Döpke, Jörg; Köhler, Tim; Tegtmeier, Lars

Are they worth it? – An evaluation of predictions for NBA 'Fantasy Sports'

Journal of Economics and Finance

Provided in Cooperation with:
Springer Nature

Suggested Citation: Döpke, Jörg; Köhler, Tim; Tegtmeier, Lars (2023) : Are they worth it? – An evaluation of predictions for NBA 'Fantasy Sports', Journal of Economics and Finance, ISSN 1938-9744, Springer US, New York, NY, Vol. 48, Iss. 1, pp. 142-165, https://doi.org/10.1007/s12197-023-09646-7

Journal of Economics and Finance (2024) 48:142–165
https://doi.org/10.1007/s12197-023-09646-7

Are they worth it? – An evaluation of predictions for NBA 'Fantasy Sports'

Jörg Döpke[1] · **Tim Köhler**[1] · **Lars Tegtmeier**[1]

Accepted: 11 September 2023 / Published online: 23 October 2023

Abstract

'Fantasy Sports' - an internet-based game in which participants chose virtual teams of real professional athletes - has recently gained in popularity. Various firms provide projections regarding athletes' future performance to help participants choose their virtual teams. We evaluate such forecasts based on 1658 projections regarding NBA basketball of four selected projection providers that were collected in February 2022. We calculate standard measures of forecast quality and find that the use of professional forecasts reduces the errors made in naïve forecasts, but only to a moderate extent. Applying regression-based tests of forecast efficiency, we find that the predictions are inefficient and, in some cases, even biased. Third, pairwise comparisons of the accuracy of the providers suggest notable differences among such providers in the short run. We use a simple optimization algorithm to choose a virtual team for each match day and feed it with the forecasts of the providers. Subsequently, we rank the providers according to the score obtained by these teams. We find small, although in one case significant, long-run differences between the providers, among whom each provides better accuracy than that of a naïve projection based on these athletes' past performances. Finally, we simulate one-on-one competition among various forecast providers to ascertain the long-term profitability of their services. Given the small magnitude of the detected differences, our results, in brief, raise doubts as to whether the forecasts provided are worth the money.

Keywords Forecast evaluation · Information efficiency · Sport

JEL Classification C53 · G14 · L83

✉ Tim Köhler
tim.koehler@hs-merseburg.de

Jörg Döpke
joerg.doepke@hs-merseburg.de

Lars Tegtmeier
lars.tegtmeier@hs-merseburg.de

[1] University of Applied Sciences Merseburg, Eberhard-Leibnitz-Straße 2, D-06217 Merseburg, Germany

1 Introduction

The statistical modelling of sports data has become an important research field in recent times. In particular, great efforts have been made in predicting the outcomes of sports events based on statistical models (Stekler et al. 2010; Štrumbelj and Vračar 2012; Manner 2016; Winston et al. 2012). Sport predictions are important to team managers, sponsors, sports fans, and the media (Spann and Skiera, 2009; Song and Shi 2020). One of the main drivers of the relevance of sports forecasting is the vast and ever-growing sports betting market (Francisco and Moore 2019; Wunderlich and Memmer 2020). Bookmakers and professional gamblers need powerful forecasting models to gain a competitive advantage over others (Goddard 2005; Koopman and Lit 2019; McHale and Morton 2011). In this context, fantasy sports as an entertainment and betting platform has also become an integral part of the sports environment and has been developed into a large industry (Kotrba 2020; Haugh and Singal 2021).There are approximately 60 million participants in the U.S. and Canada alone (Fantasy Sports & Gaming Association (FSGA) 2021). In terms of revenue, the market size of the fantasy sports services industry in the U.S. was $8.4 billion in 2021 (IBISWorld 2021). The two largest daily fantasy sports (DFS) providers, DraftKings and FanDuel, have a market share of approximately 90% between them (Easton and Newell 2019; Haugh and Singal 2021). In addition, DraftKings has been publicly traded since 2020 and achieved revenues of $615 million in 2020 (DraftKings 2022).

Several websites offer projections on the performance of athletes, i.e. projections of how many fantasy points a certain athlete will score on a given match day. In addition, these websites frequently offer an optimization model that uses the projections to build a team of athletes. In doing so, the model is constrained to obey the rules regarding team building and maintenance. The model result represents the optimal team according to the projections. Provided all of these projections were correct, the participant applying them would select the best team and thus win the competition. In reality, however, the projections would, of course, differ from the actual results. A provider whose projection error is substantially lower than that of the competition could give the participants leveraging this information a competitive advantage. The magnitude of a firms' projection error is, therefore, of central importance. For this reason, we examine the projection errors of four major analysis websites for DFS below. We focus on fantasy basketball in the National Basketball Association (NBA) using its so-called 'classic' version at the provider level (DraftKings 2022).

We examine the number and distribution of projections and the respective errors in the projections of each provider. We also test for forecast errors as indicated by zero means and medians. After calculating the forecast errors relative to the errors made by a naïve forecast, we compare the accuracy of the projections, both among each other and against the naïve projection, using the Diebold and Mariano (1995) statistic.

Relying on regression-based tests, we analyse whether the projections are efficient and unbiased. For this purpose, we investigate whether the providers have accounted for information available before the match day in their projections. In addition, we use the Stekler (1987) statistic to test for long-run differences in accuracy between

the providers. We use an optimization model to rank the best teams by relying on the projections and create a ranking for each match day. We also use the optimization model to simulate a one-on-one competition between different forecasting providers to determine the long-term profitability of their services.

Our main results show that the use of professional forecasts reduce the forecast errors made by naïve forecasts (the athletes' average fantasy points of the last five games), but only moderately (by less than 10%). Second, regression-based tests conducted to estimate forecast efficiency indicate, for some providers, inefficient predictions, that do not fully take into account the available information at the time of the forecasts. The results of conducting pairwise Diebold and Mariano (1995) tests show some noteworthy differences in forecast quality among providers. Finally, the Stekler (1987) statistics and the simulated one-on-one competition show that there are only minor differences between the various forecasters. Nevertheless, the use of forecasts from all forecasters outperformed the use of the naïve forecasting method. Given that Daily Fantasy Fuel provides its forecasts free of charge and that none of the paid forecasters perform significantly better, the value proposition involving the investment of money in forecasting services is questionable.

We organize the paper as follows. Section 2 gives some information on fantasy sports and reviews the related literature. Section 3 describes the criteria used to evaluate the projections. Section 4 presents the empirical results, starting with descriptive statistics, followed by tests for forecast efficiency, and concluding in a discussion of the long-term differences between forecasts. Finally, we conclude and discuss directions for further research.

2 Fantasy sports and the related forecast literature

2.1 A primer for the rules of the game

Daily fantasy sports is a game in which participants can take on the role of an owner or manager of a sports team. As such, they select real athletes from professional sports teams to be a player on their fantasy team. There are restrictions on the part of DFS companies regarding team composition, which may vary depending on the company, the sport, and the type of competition at play. These restrictions relate to the maximum number of each type of athlete, the number of different teams represented in a line-up, and other aspects described in greater detail below. Each athlete is also assigned a (hypothetical) salary; additionally, the total salary for each team must be lower than a specified salary cap. Athletes accumulate fantasy points during a sporting event based on their real performances on the field. Fantasy points for each athlete are based on the key statistics for that sport (for example, points or assists) and the athlete's position on the team. A fantasy team's total points equals the sum of each selected athlete's fantasy points. The participants are ranked on the basis of their respective teams. For most contests, the higher the rank achieved (or the passing of a specific threshold in the rankings), the more the prize money earned (South et al. 2019 p. 180).

2.2 Fantasy sports as an object of research

Fantasy sports represent a highly topical and interesting field of research. Statistical modelling and analyses related to the prediction of the performances of individual athletes, the formation of optimal teams, and the prediction of fluctuating athlete values based on real-world sports data have been widespread (see, for example, Sargent and Bedford 2010; Agarwal et al. 2017; Müller et al. 2017). Despite the popularity of fantasy sports, however, there has been comparatively little research on it in the scientific literature.

An example of scientific studies predicting individual athlete performance and optimal team formation in fantasy sports is the study of Fry et al. (2007). In this study a stochastic dynamic programming (DP) model is proposed for the athlete selection in a single National Football League (NFL) franchise. The best draft choice in each round is determined by DP recursion, which maximizes the sum of the values of the drafted athlete and the total expected value for the team in future rounds. To obtain a computationally tractable model, some simplifying assumptions are introduced to remove the stochastic component from the model (which is mainly the element of uncertainty regarding the behaviour of the opposing teams) and reduce the size of the state space.

Matthews et al. (2012) developed algorithms to predict the performance of soccer teams and athletes specifically for the Fantasy Premier League. They used reinforcement learning techniques to generate score predictions thereby deriving individual performance forecasts for athletes. They also used mixed-integer programming techniques to optimize the selection of athlete transfers over several weeks. Their approach outperformed the predictions of humans in 99% of the cases in both simulations and practice.

Bonomo et al. (2014) studied a fantasy sport game for the first Argentine soccer league. They presented two mathematical programming models, the forst was designed a priori and referred to as 'prescriptive', while the other was formulated a posteriori that was referred to as 'descriptive'. Both are used to identify optimal – or, at least, good – teams for a fantasy play. The descriptive model is able to find the ideal teams for achieving the highest possible score while meeting all the constraints imposed by the rules of the game. The model was also used to analyse different ways of defining the initial team in the prescriptive model, the different athlete formations allowed by the game, and the scalability of the models in terms of the number of rounds and solution times. The prescriptive model uses historical data and the characteristics of the next round of play to create a competitive team that is then tested in six fantasy tournaments. The model results placed the team in the top 0.1% of game participants in one of the tournaments (Opening 2010), in the top 4% in four other tournaments, and in the top 10% in the remaining cases.

Becker and Sun (2016) suggested a methodology to predict team and athlete performances, and developed a mixed-integer optimization model that uses such predictions for draft selection and weekly line-up management. Numerical tests of the model indicate a promising level of performance. South et al. (2019) suggested a complete system for Daily Fantasy Basketball that includes both athlete performance prediction

and team composition. They used a Bayesian random effects model to predict an aggregate measure of daily NBA athlete performance. These predictions were then used to construct teams under the constraints of the game. Next, permutation-based and K-nearest neighbours approaches are used to identify more successful teams ot those that would have been competitive more often than others based on their historical data. These predictions were then compared with those from the analysis website www.numberfire.com and daily competitions were simulated throughout the 2015–2016 season. The results od these simulations showed an expected profit of approximately $9,000 on an initial investment of $500 using the K-nearest neighbours approach, which represents a 36% increase over the use of a permutation-based approach alone.

The contribution made by South et al. (2019) comes closest to our paper. However, significant differences remain. First, South et al. (2019) focus on predicting "an aggregate measure of the daily performance of NBA players", whereas our paper aims to evaluate the forecasts of the third-party providers. Consequently, South et al. (2019) rely on only a single measure for forecast quality, do not make an explicit comparison with a naïve forecast, and do not refer to theoretical concepts such as "rational forecasts". On the other hand, South et al. (2019) provide an overview and introduction to the world of fantasy sports, and make great efforts to obtain the best possible team given a set of data, which is beyond the scope of our paper.

Beal et al. (2020) proposed several new models and algorithms for solving team line-up problems in DFS. They focused on predicting the performance of NFL athletes to form the optimal fantasy team through the use of mixed-integer programming. They tested their solutions on datasets from four seasons (2014-2017), showing that their solutions outperformed existing benchmarks and produce a win in up to 81.3% of DFS game weeks in a season. Kotrba (2020) used OLS regression models to investigate a heuristic strategy for DFS Premier League squad selection in the 2015-2016 season. The results showed that participants selected their squads based primarily on the past performances of athletes, in addition to considerations of their favourite team. By applying the heuristic, participants try to simplify their decision-making, but tend to overestimate the influence of an athlete belonging to a favoured team. Furthermore, he demonstrated that the influence of the betting odds on individual athletes' scores is statistically significant.

Haugh and Singal (2021) developed a coherent framework for DFS portfolio construction in which they explicitly modelled the behaviour of other DFS participants. They formulated an optimization problem that accurately describes the DFS problem for a risk-neutral decision-maker using both the double-up and top-heavy pay-off settings. Their formulation maximizes the expected reward subject to feasibility constraints, and they related this formulation to mean-variance optimization and the outperformance of stochastic benchmarks. In addition, they introduce a Dirichlet-multinomial data generation process for modelling opposing team choices. This allowed them to estimate the value of both insider trading and collusion in a DFS setting. They demonstrated the value of their framework by applying it to DFS contests during the 2017 National Football League season.

To the best of our knowledge, we are the the first to use a broader range of statistical measures of forecasting accuracy – thereby extending the existing literature – to examine the predictions of athlete points and related optimized line-ups based on numerous websites that specialize in DFS analysis. Such information is estimated to be used by 30% of the DFS participants, and they spend over $250 million annually to purchase such information and decision-making tools to gain a competitive advantage in these competitions Easton and Newell (2019). We focus on the DFS NBA Classic basketball league under DraftKings, as since basketball is one of the most studied sports due to its popularity and the large number of games per season (Sarlis and Tjortjis 2020; Song and Shi 2020).

3 Evaluation criteria

3.1 Measures and tests for accuracy

To assess the quality of the forecasts, we begin by applying standard statistical measures of forecast accuracy. Let FP_i denote the fantasy points (FP) score of an athlete i allocated by the fantasy league after the actual game has taken place. In addition, FP_i^e represents the FP of the same athlete forecasted *before* the event by the forecast provider. Hence, for a certain individual, the forecast error is defined as $e_i = \mathrm{FP}_i^e - \mathrm{FP}_i$. Consequently, a positive forecast error corresponds to an overestimation of the athletes' performance and *vice versa*.

We check whether the forecasters show some kind of bias by calculating the mean error (ME):

$$\mathrm{ME} = \frac{1}{T}\sum_{t=1}^{N} e_t \tag{1}$$

We assess the magnitude of errors by the mean absolute error (MAE):

$$\mathrm{MAE} = \frac{1}{T}\sum_{t=1}^{N} |e_t| \tag{2}$$

The MAE implicitly assumes a linear loss function. However, it seems plausible that forecasters may use a quadratic loss function, which punishes large forecast errors more heavily than small ones. To take this into account, we calculate the root mean squared error (RMSE):

$$\mathrm{RMSE} = \sqrt{\frac{1}{T}\sum_{t=1}^{T} e_t^2} \tag{3}$$

From the participant's perspective, it is crucial whether a forecast is worth the money or, at least, the time spent to notice them, is a crucial determination. In other words, should the participants rely on other, and in particular, cheaper forecasts, rather than listening to the advice of the firms? The arguably cheapest forecasting method

is a naïve forecast. In this method, the participant might stick to the prediction of the average of the last five match days. To assess whether the provided forecasts are at least as good as their naïve counterparts, we apply relative measures of forecast accuracy. First, we calculate Theil's inequality coefficient, which compares the mean error of the forecasts of interest with the respective number of naïve forecast:

$$U = \frac{\frac{1}{T}\sum_{t=1}^{N} e_t^2}{\frac{1}{T}\sum_{t=1}^{N}(FP_t - FP_{t-1})^2} \tag{4}$$

The closer to zero that Theil's inequality coefficient (U)) is, the more accurate the forecast is compared to a naïve forecast. A value higher than one – in contrast – implies that the naïve forecast has greater accuracy.

A related statistic is the mean absolute scaled error of Hyndman and Koehler (2006):

$$\text{MASE} = \frac{\frac{1}{T}\sum_{t=1}^{T} |e_t|}{\frac{1}{T-1}\sum_{t=1}^{T} |FP_i - FP_{t-1}|} \tag{5}$$

This number also includes the relationship with the naïve forecast as a benchmark model but with the benefit a lower level of susceptibility to data. Again, a value under one indicates a more accurate forecast than that of the benchmark model. The difference between U and the MASE reflects the distinction between a quadratic loss function and a linear loss function. Hence, in the presence of some large forecast errors, the two numbers may substantially differ.

To formally test for a different type of quality of the predictions, we use the Diebold and Mariano (1995) test. We calculate, based on the forecast error of one provider ($e_{\text{Provider 1}}$) and the forecast errors of another provider ($e_{\text{Provider 2}}$), the so-called loss differential based on squared errors ($\overline{d}_t = e^2_{\text{Provider 1}} - e^2_{\text{Provider 2}}$). The loss differential is subsequently used to build the test statistic (Diebold and Mariano 1995 p. 3):

$$DM = \frac{\overline{d}}{\sqrt{[\gamma_0 + 2\sum_{k=1}^{h-1}\gamma_k]/n}} \tag{6}$$

where γ_k represents the autocovariance at lag k. Under the null hypothesis of no difference in forecast accuracy, the statistic follows a standard normal distribution.

3.2 Tests for rational forecasts

A substantial number of forecast evaluation techniques refer to the idea of a so-called 'rational' forecast. Against this background, we test whether the providers deliver such projections. The notion of a rational forecast is usually split into three features (see, for example, Stekler 2002):

- A rational forecast should be unbiased. To test for this property, the Mincer and Zarnowitz (1969) regression is a natural starting point. In our case, we estimate

the outcome of an athlete's strength as a function of the predicted score:

$$\mathrm{FP}_i = \beta_0 + \beta_1 \mathrm{FP}_i^e + \epsilon_i \tag{7}$$

We further test for the hypothesis:

$$H_0 = \begin{cases} \beta_0 = 0 \\ \beta_1 = 1 \end{cases} \tag{8}$$

If the data reject this hypothesis, the prediction is not rational, since there is a systematic difference between the forecast and the outcome.

- A rational forecast should be weakly efficient. Since the forecast errors of the last period are known, when forming the next expectation, prior forecast errors should not provide any information on the subsequent error. We use a variant of the test that was suggested by Holden and Peel (1990), who indicate estimating the equation:

$$e_i = \gamma_0 + \gamma_1 e_{i,t-1} + \epsilon_i \tag{9}$$

and testing for

$$H_0 = \begin{cases} \gamma_0 = 0 \\ \gamma_1 = 0 \end{cases} \tag{10}$$

This approach is used to test, as the Mincer and Zarnowitz (1969) regression, for unbiasedness ($H_0 : \gamma_0 = 0$) and any information content of the lagged forecast errors for the recent ones ($H_0 : \gamma_1 = 0$). If this hypothesis cannot be rejected, then the prediction is called weakly efficient.

- Finally, a rational forecast should be strongly efficient. In other words, no exogenous information available to the forecasters prior to the forecast should contain any information relevant for the forecast error. We again use the Holden and Peel (1990) equation:

$$e_i = \gamma_0 + \gamma_1 e_{i,t-1} + \gamma_2 X_{t-1} + \epsilon_i \tag{11}$$

and test for

$$H_0 = \begin{cases} \gamma_0 = 0 \\ \gamma_1 = 0 \\ \gamma_2 = 0 \end{cases} \tag{12}$$

where X stands for any information that is available as of the forecasting date.

3.3 Use of linear programming to optimize fantasy sports picks

A participant can choose a line-up for a match day using an optimization model, which some providers offer to their customers. Since the problem is not overly complicated, we expect no quality differences among providers regarding the optimization model. Optimization models have also been developed for different types of competitions in the literature (see, for example, Fry et al. 2007; South et al. 2019). This section shows

the optimization model for the classic NBA on DraftKings (2022), which follows notations of Zamora (2022).

Our optimization model maximizes the forecasted FPs under the given constraints. Hence, if the FPs for each player were predicted correctly by the respective provider, the optimization model would give the best team for that particular match day.

Each player i has a salary c_i and an FP projection p_i. The decision regarding an athlete i is a binary decision: either this person is included in the team or not. In the first case, the variable x_i is set to the value of 1, it is set to 0 otherwise (Eq. 18). The goal is to maximize the FPs of the team, so the objective function is given by Eq. 13. In building the team, several restrictions must be taken into account.

$$\text{Max} \sum_i p_i x_i \tag{13}$$

$$\sum_i c_i x_i \leq 50000 \tag{14}$$

$$\sum_i x_i = 8 \tag{15}$$

$$\sum_{i \epsilon \text{PG}} x_i \geq 1; \sum_{i \epsilon \text{SG}} x_i \geq 1; \sum_{i \epsilon \text{SF}} x_i \geq 1; \sum_{i \epsilon \text{PF}} x_i \geq 1; \sum_{i \epsilon \text{C}} x_i \geq 1 \tag{16}$$

$$\sum_{i \epsilon \text{PG}} x_i \leq 3; \sum_{i \epsilon \text{SG}} x_i \leq 3; \sum_{i \epsilon \text{SF}} x_i \leq 3; \sum_{i \epsilon \text{PF}} x_i \leq 3; \sum_{i \epsilon \text{C}} x_i \leq 2; \tag{17}$$

$$x_i = \begin{cases} 1 \text{ if player i is selected for the Fantasy team} \\ 0 \text{ otherwise} \end{cases} \tag{18}$$

First, there is a limited budget of \$50000 (Eq. 14). Second, a total of exactly eight athletes i constitutes a line-up (Eq. 15). Each position[1] must be occupied by at least one player (Eq. 16). In addition, there are artificial player positions, such as like 'guard' (G), which is composed of 'point guard' (PG) and 'shooting guard' (SG), and 'forward' (F), which is composed of 'small forward' (SF) and 'power forward' (PF). The position 'utility' (U) taken by a player who is allowed to play in any position. The artificial positions lead to the fact that the positions 'guard' (G), 'shooting guard' (SG), 'small forwards' (SF), and 'power forwards' (PF) can be filled three times, and the position 'center' (C) can be filled twice (Eqs. 17).

3.4 Testing for long-term accuracy

To check for long-term accuracy of the forecasting accuracy of an institution, we refer the measure for long-run relative performance that was suggested by Stekler (1987). In the first step, a score (R_{it}) is assigned to every forecast, which takes the value of

[1] 'point guard' (PG), 'shooting guard' (SG), 'small forwards' (SF), 'power forwards' (PF), and 'center' (C).

the rank in accordance with to the respective criteria, e.g., the absolute forecast error. In the second step, the cumulated rank-sum of these scores is calculated:

$$S_i = \sum_{t=1}^{T} R_{it} \tag{19}$$

Under the null hypothesis that each institution has the same predictive capacity, each institution should have an expected cumulative sum of scores of:

$$S_i^e = \frac{T(N+1)}{2} \tag{20}$$

where T is the number of periods included and N is the number of forecasts considered. In our case, we have 18 match days and, thus, an expected value of $\frac{18\cdot(5+1)}{2} = 54$. To calculate the test statistic, we use the corrected standard deviation as proposed by Batchelor (1990):

$$\sigma = \sqrt{\frac{TN(N+1)}{12}} \tag{21}$$

4 Empirical results

4.1 Data and descriptive statistics

We collected data from four providers over the period spanning from 1 February 2022 to 28 February 2022. We compared our data on average results (i.e., the average actual fantasy points for each player) with the corresponding data for the entire 2021/22 season to ensure that our sample sufficiently represented the data set. To do this, we performed a Kolmogorov-Smirnov test on the null hypothesis that both data come from the same distribution. The test gives a p value of 0.17. Therefore, we cannot reject the null hypothesis. The selected forecasters offer their services at a range of prices. We have deliberately chosen these providers to represent different price segments. RotoGrinders is the most expensive service on our list, charging $64.99 per month, followed by Daily Fantasy Nerd at $29.99 per month and FantasyPros at $11.99 per month. Daily Fantasy Fuel, on the other hand, offers its predictions for free. It is worth noting that all of the providers offer various additional services that may provide justifications for their higher prices that go beyond the value of predictions. We lost some data due to technical problems. We omit all athletes who are known to have been injured before the match. Additionally, we filtered the data to account for the fact that some athletes were not considered in the match at all, resulting in a game time of precisely zero. To make the forecasts comparable, we have only considered those athletes for whom a score has been forecasted by each of the providers. Table 1 gives an overview of the data[2]. The standard deviation and range of the actual score are much higher than those of the projected scores. This property of the predictions is

[2] All calculations have been performed using R (R Core Team 2022).

reasonable since a rational forecast of a variable should have a lower variance than its actual values (see, for example, Lovell 1986).

These results are also illustrated in Fig. 1. Forecasters do not seem to predict extremely low or high values even though these values do in reality (Panel A of the exhibit). From Panel B of Fig. 1 shows that the forecast errors are more or less normally distributed. Since we have a relatively large sample, this assumption has to be further verified by a formal test.

4.2 Forecast accuracy

Since high forecast accuracy gives participants a long term advantage, such accuracy is critical. Thus, we examine how providers vary in this respect. We explore, in a first step, the statistical measures of accuracy outlined above. As a second step, we turn to a more economical approach to forecast evaluation, i.e., we analyse the consequences of forecast errors in a more realistic setting. We assume that an optimal team for fantasy sport is set in accordance with the respective forecasts of any provider. In addition to the four projection providers, we have created a naïve forecast using the average DraftKings (2022) score of each athlete over the last five matches. We chose this window size because it is a common metric provided by most data providers and is easily accessible. To assess the robustness of this approach, we conducted tests using different averaging windows ranging from 1 to 6 match days and found that none of them significantly outperformed the 5-match-day average[3].

Table 2 lists are all forecast providers and their accuracy measures.

The mean error is close to zero for each forecast series and the use of a standard t test fails to reject the null hypothesis of a zero mean for each provider, implying that the forecasts are unbiased on average. However, there is one exception: FantasyPros has the largest deviation from zero the p value indicates the rejection of the null hypothesis of a zero mean at the $p = 0.05$ significance level. This result implies that FantasyPros on average significantly slightly overestimates the athletes' performance. Its forecasts are biased. We also test the hypothesis that the errors are normally distributed by a Shapiro–Wilk test. The null of a normal distribution is overwhelmingly rejected at significance levels smaller than $p = 0.05$ for all providers. Consequently, we refer to a test that does not rely on strong assumptions about the underlying distribution of the data. In particular, we calculate the median forecast errors and use a Wilcoxon test to verify the hypothesis of a zero median that does not have to be rejected for any provider.

According to the mean absolute error, Daily Fantasy Nerd delivers the best predictions. For all providers the MAE is slightly lower or higher than 8 FPs, which shows a nontrivial magnitude of projection errors as compared to the standard deviation of the actual outcomes of slightly more than 14 FPs (see Table 1). The root mean squared error – which gives more weight to large forecast errors than the absolute mean error

[3] For the sake of brevity, we do not present the related results in detail here. They are available from the authors upon request.

Table 1 Descriptive statistics for predicted and actual fantasy points of the athletes

	Number of observations	Mean	Standard deviation	Minimum	25 % Quartile	75 % Quartile	Maximum
Forecasts of:							
Daily Fantasy Nerd (2020)	1658	25.3	11.1	6.7	16.9	31.5	63.9
RotoGrinders (2022)	1658	25.5	11.5	2.4	17.1	31.9	65.6
Daily Fantasy Fuel (2022)	1658	25.5	10.8	3.3	17.3	31.5	61.8
FantasyPros (2022)	1658	25.0	10.8	2.5	17.2	30.6	66.2
Actual data DraftKings (2022)	1658	25.5	14.3	−0.5	14.8	34.5	84.3

Notes: Per the authors' calculations

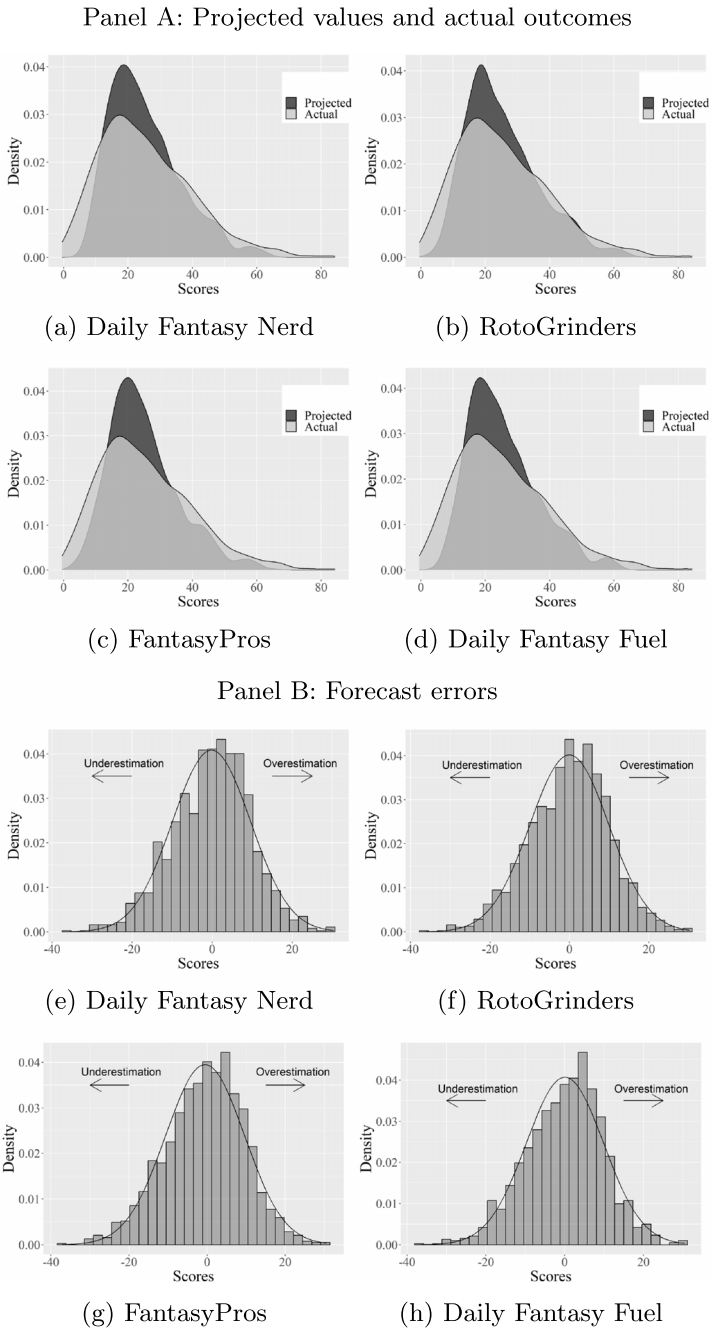

Fig. 1 Distribution of projected values, actual outcomes, and forecast errors for all providers. *Notes*: Per the authors' calculations. Kernel (Gaussian) smoothing with a smoothing parameter of $\alpha = 0.3$. The black line represents the normal distribution

Table 2 Absolute and relative accuracy measures by provider

	Daily Fantasy Nerd	RotoGrinders	Daily Fantasy Fuel	FantasyPros
	Absolute accuracy measures			
Mean error[a)] (in FP)	-0.17	-0.02	0.06	0.055**
Test for zero mean	(0.47)	(0.92)	(0.80)	(0.03)
Test for normality	(< 0.01)	(< 0.01)	(< 0.01)	(< 0.01)
Median error (in FP)	0.60	0.61	1.00	0.20
Test for zero median	(0.82)	(0.45)	(0.23)	(0.24)
Mean absolute error[b)]	7.72	7.87	7.79	8.03
Root mean squared error[c)]	9.76	9.92	9.80	10.11
	Relative accuracy measures			
Theils inequality coefficient[d)]	0.91	0.92	0.91	0.94
Mean absolute scaled error[e)]	0.91	0.93	0.92	0.95

Notes: Authors' calculations. [a)]: see Eq. 1. [b)]: see Eq. 2. [c)]: see Eq. 3. [d)]: see Eq. 4. [e)]: see Eq. 5. The p values of the receptive tests are displayed in parenthesis. The test for zero mean is a t test, the test for normality is a Shapiro–Wilk test, and the test for a zero median is a Wilcoxon rank test. ***,(**,*) denotes rejection of the null hypothesis at the 1 (5, 10) % level.

– shows an even larger magnitude of the average errors and a similar rankings of the providers' accuracy. Taken together, the providers show substantial forecast errors regarding individual athletes' performance. At first glance, however, they do not seem to be very different in terms of accuracy.

Theil's inequality coefficient and the mean absolute scaled error (see the lower part of Table 2) reveal how the providers perform against the benchmark of a naïve forecast. The one-step-ahead forecasts of each provider are better than our naïve forecast, since the respective numbers are lower than one. However, the units are close to one, which indicates that the providers are only slightly better than the naïve forecast.

Table 3 contains a pairwise comparison of the focal providers both among each other and against a naïve forecast (the average performance of an athlete during the last five match days). We report the relative mean squared error (MSE) of the forecasts, the Diebold and Mariano (1995) test statistic, and the respective p values of a test of the hypothesis test of equal forecast accuracy. To begin with good news for the providers, the relative MSE in the last row of the table, shows that all providers make projections that are substantially better than the naïve projection, which is reflected in the value of the relative MSE being greater than one. Additionally, the results of the Diebold and Mariano (1995) test reveal that the loss differential is positive, again supporting the better forecast of the providers, and that the accuracy difference is statistically significant is statistically significant ($p < 0.1$). By the same logic, the protections of the providers Daily Fantasy Nerd and Daily Fantasy Fuel are significantly better than those offered by other providers. The difference between them is very small and statistically nonsignificant. By this measure, RotoGrinders ranks third, and FantasyPros ranks fourth among the competing firms.

4.3 Test for bias and efficiency

Using the Mincer and Zarnowitz (1969) regression, we test whether the forecasts are biased. The Wald test, reported in Table 4, rejects the null of an unbiased forecast in three out of four cases. Only in the case of 'Daily Fantasy Fuel' is the null of an unbiased projection is not rejected by the data.

We also check if the forecast error of the previous match day (the lagged forecast errors) contains information that was not considered in the forecast. To this end, Table 5 reports the results of testing of Hypothesis 10 using Eq. 9. The lagged error does not contain any information that is not already included in the forecast. Therefore, the projections are at least weakly efficient for each provider.

Finally, Table 6 reports the results of a test of strong efficiency in the Holden and Peel (1990) test used to test strong efficiency,which focuses on the information efficiency of forecasts. This test is used to check whether any exogenous information available before a match day has been fully taken into account in the forecasts. If, by contrast, the projections could have been improved the inclusion of this information, they are regarded as inefficient by this criterion.

We consider two kinds of information in applying this test. First, we look at the 'past performance', i.e., the five-match-day average of each athlete. As is documented

Table 3 Pairwise relative mean squared errors and Diebold and Mariano (1995) tests

		Daily Fantasy Nerd	Roto-Grinders	Daily Fantasy Fuel	FantasyPros
RotoGrinders	Relative MSE	1.03			
	DM-statistic	3.38			
	p value	(< 0.01)			
Daily Fantasy Fuel	Relative MSE	1.01	0.98		
	DM-statistic	0.96	-2.27		
	p value	(0.34)	(0.02)		
FantasyPros	Relative MSE	1.07	1.04	1.06	
	DM-statistic	4.81	2.21	2.27	
	p value	(< 0.01)	(0.02)	(0.03)	
Naïve Forecast	Relative MSE	1.21	1.17	1.20	1.13
	DM-statistic	8.17	7.02	7.96	5.23
	p value	(< 0.01)	(< 0.01)	(< 0.01)	(< 0.01)

Notes: Authors' calculation. The first number in each row is the relative mean squared error, i.e. RMSE $= \frac{\text{MSE}_{row}}{\text{MSE}_{column}}$. A number of one indicates equal forecast accuracy. The number in parentheses represents the p value of a Diebold and Mariano (1995) test for equal forecast accuracy. Calculations have been conducted out with the Hyndman and Khandakar (2008) package

in the table, in two cases, the variable is significant ($p < 0.1$). Hence, these forecasts are not informationally efficient, since an information that is easily and virtually free of charge and available to each participant is not fully accounted for by the projection providers. Second, we use the so-called 'salaries', i.e., the virtual costs that occur when procuring a player, when an athlete is named for a fantasy team. The results presented in Table 6 show that in at least two cases, the variables impart additional information to that in the provider forecast. This may point to a situation, in which already the 'pricing' of the DraftKings (2022) webpage is already informationally efficient, and there are no gains beyond those that can be achieved.

4.4 Comparison of the teams chosen on the basis of each provider's projections

In this section, we assess, whether minor differences in forecast accuracy – as documented above – are relevant for the participants from the perspective of multiple match days. To shed light on this problem, we use the optimization model[4] described in Eqs. 13 to 18 above, to calculate the optimal team for each[5] match day and focal provider. First, we check whether the different projections offered by each provider for single athletes do, in fact, boil down to the different teams chosen by the optimization

[4] Calculations have been carried out with the package provided by Michel Berkelaar et al. (2022)

[5] Note that we had to leave out three match days due to data availability.

Table 4 Mincer and Zarnowitz (1969) regression by provider

	Daily Fantasy Nerd	RotoGrinders	Daily Fantasy Fuel	FantasyPros
Predicted Score from...	Dependent variable: Actual Score			
Daily Fantasy Nerd	0.945***			
	(0.022)			
RotoGrinders		0.904***		
		(0.021)		
Daily Fantasy Fuel			0.964***	
			(0.022)	
FantasyPros				0.936***
				(0.023)
Constant	1.568***	2.479***	0.847	2.135***
	(0.598)	(0.591)	(0.618)	(0.622)
Wald Test χ^2 value	7.00**	20.8***	2.60	12.7***
Wald Test (p value)	(0.03)	(< 0.01)	(0.27)	(< 0.01)
N	1658	1658	1658	1658
R^2	0.535	0.524	0.530	0.503

Notes: Authors' calculation. Standard errors are displayed in parentheses. *** (**, *) denotes the rejection of the null hypothesis of an unbiased forecast (Hypothesis 8 in Eq. 7) at the 1(5,10)% level

Table 5 Holden and Peel (1990) tests for weak efficiency by provider

Lagged forecast errors for...	Daily Fantasy Nerd	RotoGrinders	Daily Fantasy Fuel	FantasyPros
Daily Fantasy Nerd	0.025			
	(0.027)			
RotoGrinders		0.007		
		(0.027)		
Daily Fantasy Fuel			0.010	
			(0.027)	
FantasyPros				0.039
				(0.027)
Constant	-0.121	0.102	0.033	−0.505*
	(0.267)	(0.268)	(0.271)	(0.275)
N	1367	1367	1367	1367
R^2	0.001	0.00005	0.0001	0.001

Notes: Authors' calculation. Standard errors are displayed in parentheses. *** (**, *) denotes the rejection of the null hypothesis at the 1(5,10)% level. The F-Test refers to the hypothesis of weakly efficient forecast (Hypothesis 10 in Eq. 9)

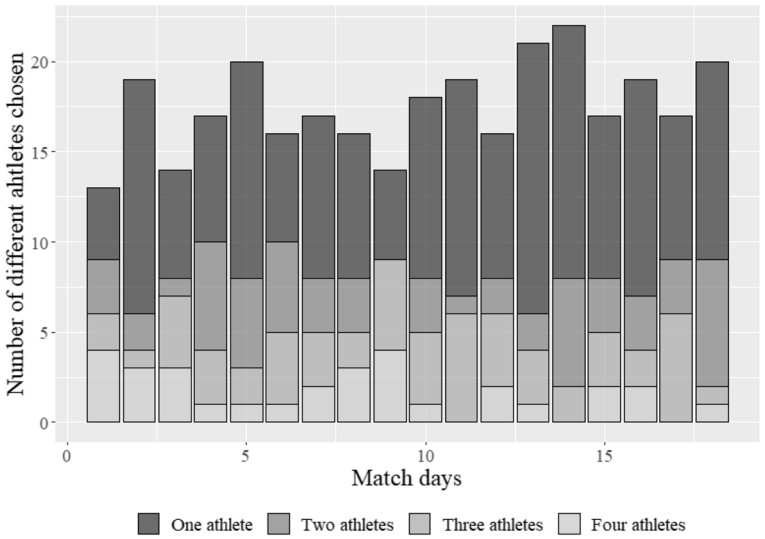

Fig. 2 Number of different athletes chosen based on providers' forecasts by match day. *Notes*: Authors' calculation

procedure. Figure 2 shows that the providers' projections indeed lead to quite different virtual teams. If all selections were completely distinct, than the four providers' forecasts would lead to the choice of 32 different athletes. Even though this is not the case, the overlap of teams remains reasonably small. On particular match day, for example, four athletes were part of the team recommended by each provider's forecasts, two athletes were in three of the teams, and so on.

Table 6 Holden and Peel (1990) tests for strong efficiency by provider

Lagged forecast errors for...	Daily Fantasy Nerd	RotoGrinders	Daily Fantasy Fuel	FantasyPros
	Based on past performance of athletes			
Daily Fantasy Nerd	0.031			
	(0.027)			
RotoGrinders		0.011		
		(0.028)		
Daily Fantasy Fuel			0.018	
			(0.027)	
FantasyPros				0.043
				(0.028)
Past performance	0.045**	0.026	0.073***	0.022
	(0.023)	(0.023)	(0.023)	(0.024)
Constant	−1.286**	-0.566	−1.842***	-1.070
	(0.645)	(0.651)	(0.653)	(0.666)
N	1367	1,367	1367	1,367
R^2	0.003	0.001	0.007	0.002
F-Statistic	2.372*	0.668	5.051***	1.435
	Based on salary for each player			
Daily Fantasy Nerd	0.025			
	(0.027)			
RotoGrinders		0.008		
		(0.027)		
Daily Fantasy Fuel			0.009	
			(0.027)	
FantasyPros				0.040
				(0.027)
Salary	0.0003**	0.0001	0.0004***	0.0001
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Constant	−1.565**	-0.689	−2.045***	-1.184
	(0.724)	(0.729)	(0.735)	(0.746)
N	1367	1367	1367	1367
R^2	0.004	0.001	0.007	0.002
F-Statistic	2.707*	0.715	4.691***	1.481

Notes: Authors' calculation. Standard errors are displayed in parentheses. *** (**, *) denotes the rejection of the null hypothesis at the 1 (5,10)% level. The F-test refers to the hypothesis a strongly efficient forecast (Hypothesis 12 in Eq. 11)

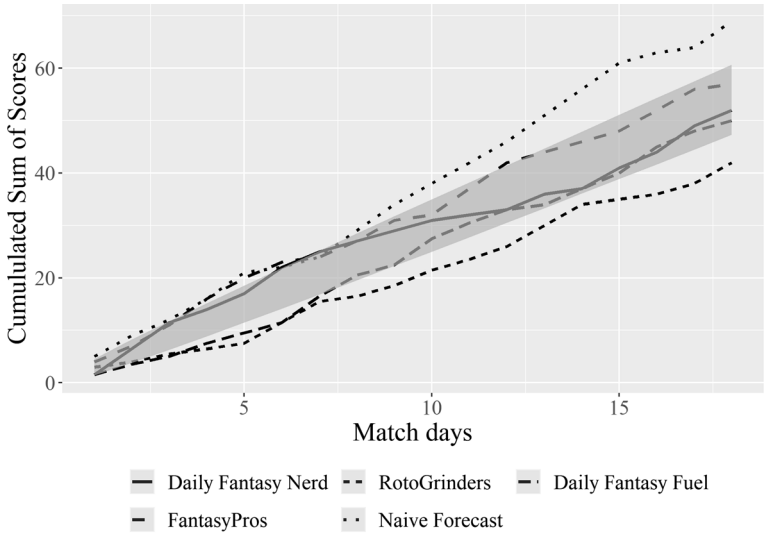

(a) Based on the best team only

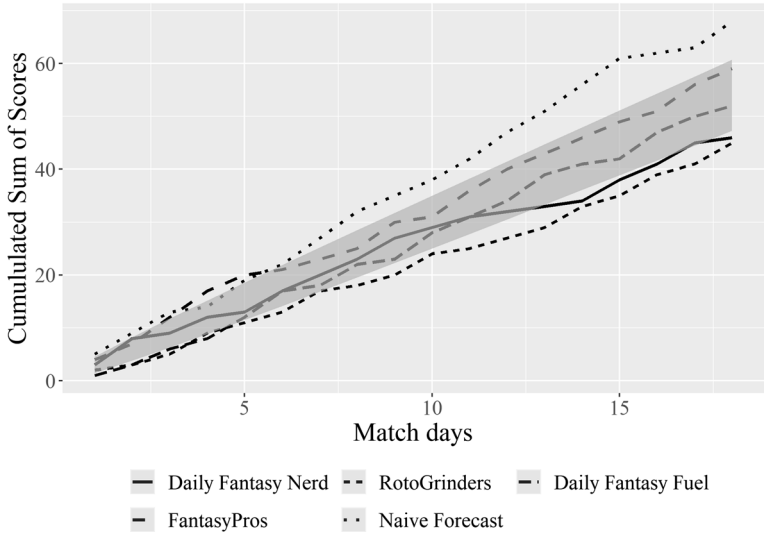

(b) Based on the average of the 3 best teams chosen

Fig. 3 Stekler (1987) cumulated rank sum test for long-term forecast accuracy. *Notes*: Authors' calculation. Calculated in accordance with Eq. 19. Shaded areas represent +/- one standard error based on the method proposed by Batchelor (1990); see Eq. 20.

Then, we use the approach proposed by Stekler (1987) to compare these teams according to the forecasts of the respective providers and then checked which team had scored the most points based on the actual results. In other words, we created a ranking for each match day and calculated the cumulative sum of the ranks of each provider.

The results are shown in Fig. 3. In the upper panel of the figure, the rankings for the best team in each case are summed. The vertical dark grey area in the figure represents the expected player performance with two standard deviations. If a provider is outside the dark grey area, it indicates that this provider performs significantly better or worse than the average. In our case, we have detected that the naïve forecast performs significantly worse and the RotoGrinders forecast performs significantly better than the other forecasts. The lower panel of Fig. 3 displays a similar comparison using the average rank of the best three teams for each provider.

4.5 One-on-one competition

In the following section, we analyse the profitability of prediction providers by simulating a one-on-one competition[6]. We follow this procedure because calculating a direct measure of profitability is somewhat difficult in the context of fantasy sports: First, there are many different tournaments with alternating rules and, second, the organizer, DraftKings, is not perfectly transparent. For example, as a rule, it is not clear from the outside (i.e., as a nonparticipant), what the necessary ranking to achieve is for getting "into the money". Thus, we adopt the so-called "head-to-head" tournament structure offered by Draftkings. In such games, for example, each participants places a $10 bet, of which Draftkings, as the host, collects a 10% fee from each player. Therefore, the winner gets $18 and the loser gets nothing. In the case of a tie, both players will receive $9 each. To not lose money in the long run, a player must achieve an average win rate of at least 55.56%, with a draw counting as half a win. The composition of the teams and the allocation of points is performed in the same way as described in the previous chapter. In the first phase, only the best-rated team from each forecaster competes over the 18 matchdays. Then, in the second stage, we repeated the experiment with the three best performing teams. As a result, we obtain 18 observations in the first stage, and 54 observations in the second stage.

Table 7 shows the results for each matchup, where all values correspond to first provider mentioned. The first column shows the percentage of games won over an 18 match day period. The second column shows whether the win rate from the first column is significantly greater than 55.56%. The results show that all forecasters outperform the naïve forecast and achieve a win rate of above 55.56%. However, in the first stage, only RotoGrinders' profit margin is statistically significant against the naïve forecast. In the second stage, all forecasters made a statistically significant profit when competing against the naïve forecast. In addition to the naïve forecast, FantasyPros also loses its statistical significance against two other forecasters. None of the forecasters were able to make a profit in the long run when competing against Daily Fantasy Fuel, which provides its forecasts for free.

[6] This line of reasoning might also be justified since it is similar to that in the literature on the economic (as opposed tom statistical) criteria of forecast accuracy. See, for example, Granger and Pesaran (2000).

Table 7 Winning rates and profits in a one-on-one competition

	Winning rate	Significance	Profit
	Best team		
RotoGrinders vs. Daily Fantasy Nerd	0.53	0.50	-27.00
Daily Fantasy Fuel vs. Daily Fantasy Nerd	0.53	0.50	-27.00
Daily Fantasy Nerd vs. FantasyPros	0.50	0.59	-18.00
Daily Fantasy Nerd vs. naïve forecast	0.67	0.24	36.00
RotoGrinders vs. Daily Fantasy Fuel	0.64	0.32	27.00
RotoGrinders vs. FantasyPros	0.67	0.24	36.00
RotoGrinders vs. naïve forecast	0.83	0.02	90.00
Daily Fantasy Fuel vs. FantasyPros	0.61	0.41	18.00
Daily Fantasy Fuel vs. naïve forecast	0.72	0.12	54.00
FantasyPros vs. naïve forecast	0.61	0.41	18.00
	Best three teams		
RotoGrinders vs. Daily Fantasy Nerd	0.52	0.83	-207.00
Daily Fantasy Nerd vs. Daily Fantasy Fuel	0.53	0.74	-81.00
Daily Fantasy Nerd vs. FantasyPros	0.60	0.15	126.00
Daily Fantasy Nerd vs. naïve forecast	0.68	0.00	378.00
RotoGrinders vs. Daily Fantasy Fuel	0.55	0.53	-18.00
RotoGrinders vs. FantasyPros	0.65	0.01	279.00
RotoGrinders vs. naïve forecast	0.78	0.00	657.00
Daily Fantasy Fuel vs. FantasyPros	0.63	0.04	216.00
Daily Fantasy Fuel vs. naïve forecast	0.71	0.00	54.00
FantasyPros vs. naïve forecast	0.64	0.02	18.00

Notes: Authors' calculations. The values correspond to the first supplier mentioned. The significance column indicates whether the values shown in the winning rate column are statistically significant and above 55.56%.

5 Conclusion

We have taken NBA basketball as an example for assessing the quality of the forecasts offered by fantasy sports forecast providers. These firms offer projections of athletes' performance to support the participants' selection of players for their virtual teams.

We begin our assessment with standard measures of forecast accuracy. The use of this method uncovers projection errors of substantial magnitude of slightly more than half of the standard errors of the actual outcomes. Moreover, the results show that professional forecasts reduce the forecast errors of naive forecasts (the athletes' previous score) only moderately (by less than 10%).

Second, a regression-based test for unbiased forecasts rejects the null hypotheses in two out of four cases. While we report no evidence of weakly inefficient forecasts, the hypotheses in relation to the concept of strong efficiency must be frequently rejected. This implies that the projections do not fully take into account the information that is available at the time of the forecasts. Third, the results of pairwise Diebold and Mariano (1995) tests show notable differences in forecast quality among the suppliers

in the short-term, i.e., based on the projection errors for athletes. However, using a simple optimization algorithm to choose a virtual team for each match day fed with the forecasts of the competing providers shows only small differences among them. All of them, however, outperform a naive selection of the athletes based on past performance. A similar result is observed in the simulated one-on-one competition.

Overall, our results cast doubt on the value of investing in paid predictions, as Daily Fantasy Fuel, which is a free service, performs comparably to paid services with no significant differences. Consequently, a natural question for further research might be as follows: Why do the participants still spend money for the services of the providers? One possible explanation for this behaviour is that paid providers may offer additional features, such as text-based news updates or integrated optimization models, which players can access to quickly set up their teams. To test these ideas, however, would require data on the participants and their bets, which are - to the best of our knowledge - not publicly available. In the context of further research, it would be interesting to develop a prediction model based on the betting odds for individual players (e.g., points, rebounds, assists, steals, blocks). For example, a large number of scientific studies reveal the highly predictive power of betting odds in real datasets for different sports (e.g. Forrest et al. (2005); Kovalchik (2016); Štrumbelj and Vračar (2012)). However, due to a lack of data on individual NBA players, this is not currently possible. Should the betting market develop further, this would be an interesting area for further research. Another line of future research might ask, how inefficiencies among the providers relate to similar findings for other sport betting related markets (see, for example, Francisco and Moore 2019).

Acknowledgements The authors thank participants of the 40^{th} EBES Conference held in Istanbul from July 6^{th} to 8^{th} 2022, for helpful comments on a previous draft of this paper. Data availability statement: The dataset generated during the current study is not publicly available as it contains information that the authors partly purchased from the projection providers. Information on how to obtain it and reproduce the analysis is available from the corresponding author on request.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declaration of interest

The authors have no relevant financial or nonfinancial interests to disclose

References

Agarwal S, Yadav L, Mehta S (2017) Cricket team prediction with hadoop: statistical modeling approach. Proc Comput Sci, 122:525–532

Batchelor RA (1990) All forecasters are equal. J Bus & Econ Stat, 8(1):143–144

Beal R, Norman TJ, Ramchurn SD (2020) Optimising Daily Fantasy Sports teams with artificial intelligence. Int J Comput Sci Sport, 19(2):21–35
Becker A, Sun XA (2016) An analytical approach for fantasy football draft and lineup management. J Quant Anal Sports, 12(1):17–30
Bonomo F, Durán G, Marenco J (2014) Mathematical programming as a tool for virtual soccer coaches: a case study of a fantasy sport game. Int Trans Oper Res, 21(3):399–414
Daily Fantasy Fuel (2022): Homepage. (last access: 5/25/2023). https://www.dailyfantasyfuel.com/
Daily Fantasy Nerd (2022): Homepage, (last access: 5/25/2023). https://dailyfantasynerd.com/
Diebold FX, Mariano RS (1995) Comparing pedictive accuracy. J Bus & Econ Stat, 13(13):235–265
DraftKings (2022): Homepage, (last access: 5/25/2023). https://www.draftkings.com/
Easton T, Newell S (2019) Are daily fantasy sports gambling? J Sports Anal, 5(1):35–43
Fantasy Sports & Gaming Association (FSGA) (2021): (last access: 5/25/2023). https://thefsga.org/
FantasyPros (2022): Hompage, (last access: 10/5/2022). https://www.fantasypros.com/
Forrest D, Goddard J, Simmons R (2005) Odds-setters as forecasters: The case of English football. Int J Forecast, 21(3):551–564
Francisco J, Moore E (2019) Betting with house money: reverse line movement based strategies in college football totals markets. J Econ & Finance, 43(4):813–827
Fry MJ, Lundberg AW, Ohlmann JW (2007): A player selection heuristic for a sports league draft, J Quant Anal Article 5. Sport 3 (2), Article 5
Goddard J (2005) Regression models for forecasting goals and match results in association football. Int J Forecast, 21(2):331–340
Granger C, Pesaran M (2000) Economic and statistical measures of forecast accuracy. J Forecast, 19(7):537–560
Haugh MB, Singal R (2021) How to play fantasy sports strategically (and win). Manag Sci, 67(1):72–92
Holden K, Peel DA (1990) On testing for unbiasedness and efficiency of forecasts. Manch Sch, 58(2):120–127
Hyndman RJ, Khandakar Y (2008): Automatic time series forecasting: the forecast package for R, (last access: 5/25/2023). https://www.jstatsoft.org/article/view/v027i03
Hyndman RJ, Koehler AB (2006) Another look at measures of forecast accuracy. Int J Forecast, 22(4):679–688
IBISWorld (2021): Fantasy sports services in the US - Market size 2005–2026, (last access: 5/25/2023). https://www.ibisworld.com/industry-statistics/market-size/fantasy-sports-services-united-states/
Koopman SJ, Lit R (2019) Forecasting football match results in national league competitions using score-driven time series models. Int J Forecast, 35(2):797–809
Kotrba V (2020) Heuristics in Fantasy sports: is it profitable to strategize based on favourite of the match? Mind & Soc, 19:195–206
Kovalchik SA (2016) Searching for the GOAT of tennis win prediction. J Quant Anal Sport, 12(3):127–138
Lovell MC (1986) Tests of the rational expectations hypothesis. The Am Econ Rev, 76(1):110–124
Manner H (2016) Modeling and forecasting the outcomes of NBA basketball games. J Quant Anal Sports, 12(1):31–41
Matthews T, Ramchurn S, Chalkiadakis G (2012): Competing with humans at fantasy football: Team formation in large partially-observable domains, in: Proceedings of the AAAI Conference on Artificial Intelligence, ed. by Association for the Advancement of Artificial Intelligence, 26:1394–1400
McHale I, Morton A (2011) A Bradley-Terry type model for forecasting tennis match results. Int J Forecast, 27(2):619–630
Michel Berkelaar et al. (2022): ipsolve: Interface to Lp_solve v. 5.5 to solve Linear Integer Programs, (last access: 5/25/2023). https://cran.r-project.org/web/packages/lpSolve/index.html
Mincer JA, Zarnowitz V (1969): The evaluation of economic forecasts, in: Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance, by Mincer JA, Zarnowitz V (ed.) NBER, 3–46
Müller O, Simons A, Weinmann M (2017) Beyond crowd judgments: Data-driven estimation of market value in association football. Eur J Oper Res, 263(2):611–624
R Core Team (2022): R: A Language and Environment for Statistical Computing, (last access: 5/25/2023). https://www.R-project.org/
RotoGrinders (2022): Homepage, (last access: 5/25/2023). https://rotogrinders.com/
Sargent J, Bedford A (2010) Improving Australian Football League player performance forecasts using optimized nonlinear smoothing. Int J Forecast, 26(3):489–497

Sarlis V, Tjortjis C (2020) Sports analytics–Evaluation of basketball players and team performance. Inf Syst, 93 Art 101562
Song K, Shi J (2020) A gamma process based in-play prediction model for National Basketball Association games. Eur J Oper Res, 283(2):706–713
South C, Elmore R, Clarage A, Sickorez R, Cao J (2019) A Starting Point for Navigating the World of Daily Fantasy Basketball. Am Stat, 73(2):179–185
Spann M, Skiera B, (2009) Sports forecasting: a comparison of the forecast accuracy of prediction markets, betting odds and tipsters. J Forecast, 28(1):55–72
Stekler HO (1987) Who forecasts better? J Bus & Econ Stat, 5(1):155–158
– (2002): The rationality and efficiency of individuals forecasts, in: a companion to economic forecasting, by Clements M, Hendry David F (ed.), 222–240 Blackwell Oxford
Stekler HO, Sendor D, Verlander R (2010) Issues in sports forecasting. Int J Forecast, 26(3):606–621
Štrumbelj E, Vračar P (2012) Simulating a basketball match with a homogeneous Markov model and forecasting the outcome. Int J Forecast, 28(2):532–542
Winston WL, Nestler S, Pelechrinis K (2012) Mathletics. Princeton University Press
Wunderlich F, Memmert D (2020) Are betting returns a useful measure of accuracy in (sports) forecasting? Int J Forecast, 36(2):713–722
Zamora R (2022): coach – R package to optimize fantasy sports lineups for a number of sports, sites, and constraints., (last acces: 5/25/2023). https://rdrr.io/github/zamorarr/coach/