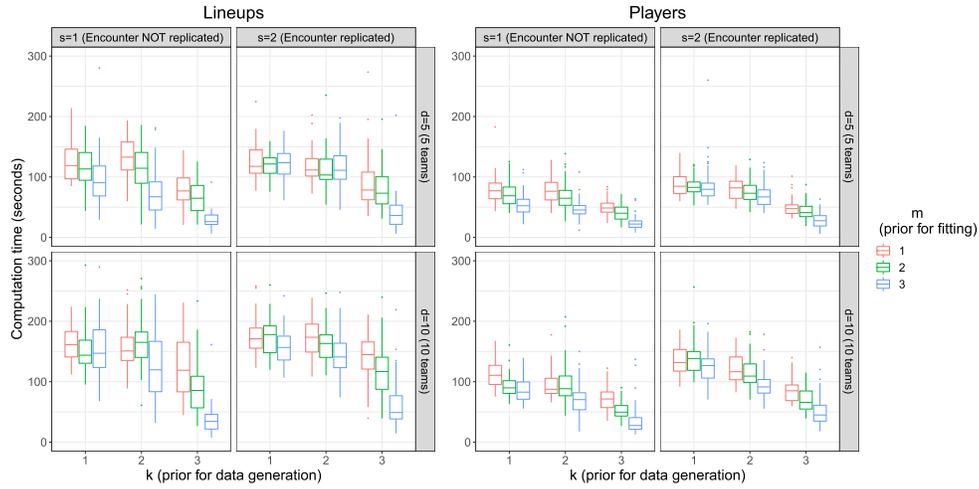
Figure 4: Wall clock times required to approximate optimal global statements without parallelisation.

tice that under the CPCR methodology the most informative statements are those for the best lineups suggesting we could alternatively consider joint statements of the form $\bigcap_{l=1}^{L_0}[\tau_{(1:l)}$ ranked 1, 2, ..., or $l]$, $L_0 = 1, \ldots, T$. Unfortunately, the probability of any of those joints statements will be upper bounded by the probability of $[\tau_{(1)}$ ranked first] which is, in the best case, around 0.3 as stated before. This strategy would limit the possibility of finding joint statements that concentrate high probability. The results for players under the CPCR methodology are similar to those obtained for lineups (see Figures 3 and 4 in the supplementary material).

The main conclusion from the simulation study is that the proposed method succeeds in finding statements having high posterior probability. The accuracy of the statements at containing the truth relies on the accuracy of the posterior distributions, and we observe that employing prior 3 produces global and local statements that maintain good frequentist properties at the cost of producing conservative statements. On the other hand, priors 1 and 2 produce statements that are less conservative, but are also more inaccurate. Thus if one does not have prior information regarding the similarity of true parameters, the recommended option is to use priors that lead to more conservative statements. If one indeed has prior information that the true parameters are relatively dissimilar, more informative priors (like prior 1) should be used as they lead to less sparse statements. We briefly mention here that the same conclusions where drawn when generating data that contained lineups that were not observed in the NBA data. Hypothetical lineups are clearly of interest to general managers and coaches as they work through the process of team building. Details are provided in Section 3 of the supplementary material.

To end this section, we provide the computational times required to optimize the global statements in Figure 4. While parts of the algorithm for finding the optimal local

and global statements can be parallelised to speed things up, we do not incorporate this into our code. Even without doing this, it takes less than two minutes to approximate the optimal global statement.

5   Application: National Basketball Association lineup data

The 2009–2010 NBA season consisted of $N = 30{,}807$ total encounters with 10,942 unique lineups based on 30 teams and 510 players. For each encounter, the points scored by each lineup is recorded along with the five players that constituted each lineup. For simplicity, we treat players that change teams during the season as different players; such changes mid-season affect a small number of players. We fit the model described in Section 4.1 using prior 3. This prior was used as the simulation study indicated that it should be used in the absence of knowledge regarding the similarity of true parameters. Similarly to Section 4.2, posterior sampling was carried out using NUTS implemented in Stan. We generated 13 independent Markov chain Monte Carlo (MCMC) chains. The initial 2,000 iterates for each chain was discarded and each chain thinned by 20 until a total of 10,000 combined posterior draws were collected. Then our method was applied to player and lineup abilities based on the posterior distributions of the player abilities.

Instead of optimizing over the set of all possible global statements, in applications one may want to impose additional constraints in order to obtain global statements with desired characteristics. Our framework allows this by simply letting us fix certain parameters while optimizing over the others. Setting the parameter t to be very small leads to many narrow local statements (that is, lineup/player statements that contain few competing lineups/players). This can be seen from the form of the function (3.6). A similar effect occurs at the global level if q is very small. To illustrate the impact that t and q have on the resulting local and global statements, we consider $t = q \in \{0, 0.1, 0.25\}$. Unless otherwise stated, we use a grid of 21 equally spaced values between 0 and 0.05 for α and γ to obtain global statements; our simulations have indicated that the results do not change by using a finer grid.

5.1   Ordering statements for lineups

We detail results for lineups in this section. For ease of illustration, we only consider lineups from teams that made the playoffs during the season under consideration. This subset of the data constitutes 16 teams, 261 players, 5,121 lineups, and 15,884 encounters.

Figure 5 displays the global statement for the 5,121 lineups based on $t = q = 0.1$ and $t = q = 0.25$. It took about eighteen minutes to compute this. The posterior probability of the global statement using $t = q = 0.1$ is 0.95 and using $t = q = 0.25$ is 1.0. As expected, each lineup's local statement for $t = q = 0.1$ is sparser in the sense that they contain less competing lineups compared to lineup local statements for $t = q = 0.25$. To illustrate this the red points in the left plot of Figure 5 correspond to lineups whose local statements are empty. That is, these lineups are no better or worse than any other
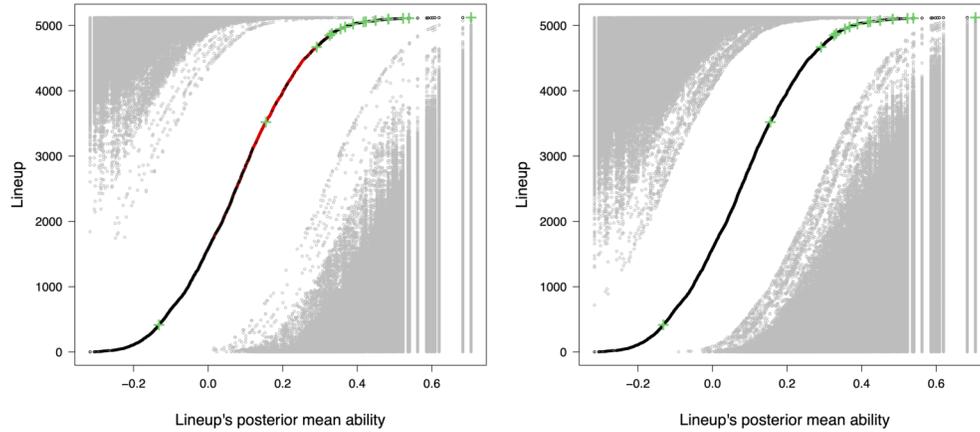
Figure 5: Visual depiction of global statement based on the 5,121 lineups from the 16 playoff teams. Left and right panels display global statements setting (α, t, γ, q) equal to $(0.025, 0.1, 0.05, 0.1)$ and $(0.05, 0.25, 0.05, 0.25)$, respectively. Each graph should be interpreted vertically, where each column of dots represents a particular lineup and its corresponding local statement. Red dots represent lineups to which no comparisons with any other lineups can be made. Gray dots below the black one define lineups that are worse than the corresponding lineups and those above define the group of lineups that are better. Green "+" points correspond to the lineups listed in Table 1.

lineup. On the other hand, all lineup local statements based on $t = 0.25$ are non-empty. This added "information" comes at the cost of accuracy as more errors are permitted for each lineup. We therefore recommend exploring global statements for a range of values for t and q, and making decisions depending on the amount of risk that a decision maker is willing to take.

The lineup that was better than the greatest number of other lineups is from the Orlando Magic and is comprised of Matt Barnes, Vince Carter, Dwight Howard, Rashard Lewis, and Jameer Nelson. This lineup was better than 4,803 of the 5,121 lineups and worse than none in their local statement for $t = q = 0.1$, and they were better than 5,019 lineups and worse than none in their local statement for $t = q = 0.25$. Table 1 contains the "best" lineups (that is, those whose local statement contains the highest number of competing lineups for which they are better) extracted from the global statement for each of the playoff teams under both $t = q = 0.25$ and $t = q = 0.1$. Interestingly, moving from $t = 0.1$ to $t = 0.25$ results in a larger increase of lineups "below" than lineups "above" for each team. Each of the lineups could be considered to be their teams "best". To visualize where each lineup in Table 1 is located in the global statement, they are highlighted by green colored "+" points in Figure 5. It should be clear why, for example, Chicago's lineup is not better than any other lineup, as it is found in the lower tail of lineup abilities. It may be tempting at first glance to use the information like that provided in Table 1 to rank order lineups. Doing so however, is problematic for two reasons. First, it is not clear how to rank lineups that have empty local statements (i.e.,

		$t = q = 0.1$		$t = q = 0.25$	
Team	Lineup	Below	Above	Below	Above
ATL	Bibby, Horford, Johnson, Josh Smith, Williams	2241	0	3230	0
BOS	Garnett, Perkins, Pierce, R. Allen, Rondo	1696	0	2624	1
CHA	Augustin, Diaw, Jackson, Murray, Wallace	3	9	51	26
CHI	Alexander, D. Brown, Law, Pargo, Richard	0	89	0	332
CLE	James, Jamison, Parker, Varejao, West	2505	0	3473	0
DAL	Dampier, Kidd, Marion, Nowitzki, Terry	1428	0	2347	0
DEN	Andersen, Anthony, Billups, Nene, Smith	555	1	1205	2
LAL	Artest, Bryant, Bynum, Fisher, Gasol	2940	0	3818	0
MIA	Alston, Haslem, ONeal, Richardson, Wade	639	0	1370	1
MIL	Bogut, Delfino, Ilyasova, Jennings, Salmons	168	1	645	2
OKC	Collison, Durant, Krstic, Sefolosha, Westbrook	890	0	1723	1
ORL	Barnes, Carter, Howard, Lewis, Nelson	4803	0	5019	0
PHX	Dudley, Frye, Nash, Richardson, Stoudemire	948	1	1774	2
POR	Aldridge, Batum, Camby, Miller, Roy	1410	0	2418	0
SAS	Duncan, Ginobili, Jefferson, McDyess, Parker	2168	0	3184	0
UTA	Boozer, Kirilenko, Korver, Millsap, Williams	452	1	1094	2

Table 1: Results when comparing lineups from all playoff teams. Each team's listed lineup corresponds to that which was better than the highest number of lineups among all playoff teams. The columns "below" and "above" refer to the number of lineups worse and better, respectively, than the stated lineup across all playoff teams.

those that are no better or worse then any other lineup). Second, it is possible that there exist two lineups whose local statements are the same size but are comprised of different lineups. Here too, it is unclear how to rank the two lineups.

It is somewhat challenging to summarize all the information found in Figure 5. From a coaching perspective, focusing specifically on their team's lineups would be more useful as this would allow coaches to more easily identify their team's best lineups and game plan for a particular opponent. To illustrate this, we extract the 305 local lineup statements from the global statement for the Utah Jazz and display them in Figure 6. We note that all local statements are guaranteed to have posterior probability at least that of the global statement. As before, setting $t = q = 0.1$ results in local statements that are sparser relative to those for $t = q = 0.25$, and the red points found in the left plot of Figure 6 correspond to lineups whose local statements are empty.

Unsurprisingly, for both $t = q = 0.1$ and $t = q = 0.25$, Deron Williams was a member of all the "best" Utah Jazz lineups. Perhaps somewhat surprisingly, Kyle Korver rather than Carlos Boozer (the supposed 2nd best player on the Jazz) was featured more regularly in the "better" Jazz lineups. Very surprisingly, Kyrylo Fesenko (a seldom used player) was a regular member of the Jazz's "better" lineups. In fact the lineup with the largest estimated ability is comprised of Carlos Boozer, Kyrylo Fesenko, Andrei Kirilenko, Kyle Korver, and Deron Williams. This lineup played sparingly (a total of five encounters), but has a local statement that is not empty (better than 305 other lineups
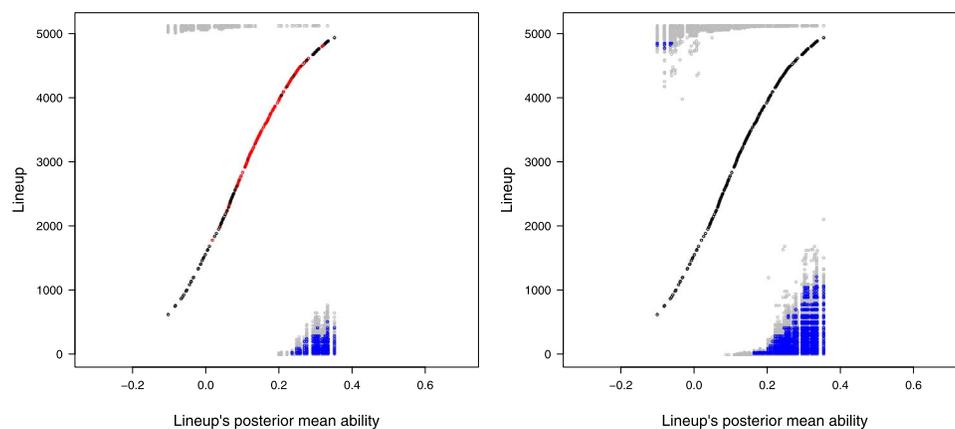
Figure 6: Graphical display of the 301 lineups that the Utah Jazz employed during the 2009–2010 NBA season. Blue dots correspond to lineups from the Denver Nuggets. Left and right panels display global statements setting (α, t, γ, q) equal to $(0.025, 0.1, 0.05, 0.1)$ and $(0.05, 0.25, 0.05, 0.25)$, respectively. Each graph should be interpreted vertically, where each column of dots represents a particular lineup and its corresponding local statement. Red dots represent lineups to which no comparisons with any other lineups can be made. Gray dots below the black one define lineups that are worse than the corresponding lineups and those above define the group of lineups that are better.

most of which didn't contain Deron Williams). Thus, it seems that the Utah Jazz coach should have explored lineups that contained Fesenko more. Contrast this to the lineup of Andrei Kirilenko, Kyle Korver, Wesley Matthews, C. J. Miles and Deron Williams which played more regularly and whose estimated ability is similar (the red dot with the largest estimated ability in left plot of Figure 6). However, the local statement associated with this lineup is empty. The reason behind this is the uncertainty associated with the lineup ability estimates. The posterior standard deviation of the latter is larger than the former. Remaining silent about this lineup highlights a novel feature of our approach. It turns out that the Utah Jazz lineup with the local statement containing the highest number of competing lineups consists of Carlos Boozer, Andrei Kirilenko, Kyle Korver, Paul Millsap, and Deron Williams. This lineup was arguably the strongest for the Utah Jazz during the 2009–2010 season.

It is valuable for a coach to know which of their team's lineups should compete against particular lineups of the opposing team. In the first round of the playoffs of the 2009–2010 NBA season, the Utah Jazz played against the Denver Nuggets. Each blue point in Figure 6 corresponds to a particular lineup from the Denver Nuggets. Thus we see that the "best" Jazz lineups are better than a number of Denver Nugget lineups for $t = q = 0.1$. However, it is only when we allow more uncertainty regarding the accuracy of the local lineups that we see some Denver Nugget lineups that are better than any Utah Jazz lineups.

5.2   Ordering statements for players

Ideally, global statements would consist of dense local statements that have high posterior probability and small values for α, γ, t and q. Unfortunately, this turns out to be impossible to achieve when comparing abilities of all 510 players. Figure 7 of the Supplementary Material displays the global statement of all 510 players for $t = q = 0.1$ and $t = q = 0.25$. It took under a minute to compute these global statements. The posterior probability is above 0.95 in both cases, but local statements are very sparse, thus limiting the usefulness of the global statement. This is a consequence of the uncertainty associated with player ability estimates (see Figure 6 of the online Supplementary Material) and producing a global statement based on all 510 players. Each additional player included in the global statement introduces more uncertainty which propagates to the local statements.

That being said, comparing all 510 players is typically not of interest as players can take on vastly different roles within a team framework. To further illustrate our method, we apply our procedure to a subset of players whose usage rate (an estimate of the percentage of a team's possessions ended by a particular player's action) was in the top twenty among all players in the 2009–2010 season. This permits us to compare players with similar roles across teams. As a brief aside, we remark that player abilities are estimated borrowing information across players that play on the same team. As a result, players who played on bad teams (for example, Golden State Warriors) were often members of lineups that lost many encounters (that is, negative difference between points scored and points allowed), and their ability was estimated to be generally lower than players on strong teams. Thus, small contributions to winning lineups result in a higher estimated ability relative to larger contributions to losing lineups.

The posterior distributions for players' abilities whose usage rate belonged to the top twenty during the 2009–2010 season are provided in the top-left plot of Figure 7. For this subset of players, we fix $t = q = 0$ as this will permit us to use the local statements to produce multiple orderings that can all be interpreted similarly to traditional rankings. Note that unlike traditional rankings, the probability that all these rankings hold simultaneously is equal to the posterior probability of the global statement. Further, the probability of any subset of these rankings holding simultaneously is lower bounded by the posterior probability of the global statement. To rank players using the global statement, we only need to look at the inherent transitivity of the local statements. That is, if player 1 is better than players 2 and 3, and player 3 better than player 2, then we can conclude that player 1 is better than player 3 is better than player 2 (i.e., $\xi_2 < \xi_3 < \xi_1$). Although the transitivity of the local statements can be used to create rankings of subsets of players when either $t \neq 0$ or $q \neq 0$ (that is, local and global errors are allowed), we refrained from exploring them as it is not evident how to quantify the inherited error for each of the derived rankings.

The global statement in the top-right plot of Figure 7 has a posterior probability of 0.87. This statement is sparse and only produces rankings that involve at most two players (for example, $\xi_{\text{James}} > \xi_{\text{Ellis}}$ and $\xi_{\text{Wade}} > \xi_{\text{Ellis}}$, et cetera.). Even though it is sparse, we can still conclude that the players with empty local statements (e.g., players with red dots "POR_Roy" – "CHI_Rose") are essentially exchangeable in terms of the
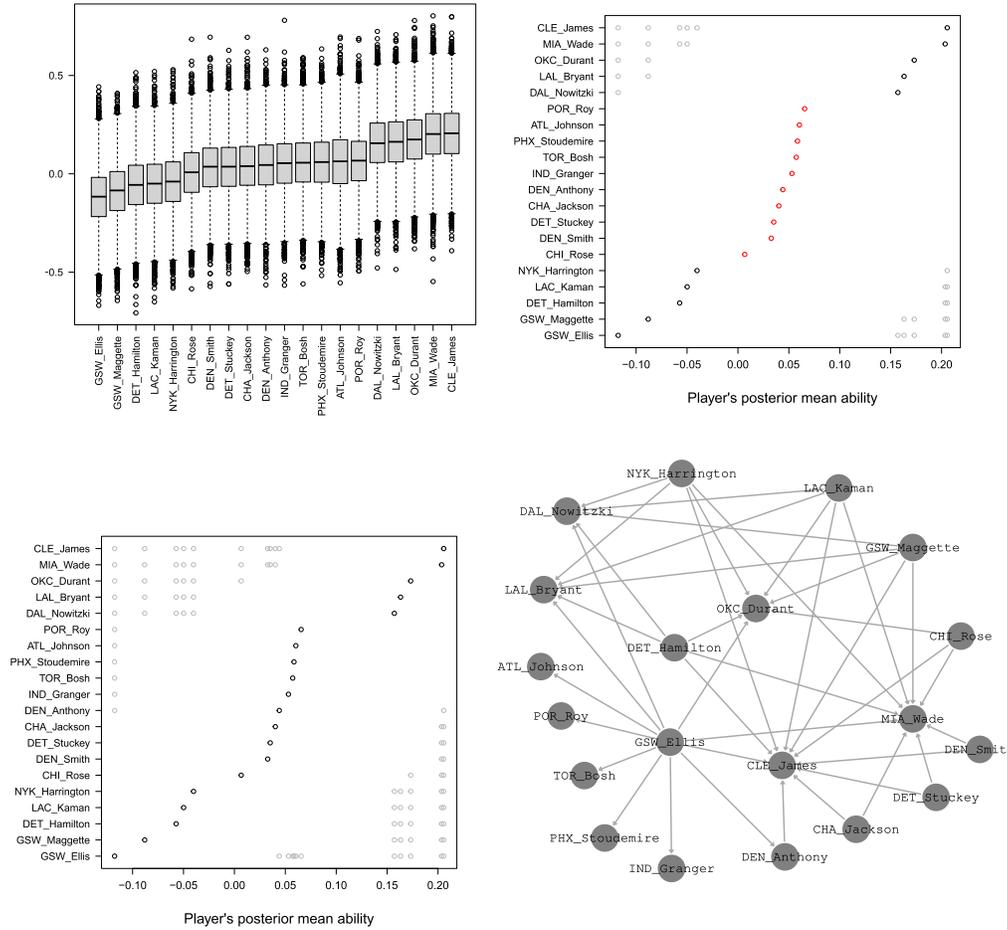
Figure 7: Results from players that were in the top twenty in terms of usage rate. The top-left plot corresponds to the posterior distributions of the player abilities for each of the twenty players. The top-right and bottom-left plots display the global statements setting (α, t, γ, q) equal to $(0.0225, 0, 0.05, 0)$ and $(0.09, 0, 0.9, 0)$, respectively. These two plots can be interpreted similarly as Figure 5. The bottom-right plot displays the global statement in the bottom left-plot using a network graph. Arrows point to nodes (players) that are better than those represented by the respective nodes.

their abilities. However, their salaries are very different ranging from 1 million to 16 million USD a year. Of course some of the discrepancy in salary is due to experience in the league. But even players like Stephen Jackson and J. R. Smith who play very similar roles, using a similar style of play, and with a similar number of years in the league, differ in salary by more than 2 million USD. Thus, our approach can provide general managers added information when making personnel decisions by indicating which players are very

similar in their contributions to lineup success. The global statement provided in the bottom left of Figure 7 is less sparse with a posterior probability of 0.34 and, given that this probability is low, we provide it more to contrast the sparser global statement. In this global statement, Lebron James and Dwayne Wade were better than 10 and 9 players, respectively, and worse than none. Monta Ellis, on the other hand, was worse than 11 players and better than none.

To help visualize the relationships between players, the global statement can be plotted using a network graph. The bottom right plot of Figure 7 displays a network graph that corresponds to the global statement in the bottom-left plot. Each node corresponds to a player. The network graph is directed in that edges connecting two players indicate which player is better by way of an arrow. Thus, any path of the graph represents a ranking of the nodes (that is, players). From the network graph, its clear that Monta Ellis is ranked as the worst player among the twenty, as many arrows are pointing away from him. Also notice that Monta Ellis is ranked below Carmelo Anthony, who is ranked below Lebron James. So, we can conclude that $\xi_{\text{James}} > \xi_{\text{Anthony}} > \xi_{\text{Ellis}}$. Similarly, we can conclude that $\xi_{\text{Durant}} > \xi_{\text{Rose}}$ and $\xi_{\text{Bryant}} > \xi_{\text{Hamilton}}$, but also that Bryant is neither better nor worse than James, et cetera. Note that all these ranking statements, or any combination of rankings from the global statement, have a joint posterior probability that is bounded below by 0.34 (which is the posterior probability of the global statement).

Twenty players are displayed in the top left plot of Figure 7, where the order of the players on the x-axis corresponds to their median ability. Given that no ranking of the twenty players appears more than once among the 10^4 MCMC iterates, any point estimate of the ranking will have posterior probability close to zero. Even if focus is placed on rankings that only require James, Wade, and Durant to be in the top three regardless of order, the posterior probability is still only 9×10^{-2}. Similarly, the posterior probability of rankings where Ellis, Maggette, and Kaman are in the bottom three regardless of order is 10^{-2}, and the posterior probability of rankings where both James, Wade, and Durant are in the top three and Ellis, Maggette, and Kaman are in the bottom three simultaneously is 7×10^{-4}. In contrast, our global statement communicates a similar message about these six players but with a much higher probability (3.4×10^{-1} compared to 2.5×10^{-3}). Notice that our global statement indicates that, for example, James, Wade, and Durant are the players who have the largest number of players that have a lower ability (see left bottom plot of Figure 7). Further, the global statement is saying that these three players are better than the seven players below Smith, and not that they have an ability that is consistently greater than the ability of all other 17 players.

6   Discussion

We have presented a novel procedure to find statements associated with orderings of parameters (for example, player and lineup abilities). Relying on a decision theoretic framework, the proposed approach generates statements associated with high posterior probability while incorporating as many players/lineups into the statements as possible. This represents a new approach that is quite different from existing ranking methods.

Although we use a prior distribution that assigns zero probability to ties among parameters in the numerical illustrations, we can incorporate ties in the statements by using prior specifications that assign positive probability to them, (for example, spike-and-slab priors of Ishwaran and Rao, 2005), or by defining ties based on practical criteria (for example, a tie is present if the maximum difference between parameters is below a given threshold). Moreover, it would be interesting to consider informative priors for player abilities. A possibility is to modify the prior mean for ξ_l in Section 4.2 to μ_l, with μ_l depending on the official NBA ranking of the players in the previous year (that is, have a different prior mean for each player ability).

From an NBA general manager's perspective, it may seem unsatisfactory that the global statements are silent about some players. This could be particularly true when employing the lineup data to make personnel decisions (for example, deciding between which free-agents to pursue). In this setting, our procedure can focus specifically on the small subset of players of interest. Posterior probabilities of the form $\mathsf{P}(\xi_l > \xi_{l'} \mid \text{data})$ can be useful for deciding whether to trade players. Since there is often a monetary cost involved in trading, this can be incorporated as well to calculate the expected gain in trading players.

Supplementary Material

Supplementary Material for "Bayesian Inferences on Uncertain Ranks and Orderings: Application to Ranking Players and Lineups" (DOI: 10.1214/22-BA1324SUPP; .pdf). Contains details associated with the algorithms employed to find local and global statements and optimal global statements in additionn to further results associated with the simulation study with regards to player and lineup abilities and posterior abilities of all the 510 players in the 2009–2010 NBA season as well as the global statement for all the players for $t = q = 0.1$ and $t = q = 0.25$.

References

Agresti, A. (2002). *Categorical Data Analysis*. John Wiley & Sons, 2 edition. MR1914507. doi: https://doi.org/10.1002/0471249688. 778

Barrientos, A. F., Sen, D., Page, G. L., and Dunson, D. B. (2022). "Supplementary Material for "Bayesian Inferences on Uncertain Ranks and Orderings: Application to Ranking Players and Lineups"." *Bayesian Analysis*. doi: https://doi.org/10.1214/22-BA1324SUPP. 785

Bradley, R. A. and Terry, M. E. (1952). "Rank analysis of incomplete block designs: I. The method of paired comparisons." *Biometrika*, 39(3/4): 324–345. MR0070925. doi: https://doi.org/10.2307/2334029. 778, 790

Caron, F. and Doucet, A. (2012). "Efficient Bayesian inference for generalized Bradley–Terry models." *Journal of Computational and Graphical Statistics*, 21(1): 174–196. MR2913362. doi: https://doi.org/10.1080/10618600.2012.638220. 778

Caron, F., Teh, Y. W., and Murphy, T. B. (2014). "Bayesian nonparametric Plackett–Luce models for the analysis of preferences for college degree programmes." *The Annals of Applied Statistics*, 8(2): 1145 – 1181. MR3262549. doi: https://doi.org/10.1214/14-AOAS717. 778

Cattelan, M., Varin, C., and Firth, D. (2013). "Dynamic Bradley–Terry modelling of sports tournaments." *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 62(1): 135–150. MR3042315. doi: https://doi.org/10.1111/j.1467-9876.2012.01046.x. 777

Chen, W., Liu, T.-Y., Lan, Y., Ma, Z.-M., and Li, H. (2009). "Ranking measures and loss functions in learning to rank." In *Advances in Neural Information Processing Systems*, 315–323. 779

Deshpande, S. K. and Jensen, S. T. (2016). "Estimating an NBA player's impact on his team's chances of winning." *Journal of Quantitative Analysis in Sports*, 12(2): 51–72. 782

Dittrich, R., Hatzinger, R., and Katzenbeisser, W. (1998). "Modelling the effect of subject-specific covariates in paired comparison studies with an application to university rankings." *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 47(4): 511–525. MR1833276. doi: https://doi.org/10.1111/1467-9876.00232. 778

Duane, S., Kennedy, A. D., Pendleton, B. J., and Roweth, D. (1987). "Hybrid Monte Carlo." *Physics Letters B*, 195(2): 216–222. MR3960671. doi: https://doi.org/10.1016/0370-2693(87)91197-x. 791

Gelman, A. (2006). "Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper)." *Bayesian Analysis*, 1(3): 515–534. MR2221284. doi: https://doi.org/10.1214/06-BA117A. 790

Hall, P. and Miller, H. (2010). "Modeling the variability of rankings." *The Annals of Statistics*, 38(5): 2652–2677. MR2722452. doi: https://doi.org/10.1214/10-AOS794. 778

Hans, C. (2009). "Bayesian lasso regression." *Biometrika*, 96(4): 835–845. MR2564494. doi: https://doi.org/10.1093/biomet/asp047. 790

Hans, C. (2010). "Model uncertainty and variable selection in Bayesian lasso regression." *Statistics and Computing*, 20(2): 221–229. MR2610774. doi: https://doi.org/10.1007/s11222-009-9160-9. 790

Henderson, N. C. and Newton, M. A. (2016). "Making the cut: improved ranking and selection for large-scale inference." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(4): 781–804. MR3534350. doi: https://doi.org/10.1111/rssb.12131. 778, 779, 780

Hoffman, M. D. and Gelman, A. (2014). "The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo." *Journal of Machine Learning Research*, 15(1): 1593–1623. MR3214779. 791

Huang, T.-K., Lin, C.-J., and Weng, R. C. (2008). "Ranking individuals by group comparisons." *Journal of Machine Learning Research*, 9(Oct): 2187–2216. MR2452625. doi: https://doi.org/10.1145/1143844.1143898. 778, 789

Hunter, D. R. (2004). "MM algorithms for generalized Bradley-Terry models." *The Annals of Statistics*, 32(1): 384 – 406. MR2051012. doi: https://doi.org/10.1214/aos/1079120141. 778

Ishwaran, H. and Rao, J. S. (2005). "Spike and slab variable selection: frequentist and Bayesian strategies." *The Annals of Statistics*, 33(2): 730–773. MR2163158. doi: https://doi.org/10.1214/009053604000001147. 803

Jewett, P. I., Zhu, L., Huang, B., Feuer, E. J., and Gangnon, R. E. (2019). "Optimal Bayesian point estimates and credible intervals for ranking with application to county health indices." *Statistical methods in medical research*, 28(9): 2876–2891. MR4000201. doi: https://doi.org/10.1177/0962280218790104. 778, 779, 780

Kalman, S. and Bosch, J. (2020). "NBA lineup analysis on clustered player tendencies: A new approach to the positions of basketball & modeling lineup efficiency of soft lineup aggregates." In *14th Annual MIT Sloan Sports Analytics Conference*. 782

Klein, M., Wright, T., and Wieczorek, J. (2020). "A joint confidence region for an overall ranking of populations." *Journal of the Royal Statistical Society Series C*, 69(3): 589–606. MR4098963. 778

Laird, N. M. and Louis, T. A. (1989). "Empirical Bayes ranking methods." *Journal of Educational Statistics*, 14(1): 29–46. MR3983320. doi: https://doi.org/10.1214/19-STS705. 777, 780

Lin, R., Louis, T. A., Paddock, S. M., and Ridgeway, G. (2006). "Loss function based ranking in two-stage, hierarchical models." *Bayesian Analysis (Online)*, 1(4): 915. MR2282211. doi: https://doi.org/10.1214/06-BA130. 777, 779, 780

Marriott, M. (2017). "Caste ranking and food transactions: a matrix analysis." In *Structure and Change in Indian society*, 133–172. Routledge. 777

Millot, B. (2015). "International rankings: Universities vs. higher education systems." *International Journal of Educational Development*, 40: 156–165. 777

Mohamad, D. A., Zwet, E. V., Solari, A., and Goeman, J. (2021). "Simultaneous Confidence Intervals for Ranks Using the Partitioning Principle." *Electronic Journal of Statistics*, 15: 2608–2646. MR4255330. doi: https://doi.org/10.1214/21-ejs1847. 778

Page, G. L., Fellingham, G. W., and Reese, C. S. (2007). "Using Box-Scores to Determine a Position's Contribution to Winning Basketball Games." *Journal of Quantitative Analysis in Sports*, 3(Article 1): 337–345. MR2350518. doi: https://doi.org/10.2202/1559-0410.1033. 782

Park, L. S., Yu, J., Yu, S., and Liao, L. (2015). "Search engine optimization for category specific search results." US Patent 9,116,994. 777

Park, T. and Casella, G. (2008). "The Bayesian lasso." *Journal of the American Statistical Association*, 103(482): 681–686. MR2524001. doi: https://doi.org/10.1198/016214508000000337.   790

Polson, N. G. and Scott, J. G. (2012). "On the half-Cauchy prior for a global scale parameter." *Bayesian Analysis*, 7(4): 887–902. MR3000018. doi: https://doi.org/10.1214/12-BA730.   790

Rendle, S., Freudenthaler, C., Gantner, Z., and Schmidt-Thieme, L. (2009). "BPR: Bayesian personalized ranking from implicit feedback." In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, 452–461. AUAI Press.   779

Rodríguez, A. and Moser, S. (2015). "Measuring and accounting for strategic abstentions in the US Senate, 1989–2012." *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 64(5): 779–797. MR3415950. doi: https://doi.org/10.1111/rssc.12099.   778

Rücker, G. and Schwarzer, G. (2015). "Ranking treatments in frequentist network meta-analysis works without resampling methods." *BMC Medical Research Methodology*, 15(1): 58.   777

Sill, J. (2010). "Improved NBA Adjusted +/- Using Regularization and Out-of-Sample Testing." In *MIT Sloan Sports Analytics Conference*, 1–7.   780

Soliman, M. A. and Ilyas, I. F. (2009). "Ranking with uncertain scores." In *2009 IEEE 25th International Conference on Data Engineering*, 317–328. IEEE.   778

Stan Development Team (2018). "RStan: the R interface to Stan." R package version 2.18.2. URL http://mc-stan.org/.   791

Vitelli, V., Sørensen, Ø., Crispino, M., Frigessi Di Rattalma, A., and Arjas, E. (2018). "Probabilistic preference learning with the Mallows rank model." *Journal of Machine Learning Research*, 18(158): 1–49. MR3813807.   778, 779, 791, 792, 794

Wright, D. L., Stern, H. S., and Cressie, N. (2003). "Loss functions for estimation of extrema with an application to disease mapping." *Canadian Journal of Statistics*, 31(3): 251–266. MR2030123. doi: https://doi.org/10.2307/3316085.   778

Wright, T., Klein, M., and Wieczorek, J. (2018). "A primer on visualizations for comparing populations, including the issue of overlapping confidence intervals." *The American Statistician*, 1–14. MR3953629. doi: https://doi.org/10.1080/00031305.2017.1392359.   778

Xie, M., Singh, K., and Zhang, C.-H. (2009). "Confidence intervals for population ranks in the presence of ties and near ties." *Journal of the American Statistical Association*, 104(486): 775–788. MR2541594. doi: https://doi.org/10.1198/jasa.2009.0142.   778