

A Market for Trading Forecasts: A Wagering Mechanism*

Aitazaz Ali Raja^{a,*}, Pierre Pinson^{b,c}, Jalal Kazempour^d, Sergio Grammatico^a

^a*Delft Center for Systems and Control, TU Delft, The Netherlands*

^b*Imperial College London, Dyson School of Design Engineering, London, United Kingdom*

^c*Technical University of Denmark, Department of Technology, Management and Economics, Kongens Lyngby, Denmark*

^d*Technical University of Denmark, Department of Wind and Energy Systems, Kongens Lyngby, Denmark*

Abstract

In many areas of industry and society, e.g., energy, healthcare, logistics, agents collect vast amounts of data that they deem proprietary. These data owners extract predictive information of varying quality and relevance from data depending on quantity, inherent information content and their own technical expertise. Aggregating these data and heterogeneous predictive skills, which are distributed in terms of ownership, can result in a higher collective value for a prediction task. In this paper, we envision a platform for improving predictions via implicit pooling of private information in return for possible remuneration. Specifically, we design a wagering-based forecast elicitation market platform, where a buyer intending to improve their forecasts posts a prediction task, and sellers respond to it with their forecast reports and wagers. This market delivers an aggregated forecast to the buyer (pre-event) and allocates a payoff to the sellers (post-event) for their contribution. We propose a payoff mechanism and prove that it satisfies several desirable economic properties, including those specific to electronic platforms. Furthermore, we discuss the properties of the forecast aggregation operator and scoring rules to emphasise their effect on the sellers' payoff. Finally, we provide numerical examples to illustrate the structure and properties of the proposed market platform.

Keywords: mechanism design, wagering mechanism, predictive distribution, elicitation of probabilities, value of forecast, scoring rules

1. Introduction

Forecasting plays a central role in planning and decision-making. Thus, it has always received substantial attention from researchers and practitioners. For a comprehensive review of forecasting and methodological advances, we refer the reader to an encyclopedic article by [Petropoulos et al. \(2022\)](#). To produce good quality predictions, forecasters rely on high-quality data and sophisticated mathematical models. Often, the data are collected and held by different owners at different locations, i.e., distributed in terms of geography and ownership. The pooling of this distributed data can generate additional value. For example, logistics companies can exchange their data on consumer behavior to improve their forecast of future inventory demand. Such a forecast improvement by combining or accessing more data from distributed sources is demonstrated in several studies, see [Andrade & Bessa \(2017\)](#) and [Messner & Pinson \(2019\)](#), for the example of energy applications. The general results such that forecasts can be improved through combination is already well-known within the forecasting community. However, in practice, the data owned by firms or individuals are perceived to have a cost when exposed. For businesses, this cost can be in terms of competitive disadvantage, and for individuals, in terms of privacy loss. Therefore, to utilize the distributed data, we aim at designing platforms for the pooling of predictive information. Such platforms allow for a monetary transfer from the buyer to the sellers, who are then compensated for the costs incurred for data collection, processing, modelling, etc., without explicit exposure of their private data. Because of the market context, in this work, we do not consider the infrastructural cost associated with the data.

We position our work in the area of market-based analytics, which can be broadly categorised into *data markets* and *information markets* depending on whether the traded product is the raw data or extracted information. Both these platforms have received increasing attention in the last few decades. In data markets, the key task is data valuation based on the contribution of each data seller to a learning task posted by a data buyer (the *client*), typically at a central platform ([Agarwal et al., 2019](#); [Ghorbani & Zou, 2019](#)). The market platform determines the monetary

*This work was partially supported by NWO under research project P2P-TALES (grant n. 647.003.003), the ERC under research project COSMOS (802348) and by COST under European Network for Game Theory (action CA16228). Pierre Pinson and Jalal Kazempour were additionally supported through the Smart4RES project (European Union's Horizon 2020, No. 864337). The sole responsibility of this publication lies with the authors. The European Union is not responsible for any use that may be made of the information contained therein.

*Corresponding author

Email address: a.a.raja@tudelft.nl (Aitazaz Ali Raja)

compensation that corresponds to the data value. Another significant factor in designing data markets is the cost of seller’s privacy loss (Ghosh & Roth, 2011), which plays an important role in determining the value of data, see Spiekermann et al. (2015) and Acemoglu et al. (2019). For details on data markets, we refer the reader to a comprehensive review by Bergemann & Bonatti (2019). Data markets empower data owners (sellers) to have control over the exposure of their private resources and allow buyers to obtain high-quality training data for their learning algorithms and prediction tasks. Despite their huge potential, data markets are not free from limitations and challenges. First, determining the contribution of a particular data set for a buyer is, in principle, a combinatorial problem because of the possible overlap of information among the data sets (Agarwal et al., 2019). Thus, the computational requirements for data valuation grow exponentially with the increase in the number of sellers and consequently for the evaluation of remuneration. Second, each seller can have different sensitivity to their data privacy, which makes it challenging to design a privacy-preserving mechanism. Both these issues can be addressed, to some extent, by so-called information markets.

Information markets (Linde & Stock, 2011), encompass the trade of a much broader category of information goods like news, translations, legal information, etc. However, here, we focus on the frameworks of forecasting platforms that can be categorised under the information markets. In this direction, *prediction markets* gained popularity beyond the academic circles (Wolfers & Zitzewitz, 2006; Berg et al., 2008). Prediction markets generate aggregate forecasts of uncertain future events, from dispersed information, by utilizing the notion of “wisdom of crowds”. For example, in a prediction market designed for forecasting the election result, the share price of political candidates indicates the aggregate opinion on the probability of a candidate’s win. Different from the structure of prediction markets, we design an information market for the improvement of buyer’s forecast. This improvement offered by the forecasters is remunerated via a mechanism with formal mathematical guarantees on *desirable* economic properties like budget balanced, truthfulness, etc. (Kilgour & Gerchak, 2004). Thus, in terms of design, our work is closer to the markets proposed for forecast elicitation with formal guarantees. In these works, typically, the sellers report their beliefs about a future event. Then, after the event occurs, the sellers are ranked according to the quality of their forecasts, evaluated by a *scoring rule* (Kilgour & Gerchak, 2004), (Gneiting & Raftery, 2007). An approach different from contribution-based reward, i.e., “winner takes it all”, is proposed in

Witkowski et al. (2018). Interesting to note that rewarding the best encompasses many real-world forecasting settings. For example, Netflix offered 1M USD to the team with the best prediction on how users would rate movies Witkowski et al. (2018). Even though popular in forecasting competitions, the “winner takes it all” approach ignores the fact that the forecasts other than the best one can still provide additional information. Therefore, in line with the idea of pooling the distributed information, we pursue the mechanisms that aggregate information provided by all the sellers and reward according to the quality.

In our work, we particularly take inspiration from the self-financed *wagering* market setup of Lambert et al. (2008) that features a weighted-score mechanism. In their setup, each player posts a prediction report for an event and wagers a positive amount of money into a common pool. After the occurrence of the event, the wager pool is redistributed among the players according to their relative individual performance. The payoff function is a weighted mixture of strictly proper scoring functions that satisfies several desirable economic properties. Such self-financed mechanisms create a competition in terms of forecast skill, though it does not include criteria related to the use of the forecasts, thus ignoring their value for a particular application or for an observer. In other words, there is no external agent who is aggregating, utilizing, and rewarding the resulting forecast based on the utility it generates. Differently from the setting in Lambert et al. (2008), in this paper, we design a mechanism that considers both the forecast skill of the players and the utility of the forecasts for a decision-maker.

We consider a situation where a client (see Kilgour & Gerchak (2004)) posts a forecasting task on the market platform, along with the monetary reward they are willing to pay for an improvement in their own belief. In response, the sellers report their forecasts along with their wagers. A central operator then aggregates these forecasts, considering the wagers as corresponding weights, and passes to the client for planning or decision making. We note that, unlike prediction markets, where their mechanism inherently elicits an aggregated information in terms of stock prices, the aggregation of forecasts here has to be performed methodically (Winkler et al., 2019). Thus, our first goal is to select a suitable aggregation method that reflects players’ wagers into the aggregated forecast. Next, a central operator evaluates the quality and contribution of each reported forecast and their corresponding payoffs. Our framework requires a payoff function with a utility component that rewards a contribution to the forecast improvement and a competitive component that

evaluates the relative performance of sellers to reward or penalize accordingly. Thus, our second goal is to design a collective payoff function, with utility and competitive components that enjoys desirable economic properties.

Contribution: We propose a marketplace for aggregate forecast elicitation using a wagering mechanism focused on improving the client’s utility in terms of an improvement in their forecast. The proposed market model (Section 3.1.3) is general and history-free. It is general in the sense that tasks from any application area can be posted in the form of binary, discrete, or continuous random variables. History-free implies that we do not utilize past data on sellers’ performance or market outcome, i.e., each instance of the market is set up independently. Then, we provide requirements for the aggregation of forecast reports by utilizing corresponding wagers and compare the quantile averaging to the linear pooling method as examples. (Section 3.2.1). Finally, we design a payoff function that rewards the skill of forecasters relative to each other as well as their contribution to the improvement of the utility of the client. We show that the proposed payoff function satisfies the desirable economic properties (Section 3.2.3).

2. Preliminaries

2.1. Forecasting task

Forecasting is a key requisite for decision making and planning employed in diverse situations for example, to predict a candidate’s probability of winning the election, to project an economic condition of a country, businesses forecast their sales growth for production planning, renewable energy producers make an energy generation forecast for bidding in the market, etc. The diversity in the purpose of forecasting also translates into the types of forecasting tasks faced by a decision-maker. Broadly speaking, we can categorize forecasts into point forecasts, probabilistic forecasts, and scenarios (Gneiting, 2011; Morales et al., 2014).

Point forecasts do not communicate the uncertainty associated with the possible outcomes of an event, hence an incomplete picture is delivered to a decision-maker. This shortcoming of point forecasts is resolved by probabilistic forecasts that provide decision-makers with the comprehensive information about potential future outcomes. Thus, in this paper we focus on probabilistic forecasts. A probabilistic forecast consists in a prediction of the probability distribution function (PDF) or of some summary measures of a random variable Y . These summary measures can

be quantile forecasts or prediction intervals (Gneiting & Katzfuss, 2014). The market framework proposed in this paper covers all types of probabilistic forecasts, given that the forecast evaluation method satisfies the property of being *strictly proper*. However, in the sequel, we focus on forecasting tasks in terms of PDFs for better exposition. We consider the single category, multi-category, and continuous forecasting tasks. Mathematically, these types of forecasts relate to forecasting of binary, discrete, and continuous random variables, respectively. Therefore, these cases suffice to cover most forecasting tasks we find in practice. Let us describe these forecasting tasks for uncertain events and provide relevant examples.

A single category task covers the binary events where the probability of an event happening is forecasted. For example a hedge fund predicting a return from a prospective investment has a single category forecasting task, i.e., will the quarterly growth of prospective investment be greater than $x\%$? In the probabilistic forecasting framework the task will translate into; “the probability of the quarterly growth being greater than $x\%$ ”. For a multi-category forecast, take an example of a farming company that wants to predict seasonal rainfall in categories of light, moderate and heavy. Here, the forecast is in the form of a discrete probability distribution, e.g., the rainfall in the upcoming season being {light, moderate, heavy} has probability distribution {0.2, 0.5, 0.3}. An even more comprehensive probabilistic information can be obtained by forecasting an event in terms of a continuous probability distribution. For example, a wind energy producer bidding in an electricity market can obtain the whole uncertainty associated with the day-ahead energy generation event by obtaining a forecast in terms of a probability density function (Pinson, 2012; Zhou et al., 2013).

In all three forms of forecasting presented above, the decision-makers, i.e., the hedge fund, farming company, and energy producer, can also have the in-house capability of forecasting. However, they expect that additional data and expertise can help them improve the quality of their forecasts for better planning and decision making, which in turn can lead to a higher utility. One way to achieve such a quality improvement is by designing a forecasting market platform where the data and expertise of the expert forecasters can be pooled in return for a competitive reward, depending on the contribution of each expert. When a decision-maker utilizes such a platform for forecast improvement, they expect experts to report their beliefs truthfully instead of gaming the market for higher rewards. Furthermore, the decision-maker requires the improvement offered

by the experts to be measurable by formalized criteria. Both, the guaranteed truthful reporting and numeric evaluation of the quality of probabilistic forecast can be achieved by so-called *scoring rules*.

2.2. Quality, skill and scoring

At a forecast pooling platform, a scoring rule is required for quantifying the improvement in the forecast to be used by the decision-maker. Furthermore, it allows us to rank the forecasters to assign rewards according to their contributions. We note that this assessment is performed in an ex-post sense, i.e., after the event has occurred.

Definition 1 (Scoring rule) *Let r be a reported probabilistic forecast and ω represent the event observed eventually. Then, a scoring rule $s : (r, \omega) \rightarrow \mathbb{R}$ provides a summary measure that assigns a real value for the evaluation of a probabilistic forecast r in view of the realization ω .*

In context of a marketplace for forecast elicitation, the role of the scoring rule $s(r, \omega)$ is to encourage the players to do their best in generating valuable predictive information, as well as in incentivizing their honest reporting. These tasks can be achieved by selecting scoring rules that satisfy certain properties. Next, we discuss the properties of scoring rules that we need in this work.

2.2.1. Properties of scoring rules

First, we can incentivize that the forecasters report their beliefs truthfully, by rewarding them according to a strictly proper scoring rule (Gneiting & Raftery, 2007).

Definition 2 (Strictly proper scoring rule) *Let a player report a probabilistic forecast r of an uncertain event Y . Let an outcome ω of an event be distributed according to the probability distribution p . Then, a real-valued function $s(\cdot, \omega)$ is called strictly proper when*

$$\mathbb{E}_p[s(r, \omega)] < \mathbb{E}_p[s(p, \omega)], \text{ for all } r \neq p.$$

Here, let ϱ be the support of p and f_{PDF} be the probability density function. Then, $\mathbb{E}_p[s(p, \omega)] = \int_{\varrho} s(p, \omega) f_{PDF}(p) dp$.

Later, we utilize a strictly proper scoring rule for our payoff criteria to measure the quality of the probabilistic forecasts and reward the players accordingly. There are many such rules reported in

the literature, e.g., Brier score, logarithmic score, quadratic score, etc. (Winkler et al., 1996). In principle, a scoring rule is chosen based on the properties suitable for the application. Here, for a strictly proper score rule we consider two more properties of non-local and sensitivity to distance (Gneiting & Raftery, 2007). These properties consider a complete PDF, while ranking, and allocate a higher reward to a forecaster that concentrates the probability more around the realized event. This corresponds to rewarding a higher forecasting skill on forecaster’s behalf. Next, we describe two other properties of scoring rules which we relate later on to show the effect of the choice of scoring rules on the payoff mechanism. This choice is important for implementing our proposed market design in practical scenarios.

Definition 3 (Non-local scoring (Winkler et al., 1996)) *Let the forecasters report a PDF of an event Y and we observe the corresponding outcome ω . Then, a scoring rule is called local if the score depends only on the probability (for a categorical event), or likelihood (for a continuous variable), assigned to ω . Conversely, the rule is not local if it depends on the entire reported PDF.*

Definition 4 (Sensitivity to distance (Jose et al., 2009)) *Let r be a predictive PDF and R the corresponding cumulative distribution function (CDF). Then, a CDF R' is more distant from the value x than R if $R' \neq R$, $R'(y) \geq R(y)$ for $y \leq x$, and $R'(y) \leq R(y)$ for $y \geq x$. Consequently, a scoring rule s is said to be sensitive to distance if $s(r, \omega) > s(r', \omega)$, whenever R' is more distant from R .*

In other words, a scoring rule that allocates a higher score to the player whose report has assigned higher probability to the values closer to the observed value as compared with probabilities assigned to the values farther from the true value is said to be sensitive to the distance (Winkler et al., 1996). Later in Section 4.4, we numerically illustrate the properties of locality and sensitivity to distance for building a better intuition and providing a comparison between scoring rules.

3. Proposed forecast elicitation market design

We consider a setting of a market with a single buyer and multiple sellers for eliciting a probabilistic forecast in the form of a probability distribution of an uncertain future event. In our setting, we refer to a buyer as a client and sellers as players or forecasters. A client posts a forecasting task on the market platform and announces a rate of monetary compensation for improvement in their

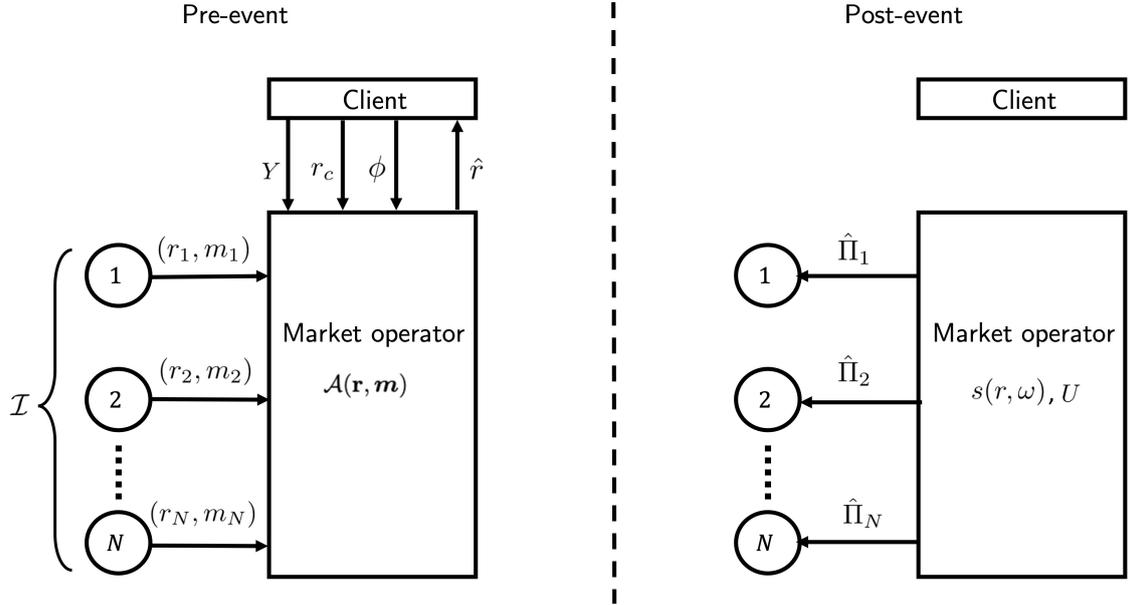


Figure 1: Market structure showing information flow and pre and post event evaluations. The delivery of \hat{r} occurs after all the inputs are received.

own belief. Players with resources and expertise in forecasting the posted task respond by reporting their forecasts along with the wagers. The market then aggregates the received information and delivers it to the client. This aggregated forecast, in turn, is expected to generate a utility for a client in terms of operational improvement. The resulting utility, considering the announced reward rate, is then distributed among the players such that it corresponds to their contribution. We note that the proposed mechanism can generally be used for the forecast eliciting of any event that can generate utility, such as the movement of a stock. Next, we formally describe our market model, and later we show the properties of the corresponding payoff distribution function.

3.1. Market model and participants

3.1.1. Client

Let there be a client i_c who is interested in improving their forecast (e.g., a generation forecast for their renewable energy asset). We parameterize a client through the following quantities:

- *Forecasting task* Y , an uncertain event that the client wants to predict better;
- *Forecast report* r_c , client's own forecast which is used as a reference for improvement;

- *Reward rate* $\phi > 0$, a monetary value that the client offers for per unit improvement in prediction.

A client can post a task Y in the form of a single category forecast (e.g., probability of energy generation being $[0.4, 0.6]$ per unit), a multi-category forecast (e.g., discrete probability distribution of energy generation in the intervals $\{[0.4, 0.6], [0.6, 0.8]\}$ per unit) and a continuous forecast (e.g., probability density function of energy generation). We note that the market design can also accommodate reports in the form of cumulative distribution functions. In the sequel, we represent the forecast reports of all three forms by r to keep the focus primarily on proposed mechanism, which holds for all forms of predictive distributions.

3.1.2. Players

Let $\mathcal{I} = \{1, \dots, N\}$ be the set of players that are forecasting experts in the area of a prediction task. We parameterize a player through the following quantities:

- *Forecast report* r_i , a prediction of forecasting task Y generated using player i 's data resources and expertise; players try to improve r_c in return for a monetary reward;
- *Wager* $m_i > 0$ which accompanies the report r_i and expresses player i 's confidence on their forecast.

A wager is associated with the player's confidence because it decides the level of impact their prediction has on the resulting forecast. Furthermore, in proposed payoff function, wagers also influence the reward (penalty) of the players.

3.1.3. Market operator

A central market operator manages the platform where a client and the players arrive with respective parameters. This operator is also responsible for maintaining transparency in the market process and is assumed to be honest. The functions of a market operator are:

- evaluation of an aggregated forecast $\hat{r}(\mathbf{m}, \mathbf{r})$, where \mathbf{r} represents a set of predictive distributions $\{r_i\}_{i=1}^N$ posted by the players and \mathbf{m} is the vector of corresponding wagers;
- evaluation of the score $s(r_i, \omega)$ of each player $i \in \mathcal{I}$, after observing the outcome ω ;

- evaluation of the utility U that corresponds to the improvement in client’s forecast; thus, in case of improvement the utility $U \propto \phi(s(\hat{r}, \omega) - s(r_c, \omega))$ and is zero otherwise.
- evaluation of the payoff $\hat{\Pi}_i$ of each player $i \in \mathcal{I}$.

Here, after the occurrence of the event, the market operator observes the true outcome ω and evaluates the score $s(r_i, \omega)$ of each player $i \in \mathcal{I}$, which shows how “good” was the forecast reported by player i . Then, the operator evaluates the utility $U(s(\hat{r}, \omega), s(r_c, \omega), \phi)$ allocated by the client and distributes it among the players that have contributed to the improvement. For transparency, the market operator publicly posts the reward rate, forecast aggregation method, scoring rule and utility evaluation method, as agreed with the client. The individual predictions posted by the players can be kept private and only an aggregated forecast is delivered to the client. In Figure 1, we show the schematic structure of the proposed market with all participants and stages. Note that the allocated utility U depends on the improvement a client has made and, for the purposes of this work, we treat it as an exogenously specified value. Further details of the forecast aggregation methods, the payoff function and their properties are discussed in the sequel.

Remark 1 *An important benefit of the proposed market architecture is that the client cannot access the underlying features; instead they only receive an aggregated forecast. This mitigates a key challenge faced by data markets where sellers are hesitant to release their proprietary data streams as they are freely replicable.*

The mechanism design of this market model requires three main components: (i) an aggregation operator (to combine forecasts), (ii) a scoring rule, and (iii) a payoff allocation mechanism. Our goal is to design a history-free mechanism, i.e., a mechanism that does not require the past data or reputation of the players to compute a solution. This allows us to keep our market general, where clients can post diverse tasks in various forms without an assumption of a repetitive market with a pre-specified task. We note that, in the sequel, we use and drop the arguments from the notations depending on the necessity. Next, we present the components of our market mechanism and discuss their properties.

3.2. Mechanism design

3.2.1. Aggregation operator

After the players have submitted their reports and wagers, in response to the client’s forecasting task, the market operator creates a collective forecast \hat{r} , using an aggregation operator. Then, the client utilizes the resulting aggregated forecast for the decision making which in turn generates some utility. An improvement in the client’s forecast r_c is rewarded at a pre-announced rate ϕ by the client. Therefore, the selection of the forecast aggregation operator constitutes an important part of the mechanism design.

Combining of probabilistic forecasts can be achieved via weighted averaging of predictive distributions. In this method, a weight assigned to a prediction reflects its relative accuracy determined by the historical data (Knüppel & Krüger, 2022). In other words, the predictions of players are weighted by their historical *performance* and have a corresponding impact on the evaluation of an aggregated forecast. Although logical, such methods are not useful for history-free mechanisms. Thus, in our proposed mechanism, the performance of a player is associated with their confidence in the reported prediction. Here, the players quantify this confidence via a wagering amount. This allows assigning an appropriate weightage to the individual forecasts while combining, which can improve the quality of an aggregated forecast. It also allows our mechanism to penalize (reward) forecasters for low (high) quality predictions, proportional to their influence on the aggregated forecast via wagers. We present this penalizing property of the payoff function named stimulant in the sequel.

Definition 5 (Aggregation operator) *An aggregation operator $\mathcal{A} : (\mathbf{r}, \mathbf{m}) \rightarrow \hat{r}$ takes a set of predictive reports $\{r_i\}_{i=1}^N$ and a vector of corresponding wagers $\mathbf{m} \in \mathbb{R}^N$ as inputs, to evaluate a combined prediction \hat{r} .*

Two candidate methods that fulfil the criteria of aggregation operator are the so-called *linear opinion pool* (LOP) and the *quantile averaging* (QA). In terms of distributional forecasts, linear averaging of the probability forecasts can be viewed as vertical combining and averaging the quantiles can be seen as horizontal combining (Lichtendahl Jr. et al., 2013). Therefore, these two methods can be regarded as two extreme cases in averaging. The first method LOP is the most widely used method in literature (Knüppel & Krüger, 2022), as well as in practice and has several extensions such as weighted linear opinion pool and optimally weighted linear opinion pool.

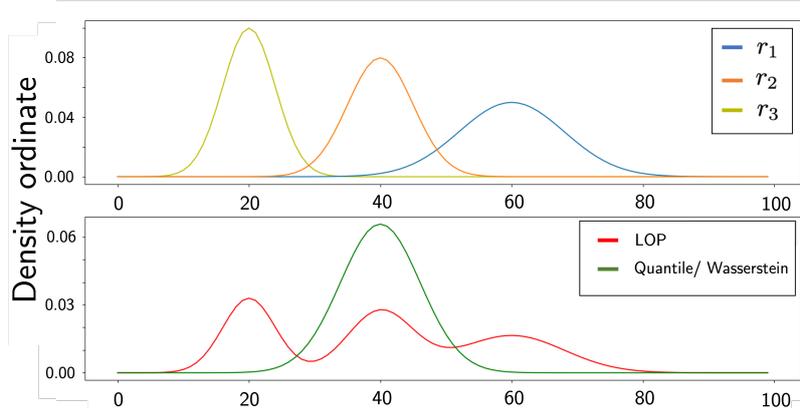


Figure 2: Comparison of LOP and quantile averaging/Wasserstein barycenter as aggregation operator

Definition 6 (Linear opinion pool) Let $\mathcal{I} = \{1, \dots, N\}$ be a set of players. Let r_i be the forecast report of player $i \in \mathcal{I}$ and m_i be the corresponding wager. Then LOP is merely an average of all the reports weighted by wagers as $\sum_i \hat{m}_i r_i$ where $\hat{m}_i = \frac{m_i}{\sum_{j \in \mathcal{I}} m_j}$.

For the optimally weighted extension, the weights m_i for all $i \in \mathcal{I}$, are evaluated by setting up an optimization problem considering the past data of the same market. However, even with optimized weights, the LOP suffers the problem of over-dispersed (under-confident) forecasting, meaning that the aggregate forecast evaluated via LOP has higher dispersion than the individual reports (Ranjan & Gneiting, 2010). The authors in Ranjan & Gneiting (2010) propose a re-calibration method to improve the combined forecast resulting from the LOP, where the re-calibration parameters are evaluated by utilizing past data. Thus, this re-calibration method is not suitable for our history-free market mechanism. Next, we explore the quantile averaging which, interestingly, also corresponds to the Wasserstein barycenter (Agueh & Carlier, 2011) of the reported forecasts.

Definition 7 (Quantile averaging) Let $\mathcal{I} = \{1, \dots, N\}$ be a set of players. For each player $i \in \mathcal{I}$, let r_i be the forecast report in terms of probability distribution function and R_i be the corresponding cumulative distribution function. Then, the average quantile forecast is given by $\hat{r}_{QA} = \sum_i \hat{m}_i R_i^{-1}$.

In Figure 2, we present an illustration for the comparison of the aggregate forecasts evaluated via LOP and QA with equal weights (wagers). It provides an intuition for how the QA keeps the shape of individual forecasts reported as widely used parametric families of distributions, e.g., normal distribution. Consequently, it also maintains the properties of those parametric families

that can comparatively provide more meaningful aggregation for decision-makers. [Lichtendahl Jr. et al. \(2013\)](#) show some useful properties of the aggregated forecast evaluated via QA. For instance, an aggregated forecast attained by QA is sharper than that by LOP and each of its even central moments is less than or equal to those of the LOP ([Lichtendahl Jr. et al., 2013](#), Prop. 8). In a memory less market, like the proposed one, a prediction which is sharper around the observation can provide better information to the decision makers and thus is regarded as of higher quality.

We note that the QA can also be interpreted as the report that minimizes the Wasserstein distance $W(\cdot, \cdot)$ from all the forecast reports, i.e., $\hat{r} = \min_r \sum_{i=1}^N W(r, r_i)$, which corresponds to the Wasserstein barycenter. We refer the reader to [Agueh & Carlier \(2011\)](#) for further details on the Wasserstein distance and barycenter.

Remark 2 *The preference of one forecasts aggregation method over the other is primarily an empirical design choice that is largely application dependent.*

3.2.2. Scoring rules

In this subsection, we specify a scoring function $s(r, \omega)$ to evaluate the quality of the forecast in an ex-post sense. We present a continuous ranked probability score (CRPS), as a strictly proper score function for elicitation of a forecast in terms of a probability density function. CRPS is non-local and sensitive to distance (see Section 2.2). For single category and multi-category prediction tasks, we present scores with same properties as of CRPS in [Appendix A](#). Note that, to stay consistent with the literature, we define scoring rules as negatively oriented, i.e., the lower, the better. However, for our design of the payoff function, presented later, we need a positively oriented scoring. Thus, in the sequel, we re-orient scoring rules for illustrative examples.

Definition 8 (Continuous ranked probability score) *For an event of interest x , let the probability density function reported by a player be r , and let ω be the event that actually occurred. Let R denote the cumulative distribution. Then, the continuous ranked probability score is defined as*

$$CRPS(R, \omega) = \int_{-\infty}^{\infty} [R_r(x) - R_\omega(x)]^2 dx \quad (1)$$

where

$$R_\omega(x) = \begin{cases} 0 & \text{if } x < \omega \\ 1 & \text{if } x \geq \omega \end{cases}$$

In words, the CRPS presents a distance between the probabilistic forecast r and the truth ω .

Note that we can conveniently re-orient the CRPS depending on the application. For example, renewable energy production can be normalized to obtain a continuous random variable $P_g \in [0, 1]$. Then, we can re-orient the scoring function by defining $s(r, \omega) = 1 - \text{CRPS}$ and consequently $s(r, \omega) \in [0, 1]$. With all the components defined, we are now ready to propose a wagering-based payoff mechanism and its desired economic properties.

3.2.3. Payoff allocation mechanism

A payoff function is central to the design of a market mechanism as it distributes the pool of wagers $\sum_j m_j$ and the generated utility U among the market players according to their *performance*. Therefore, it is critical for the design of a payoff function that it encourages market participation, on one hand by clearly reflecting the player's relative contribution, and on the other hand by enabling the delivery of valuable information to the client. The payoff functions are characterized by several desirable properties that can be proven mathematically, e.g., budget balanced, individual rationality, etc.

For the design of a payoff function, we take inspiration from [Lambert et al. \(2008\)](#), where the authors present a self-financed wagering mechanism for competitive forecast elicitation. The payoff function in [Lambert et al. \(2008\)](#) rewards the skill of the player relative to the other players by re-distributing the wagers and is shown to satisfy several interesting properties. Such self-financed markets work in the absence of a particular client with a task hence, the payoff is only based on the skill component of the players and does not involve any utility component. In other words, a player is rewarded for being better than other players regardless of the value or utility of their forecast. However, our market model in [Section 3.1.3](#) involves a client with a specified task and, therefore, our model involves an external payment associated with the utility of the client. Consequently, we need a payoff function that distributes the utility generated by the forecast, i.e., a monetary gain corresponding to an improvement in client's operational decisions, apart from rewarding the forecasting skill of the players. Let us first propose a payoff function, and then we present its

desirable economic properties.

We divide the payoff function in two parts, one representing the allocation from the wager pool and another from the client's allocated utility. The former evaluates the relative forecasting skill of a player, and the latter compensates for their contribution to an improvement of the client's utility U . Let the wager payoff of a player i be

$$\Pi_i(\mathbf{r}, \mathbf{m}, \omega) := m_i \left(1 + s(r_i, \omega) - \frac{\sum_j s(r_j, \omega) m_j}{\sum_j m_j} \right). \quad (2)$$

This term evaluates the relative performance of the players, considering the relative quality of the forecasts and the amounts wagered. It shows that the reward of player i , i.e., $\Pi_i(\mathbf{r}, \mathbf{m}, \omega) - m_i$ equals the difference between its performance (confidence and quality) and the average performance of the players. Now, let us define an indicator $\mathbb{1}_{\{a>b\}}$ that takes value 1 if $a > b$ and 0 otherwise. Then, an overall payoff is given as

$$\hat{\Pi}_i = \underbrace{\Pi_i}_{\text{skill component}} + \underbrace{\mathbb{1}_{\{U>0\}} \left(\frac{\tilde{s}(r_i, \omega) m_i}{\sum_j \tilde{s}(r_j, \omega) m_j} U \right)}_{\text{utility component}}, \quad (3)$$

where $\tilde{s}(r_i, \omega) = \mathbb{1}_{\{s(r_i, \omega) > \bar{s}\}} s(r_i, \omega)$ and $\bar{s} := s(r_c, \omega)$. Here, the utility component depends on an improvement offered by the player beyond the client's own resources r_c . Thus, to be eligible for a share of an allocated utility U , first, there should be an improvement in the client's resulting forecast, i.e., $U > 0$, and second, the score of player i , $s(r_i, \omega)$ should be greater than the score of the client. Here, the utility payoff of a player is always non-negative but a wager payoff can also create a loss, i.e., $\Pi_i - m_i < 0$ is possible. The possibility of a loss encourages players to compete in improving the forecast by employing better models and acquiring more meaningful data. We note that the client can achieve negative utility as well, i.e., the forecast becomes worst than their own prediction. However, again with a penalty imposed by the wagering part of the payoff function, it is expected from risk-averse players to report high-quality forecasts. Next, we provide a brief explanation of some desirable properties of a payoff function.

Desirable properties: The properties are adapted from [Lambert et al. \(2008\)](#) and here we include their explanations in context of the payoff function in (3).

- i) *Budget-balance:* A mechanism is budget-balanced if the market generates no profit and creates

no loss, i.e., $\sum_{i \in \mathcal{I}} \hat{\Pi}_i = \sum_{i \in \mathcal{I}} m_i + U$. In other words, the generated utility and the wager pool must be completely distributed, as a payoff, among the players.

ii) *Anonymity*: A mechanism satisfies anonymity if the payoff received by a player does not depend on their identity; rather it depends only on the forecast reports and the realization of an uncertain event. Formally, for any permutation σ of \mathcal{I} , the payoff $\hat{\Pi}_i((r_i), (m_i), \omega, U) = \hat{\Pi}_{\sigma(i)}\left((r_{\sigma^{-1}(i)}), (m_{\sigma^{-1}(i)}), \omega, U\right)$ for all $i \in \mathcal{I}$.

iii) *Individually rational*: Let the belief of a player $i \in \mathcal{I}$ about an event be p . Then, a mechanism is individually rational if for any wager $m_i > 0$ there exists r_i^* such that an expected profit of a player is non-negative, i.e., $\mathbb{E}_p[\hat{\Pi}_i((\mathbf{r}_{-i}, r_i^*), \mathbf{m}, \omega, U) - m_i] \geq 0$, for any vector of wagers \mathbf{m}_{-i} and reports \mathbf{r}_{-i} .

Individual rationality encourages the participation of players by ensuring a non-negative expected profit according to their beliefs.

iv) *Sybilproofness*: A truthful mechanism is sybilproof if the players cannot improve their payoff by creating fake identities and copies of their identities. Formally, let the reports \mathbf{r} and vectors of wagers \mathbf{m} and \mathbf{m}' be such that for a subset of players $S \subset \mathcal{I}$ the reports $r_i = r_j$ for $i, j \in S$, the wagers $m_i = m'_i$ for $i \notin S$ and that $\sum_{i \in S} m_i = \sum_{i \in S} m'_i$. Then, the sybilproofness implies that, for all $i \notin S$, $\hat{\Pi}_i(\mathbf{r}, \mathbf{m}, \omega, U) = \hat{\Pi}_i(\mathbf{r}, \mathbf{m}', \omega, U)$ and that $\sum_{i \in S} \hat{\Pi}_i(\mathbf{r}, \mathbf{m}, \omega, U) = \sum_{i \in S} \hat{\Pi}_i(\mathbf{r}, \mathbf{m}', \omega, U)$.

We note that the Shapley value, a solution used to evaluate data in market setting, suffers the drawback of being prone to replication, i.e., players can increase their payoff by creating fake copies of themselves (Agarwal et al., 2019). This consideration takes special importance in markets dealing with forecasts as the data are a freely-replicating good.

v) *Conditionally truthful for players*: A mechanism is conditionally truthful if the player does not have enough information or influence over the payoff function to manipulate it for their benefit. Thus, reporting their true belief becomes the best strategy for a risk averse player.

This definition of conditional truthfulness considers practical situations for the players and the market operation. Truthfulness of a mechanism encourages the players to post their true belief at the market platform thus, fulfilling the client's expectation of having an access to the honest assessments of the experts about an event.

vi) *Truthful for the client*: A mechanism is truthful for a client, in terms of reported prediction, if the client's expected payment (allocated utility U) is minimized by reporting their true belief p as their own forecast, i.e., $\mathbb{E}_p [U(s(\hat{r}, \omega), s(r_c, \omega), \phi)] > \mathbb{E}_p [U(s(\hat{r}, \omega), s(p, \omega), \phi)]$ is satisfied for all $r_c \neq p$. We note that the truthfulness of the client concerns the prediction report r_c and not the reward rate ϕ . Since, with our single-buyer design, it is not possible to elicit their true willingness to pay.

vii) *Stimulant*: Let a player i 's payoff be the sum of skill and utility components, i.e., $\pi_i(\mathbf{r}, (\mathbf{m}_{-i}, m_i), \omega, U) = \pi_i^s(\mathbf{r}, (\mathbf{m}_{-i}, m_i), \omega) + \pi_i^u(\mathbf{r}, (\mathbf{m}_{-i}, m_i), \omega, U)$. Let the wager be $m'_i > m_i$. Then, this payoff is monotonic if it holds that for the skill component, either

$$0 < \mathbb{E}_p [\pi_i^s(\mathbf{r}, (\mathbf{m}_{-i}, m_i), \omega) - m_i] < \mathbb{E}_p [\pi_i^s(\mathbf{r}, (\mathbf{m}_{-i}, m'_i), \omega) - m'_i]$$

or

$$0 > \mathbb{E}_p [\pi_i^s(\mathbf{r}, (\mathbf{m}_{-i}, m_i), \omega) - m_i] > \mathbb{E}_p [\pi_i^s(\mathbf{r}, (\mathbf{m}_{-i}, m'_i), \omega) - m'_i].$$

In words, a mechanism is monotonic if a player's expected profit, as well as loss from the skill component, increases by increasing the wager. Now, for a utility factor, let $U > 0$ and $s(r_i, \omega) > \bar{s}$. Then,

$$\pi_i^u(\mathbf{r}, (\mathbf{m}_{-i}, m_i), \omega) < \pi_i^u(\mathbf{r}, (\mathbf{m}_{-i}, m'_i), \omega).$$

These properties encourage the players to post higher wagers considering their confidence in their forecasts thus we refer to them as stimulant. Importantly, it also justifies weighting the forecasts by the corresponding wagers while creating an aggregate forecast. We note that, for real-world applications, the market operator can place lower and upper bounds on the amounts of wagers considering the viability of the market.

Now, we show that the proposed payoff criterion in (3) satisfies all the desirable properties described above.

Theorem 1 (Characteristics of payoff allocation) *Let $s(r, \omega) \in [0, 1]$ be a strictly proper score function. Then, the payoff function given in (3) is (i) budget-balanced, (ii) anonymous, (iii) individually rational, (iv) sybilproof, (v) conditionally truthful for players, (vi) truthful for the client, and (vii) stimulant.*

We provide the proof of the theorem in Appendix to improve the reading flow.

4. Illustrative examples

In this section, we illustrate several numerical examples to provide some intuition on the proposed market model and to numerically demonstrate the properties of the proposed payoff function in (3). For all the illustrations, we use a beta distribution, with parameters (α, β) , as a base predictive density. We then vary its parameters to simulate potential forecast reports of different players. We acknowledge that these reports might not represent a real-world scenario. However, these examples are sufficient to illustrate and discuss interesting properties of the payoff function.

4.1. Effect of wager amount

Let a client post a prediction task Y on a market platform along with their own forecast that has a score of 0.5, i.e., $s(r_c, \omega) = 0.5$. In response, let the players $\mathcal{I} = \{1, 2, 3\}$ post the predictive densities of a random variable $Y \in [0, 1]$, as shown in Figure 3. Though, in reality we expect the reports by expert forecasters to be concentrated around nearby values but here we consider an extreme case to emphasize our observations. First, we evaluate the players' payoff for equal wagers and then increase the wager of player 3 to demonstrate the stimulant property of the payoff function, defined in Section 3.2.3. Suppose the market operator announces the cap on the wager amount, i.e., maximum value a player can wager, $\bar{m} = 500$. The case of equal wagers in Table 1a shows a loss for player 3, taken from their wager, for posting a sharp predictive density concentrated far from the realized event, $\omega = 0.8$. The corresponding aggregate prediction \hat{r}_a , shown in Figure 3, has a score of 0.867. Here, the score of player 3 is lower than client's score and thus it doesn't receive any share from utility payoff. We note that the score of each player is given by a positively oriented scoring rule (1-CRPS) and the utility of a client is assumed to be specified exogenously. Next, for the case in Table 1b, the wager of player 3 is increased to maximum, which results in the increase of loss. This implies that showing more confidence via higher wager on a "bad" forecast will result in a higher loss which is an important consequence as a higher wager by player 3 resulted in the reduced quality of the aggregated prediction \hat{r}_b , as shown in Figure 3, with $s(\hat{r}_b, \omega) = 0.822$. This example illustrates the justification for using wagers as weights in the aggregation method. It also demonstrates how using a wager as a player's confidence results in a fair penalty or reward for them.

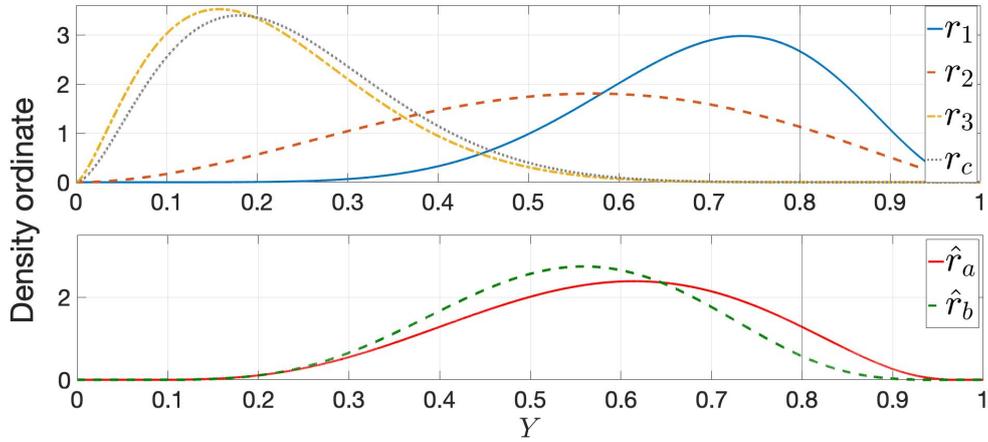


Figure 3: Plot on the top shows the reports on density forecast of random variable $Y \in [0, 1]$ by market participants and the bottom plot shows aggregate density forecasts for wagering case (a) and (b) as in Table 1a and 1b, respectively. The vertical line is at the realization, $\omega = 0.8$.

Table 1: Profit (payoff - wager) evaluation for forecast reports in Figure 3

and its sensitivity to wagers							
(a)				(b)			
Players	1	2	3	Players	1	2	3
Wager	100	100	100	Wager	100	100	500
Scores	0.9430	0.8450	0.4830	Scores	0.9430	0.8450	0.4830
Profit	546	481.39	-27.40	Profit	552.85	488.24	-41.10

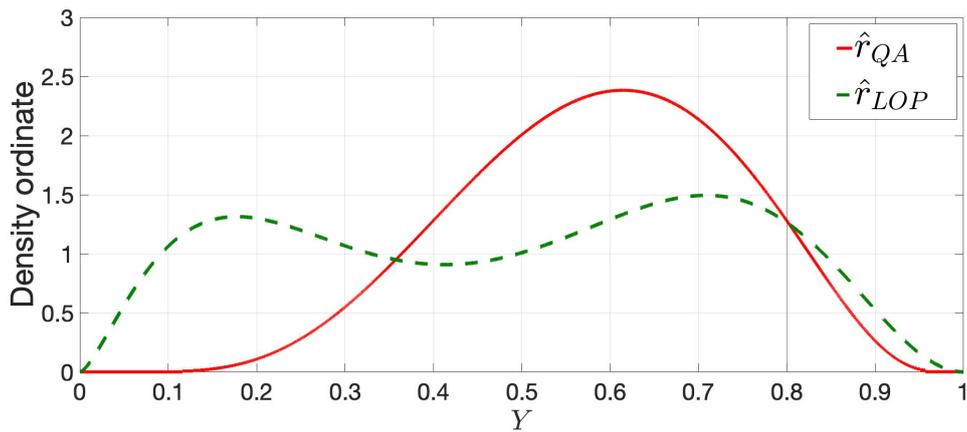


Figure 4: Plots of aggregate predictive densities obtained by quantile averaging \hat{r}_{QA} and linear pooling \hat{r}_{LOP} in an equal wagering case.

Table 2: Sybilproofness of profit (payoff - wager) in proposed mechanism

(a)			(b)			
Players	1	2	Players	1	2(a)	2(b)
Wager	100	100	Wager	100	40	60
Scores	0.9430	0.8450	Scores	0.9430	0.8450	0.8450
Profit	532.30	467.69	Profit	532.30	187.07	280.61

4.2. Comparison of QA and LOP

In Figure 4, we present the comparison of aggregate predictive distributions obtained via quantile averaging \hat{r}_{QA} and linear pooling \hat{r}_{LOP} . It is evident how \hat{r}_{LOP} can be problematic for a decision-maker. The loss of sharpness translates into lower scores for linear opinion pool as well where, $s(\hat{r}_{LOP}, \omega) = 0.817$ compared with $s(\hat{r}_{QA}, \omega) = 0.867$. Furthermore, for commonly used parametric distributions quantile averaging maintains the shape of the distribution, while linear pooling does not.

4.3. Demonstration of sybilproofness

Now, we illustrate the property of sybilproofness (see Section 3.2.3), which in truthful mechanisms prevents players from manipulating identities. Sybilproofness of payoff function is specially important for electronic platforms. Table 2a shows profit and scores of two players with reported predictive densities r_1 and r_2 , as in Figure 3. Now, let the player 2 create a fake identity and appear in the market as 2(a) and 2(b) with different wagers, as reported in Table 2b. We note that, even after identity manipulation, the collective profit of both identities of player 2 remained the same as with the true identity. Consequently, it does not affect the player 1 as well.

4.4. Sensitivity of scoring rules

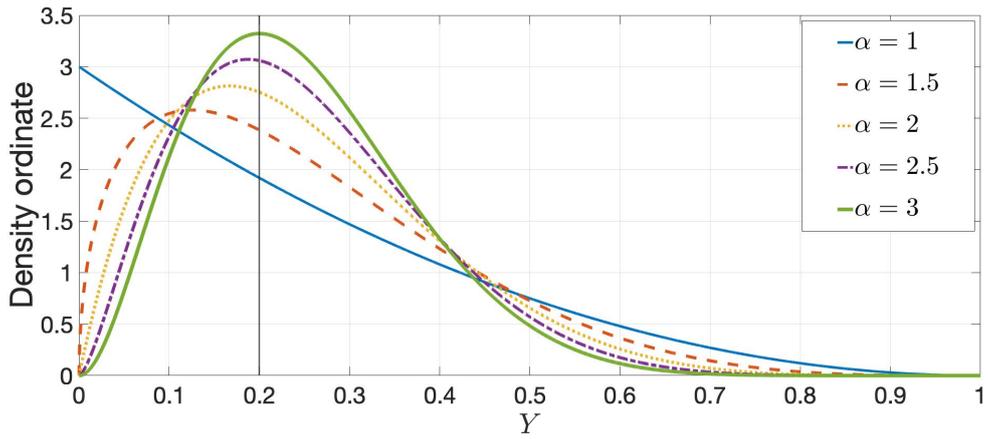
In this section, we demonstrate various properties of scoring rules to emphasize their effect on the design of a payoff function. Generally, the choice of a scoring rule depends on the application area of the prediction task. Thus, these illustrations are important to provide useful insights to the practitioners for adopting the proposed mechanism to a particular application. The choice of scoring rules can also affect the willingness of players to participate and constitute an important part of the design.

4.4.1. Local vs. non-local scoring

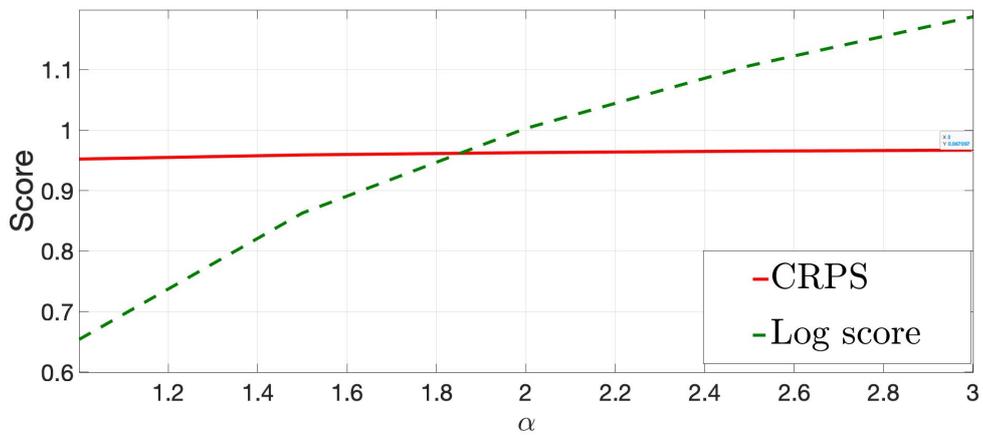
Different scoring rules differ in their sensitivity to the variation in prediction quality. For applications where sharp predictions are required because of the high stakes, the scoring rules with higher sensitivity can perform better. Let us now compare the sensitivity of CRPS and log score by varying parameters (α, β) of predictive densities. To illustrate these effects across the variation in single parameter α , we fix the mean of densities and then evaluate β as $\beta = \frac{\alpha(1-\text{mean})}{\text{mean}}$. We note that in parametric case the variation in parameters simulates the varying quality or features utilized to construct the predictive densities. In Figure 5b, we show the predictive beta distributions for different values of α and the corresponding CRPS and log scores. As the log score depends only on the realization ω , it has considerable variation for given predictive densities. Whereas, CRPS takes complete information into account thus varies slightly with the slight change in densities. The scoring rules are selected essentially by considering the nature of the prediction task at hand. We note that our results hold for all strictly proper scoring rules, including the normalized log score.

4.4.2. Sensitivity to distance

In this example, we illustrate the impact of the scoring rules' sensitivity to distance (see Definition 4). Let the three forecasters E_1, E_2 and E_3 provide a normalized multi-category probabilistic forecasts for the energy generation y of a wind producer for intervals $\{[0 - 0.2], (0.2 - 0.4], (0.4 - 0.6], (0.6 - 0.8], (0.8 - 1]\}$ per-unit represented by $\{1, 2, 3, 4, 5\}$. Let the reported probabilistic forecasts of E_1, E_2 and E_3 be $\{0.1, 0.1, 0.6, 0.1, 0.1\}$, $\{0, 0.2, 0.6, 0.2, 0\}$ and $\{0.2, 0, 0.6, 0, 0.2\}$, respectively. Suppose we observe the actual wind production in the third interval, i.e., $y = 3$. Let us now assess the quality of the forecasts using quadratic and ranked probability scoring (RPS) rules (see Winkler et al. (1996) and Appendix A for mathematical expressions). Here, E_1 receives a quadratic score of 0.8 while E_2 and E_3 receive 0.76. We first observe that all three forecasters have assigned a probability of 0.6 to the realized value of y . Next, we note that E_2 assigns the remaining probability of 0.4 to the intervals 2 and 4, that are adjacent to the realized interval, i.e., 3, while E_3 assigns it to the farthest (more distant) intervals. This probability assignment shows comparatively a better forecasting skill on behalf of E_2 . However, their scores are same, which shows that the quadratic scoring is not sensitive to distance. In comparison, RPS assigns 0.975, 0.98 and 0.96 to the predictions of E_1, E_2 and E_3 , respectively. We note that RPS acknowledges the concentration



(a)



(b)

Figure 5: (a) Predictive beta distributions with same mean = 0.25 where, for each given α , $\beta = \frac{\alpha(1-\text{mean})}{\text{mean}}$. (b) Comparison of scores assigned to predictive distributions via CRPS and log-score.

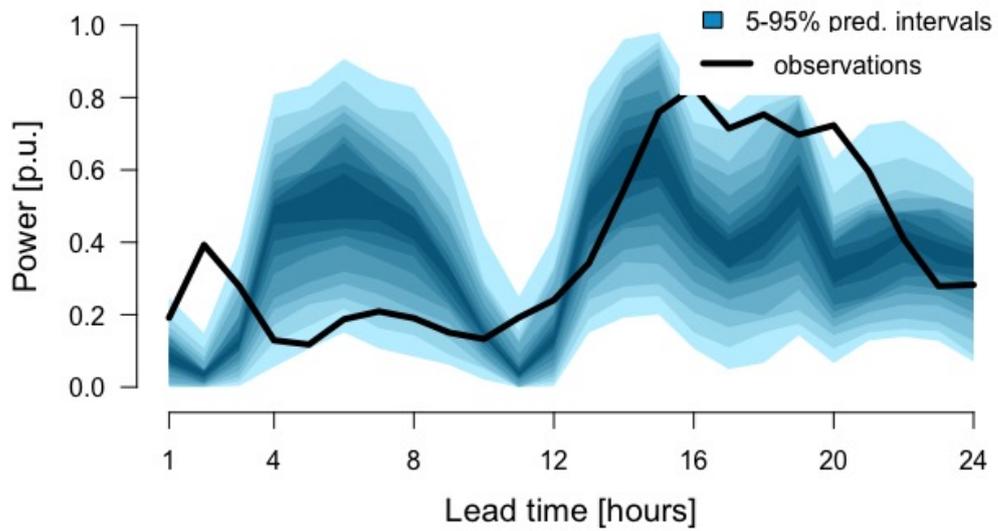
of probability around the observation and assigns highest score to E_2 . Thus, RPS is sensitive to distance which can be important for the practitioners while designing a payoff function.

5. Wind energy forecasting: A case study

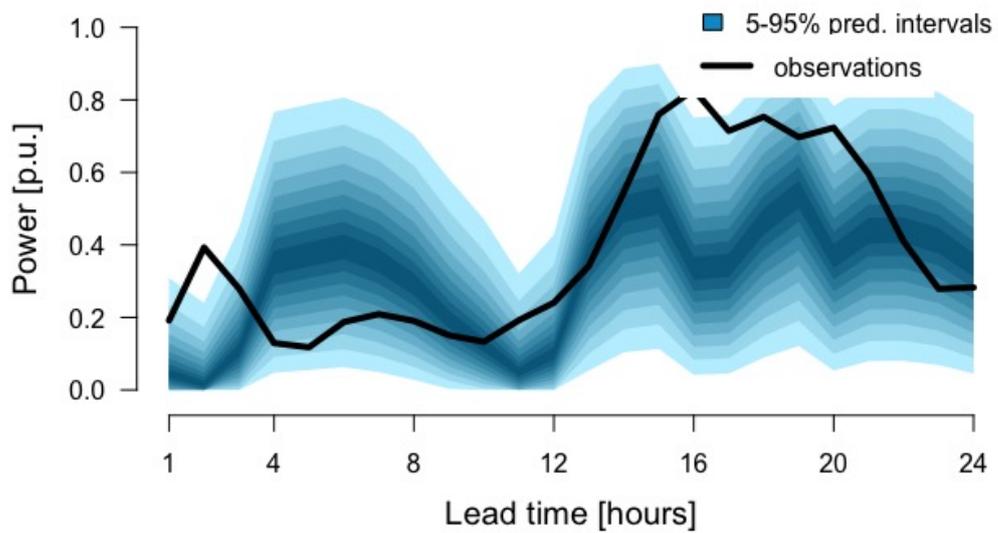
In this section, we present an energy forecasting application of the proposed market mechanism. Here, we differentiate forecasters based on their forecasting skill and resourcefulness. In former, the players utilize same data but different models (forecasting skill) to construct predictive densities and vice versa in the latter. This differentiation criteria covers an important feature of forecasting market that it creates a competition of both resourcefulness (data) and forecasting skill among the players. The aim of this case study is to demonstrate the compensation allocated by our market mechanism for eliciting forecasts evaluated by experts based on their private information and skills. Elicited forecasts are aggregated and delivered to the client.

5.1. Simulation setup

Consider a wind energy producer who wants to improve its generation forecast for more informed bidding in an electricity market, thereby avoiding a penalty for causing an imbalance. For this purpose, the energy producer arrives at the wagering based forecasting market, described in Section 3, as a client. We assume that the client submits the task of forecasting the next 24-hours of wind energy generation. In response, let the forecasters \mathcal{I} submit the probabilistic forecasts along with their wagers. The market operator evaluates the scores of submitted forecasts on hourly basis and compensates accordingly. For our case study, we use an open data set from the Global Energy Forecasting Competition 2014, GEFcom2014 (Hong et al., 2016) and an open-source toolkit ProbCast by Browell & Gilbert (2020). The wind power measurements are normalized and thus take values in $[0, 1]$. For the market setup, we assume fixed utility U , offered by the client, to analyse scores and the share of each player's payoff $\hat{\Pi}_i$ in $\sum_i m_i + U$. We note that, in reality, the compensation provided by the client depends on the operational benefits that they receive through an improvement in their forecast. Next, we first present a simpler case of wind energy forecasting with 2 players evaluate the resulting payoff allocation, as in (3), and later we move to more extensive cases.



(a)



(b)

Figure 6: (a) Wind energy generation forecast reported by player 1 via inflated beta distribution, i.e., r_1 . (b) Wind energy generation forecast reported by player 2 via non-parametric predictive density, i.e., r_2 . Observations represent realization ω .

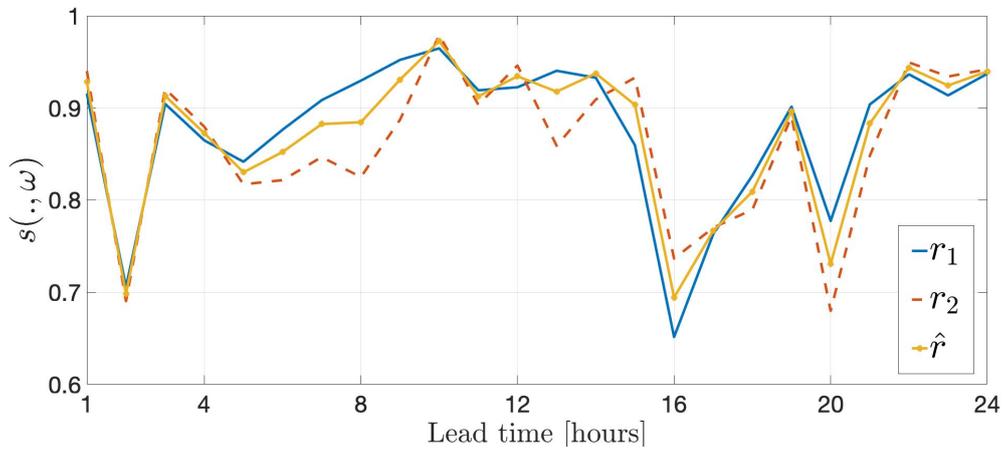


Figure 7: CRPS of forecasts reported by player 1 (r_1), player 2 (r_2) and an aggregate forecast \hat{r} .

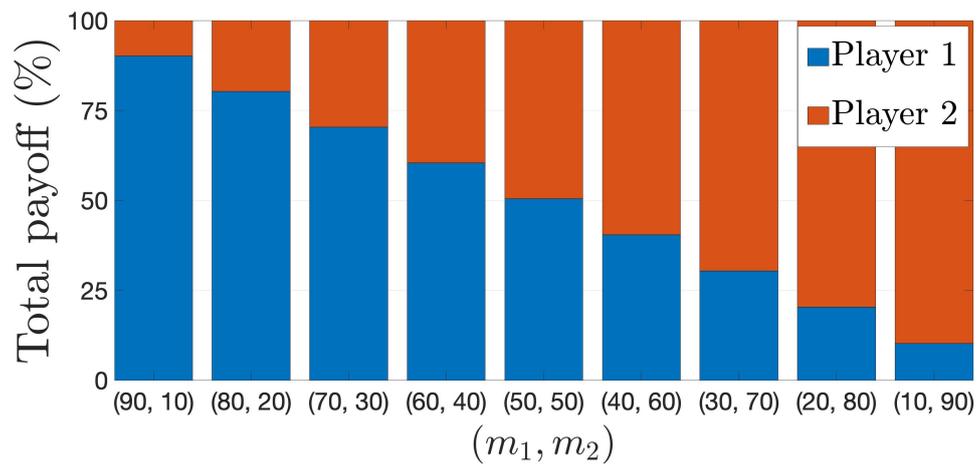


Figure 8: Players' total payoff of 24 hours as a share of money pool $\sum_i m_i + U$ for different wagers.

5.2. Forecasting market with 2 players

Let the players $\mathcal{I} = \{1, 2\}$ provide wind energy generation forecast for the next 24 hours. Here, we assume that both forecasters have the same data but they utilize different models to generate predictive densities for wind energy forecasting. Selection of a particular forecasting model can be seen as a forecasting skill of a player thus, the players have different forecasting skills. In this case, player 1 provides their wager m_1 and the forecast report r_1 as a parametric distribution, i.e., an inflated beta distribution as proposed by [Ospina & Ferrari \(2010\)](#) generated by using a generalised additive model GAMLSS. Whereas, player 2 utilizes gradient boosted regression trees to generate non-parametric predictive densities and submits the forecast report r_2 along with the wager m_2 . Let the market operator announce wager bounds such that $m_1, m_2 \in [10, 100]$. We assume that the score of the client's own forecast is constant at 0.5 for all 24 hours. Such a low score shows that the client has a low-quality forecast and consequently, for our data, the players will be eligible for utility payoff at each hour. After receiving the reports, the market operator evaluates an aggregate forecast \hat{r} and delivers it to the wind energy producer (client), who in turn uses it for operational planning. [Figures 6a](#) and [6b](#) show the reports of player 1 and player 2, i.e., r_1 and r_2 , respectively. The hourly observations represent the realization ω , i.e., the actual wind energy generation during the corresponding hour. After the forecasting period has passed, the market operator evaluates the score of each player and that of an aggregate forecast. [Figure 7](#) shows the scores (CRPS) of r_1, r_2 and \hat{r} . We note that the aggregate forecast \hat{r} evaluated via quantile averaging, as in [Definition 7](#), depends on the wagers of the players, and [Figure 7](#) is the case of equal wagers. The difference in the scores of both players is not much as their reported predictive densities follow a similar trend. Though the score rank of players varies at different hours the parametric forecaster performs slightly better in a cumulative sense, for this particular instance of market. If this variation in score rank is considerable the aggregate forecast can score better than both players. We illustrate this fact later in our case study. Next, we show players' total payoff for 24 hours as a share of money pool $\sum_i m_i + U$. The payoff, as in [\(3\)](#), also depends on wagers m_i and in the case of equal wagers it corresponds directly to the scores. To observe the effect of wager, in [Figure 8](#) we plot payoff across different wager pairs. As both players offer improvement and the scores of both players do not differ much, the stimulant property of our payoff function, explained in [Section 3.2.3](#), allocates higher payoff to high wagering player.

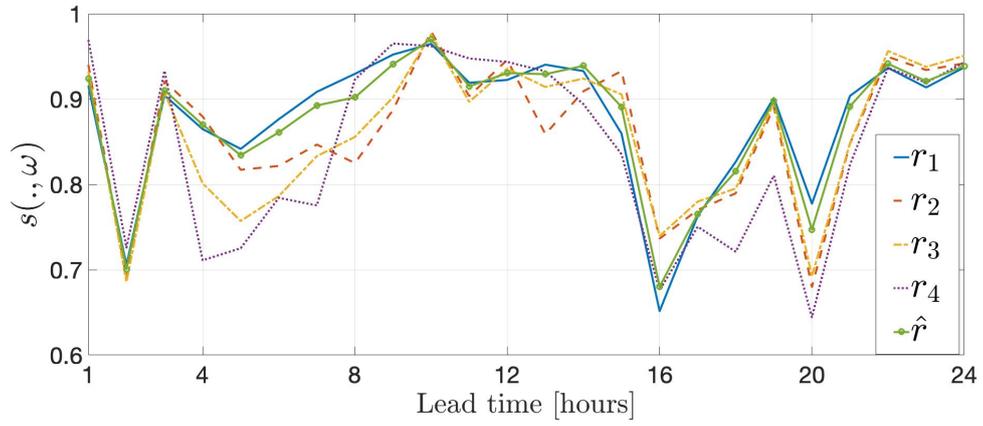


Figure 9: CRPS of forecasts reported by players (r_1, r_2, r_3 and r_4) and an aggregate forecast \hat{r} .

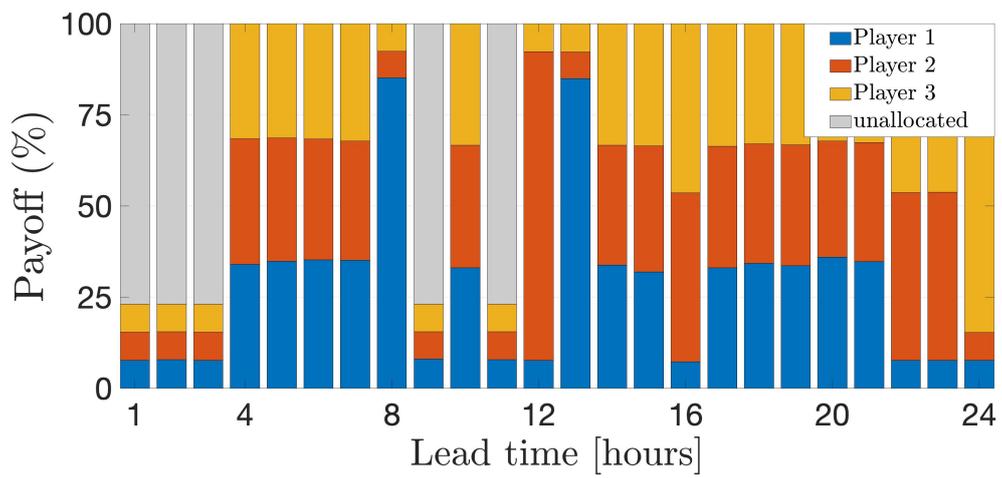


Figure 10: Players' hourly payoff as a share of money pool $\sum_i m_i + U$, assuming equal wagers.

Table 3: Total score (CRPS) of reported forecasts over 24 hour period.

Report	r_1	r_2	r_3	r_4	\hat{r}
Total score	21.0480	20.6978	20.6090	20.2514	21.0074

5.3. Forecasting market with 4 players

Now, let two more players join the market referred as player 3 and player 4. We assume that these new players have the same forecasting skill, i.e., both players utilize same forecasting method. However, the data held/ collected by the players is different. Player 3 holds the data of wind forecasts, as predictor, at the height of 10 m above the ground level whereas player 4 has the data of wind forecasts at 100 m above ground level. Wind forecast being a key predictor effects the quality of energy generation forecasts. The quality of all 4 reports is evaluated by CRPS and is presented in Figure 9 along with the score of an aggregate forecast. In Table 3, we report total scores of all forecast reports over the period of 24 hours. Interestingly, for this market instance, the score of aggregate forecast $s(\hat{r}, \omega)$ is higher than that of individual forecast reports of all the players.

To analyse the hourly payoff allocation when the client has a forecast report of a reasonable quality, we assume player 4 to be the client, i.e., $r_c = r_4$, as in (3). Consequently, according to proposed payoff function in (3), a player becomes eligible for a utility payoff only when it offers an improvement to the client, i.e., scores higher than the client. Assuming a fixed utility payoff U , we present players' payoff allocation in Figure 10. We can observe that for the first 3 hours the score of client's forecast report (r_4) in Figure 9 is higher than the players thus, the payoff distribution occurs only from the wager pool $\sum_i m_i$. As we consider a fixed utility component U , there remains an unallocated utility payoff component which is returned back to the client. In contrast, if utility component depends on the forecast improvement of the client then $U = 0$ in case of first 3 hours. Next, observe that at the 12th hour only player 2 offers slight improvement, i.e., scores higher than the client (see Figure 9) thus, they receive the whole offered utility payoff.

6. Conclusion

We have designed a marketplace for revealing an aggregate forecast by eliciting truthful individual forecasts from a group of forecasters. In the proposed model, a client with a prediction task

calls for forecasts on a market platform and announces a monetary reward for it. The forecasters respond with predictive reports and wagers showing their confidence. The platform aggregates the forecasts and delivers them to the client. Here, the utilized aggregation criteria allows us to make our mechanism a one-shot history-free method that does not account for the forecaster’s performance in the past. Next, upon the realization of the event, it allocates payoffs to the forecasters depending on the quality of their forecasts. We have proposed a payoff function with skill and utility components that depend on the relative forecast quality of a forecaster and their contribution to improving the forecast of the client, respectively. We show that the proposed payoff allocation satisfies several desirable economic properties, including budget balance, anonymity, conditional truthfulness, sybilproofness, individual rationality, and stimulant. The simplicity of scoring-based market design, with a wagering mechanism, allows it to cater diverse forecasting tasks with forecasting reports taking forms of discrete to continuous probability distributions.

From the success story of platforms like NUMERAI (NUMERAI, 2022), we see a high potential for real-world aggregative forecasting marketplaces. Differently from current implementations, the mechanism proposed in this paper is designed for the improvement of predictions and provides theoretical guarantees on the monetary compensation that can encourage and retain the participation of experts. Next, we envision a competition platform to test the performance of the proposed market model and the behavior of players in practical scenarios. Such an experimental setup would help us gain further insights for real-world implementation. Furthermore, our market setup opens several paths for applied modelling of information eliciting platforms and their analysis. An important step is to design a mechanism for online predictions based on streaming data and in turn analyse if it maintains the economic properties discussed in this paper. Another interesting research avenue is to design models that value the reputation of forecasters (historic credits) as well.

Appendix A. Scoring rules

Let us present strictly proper scoring rules for single-category and multi-category reporting that are non-local and sensitive to distance (see Section 2.2). A strictly proper scoring rule which is non-local and can be used for eliciting a single-category forecast for binary events, is the Brier score.

Definition 9 (Brier score) *Let the probability of occurrence of an event x , reported by a player, be r and let ω be the actual outcome. Then, the Brier score is given as*

$$\text{BS} = (r - \omega)^2. \quad (\text{A.1})$$

Interestingly, a generalization of the Brier score known as ranked probability score (RPS), which is also non-local and sensitive to distance, can be used for multi-category forecasting tasks where the reports are in the form of discrete probability distributions.

Definition 10 (Ranked probability score) *Let the multi-category forecasting task have J categories. Let $r(i)$ be the forecasted probability of outcome i and $\omega(j)$ represents if the category j has occurred. Then, the ranked probability score is defined as*

$$\text{RPS} = \sum_{i=1}^J (R(i) - O(i))^2 \quad (\text{A.2})$$

with $R(i) = \sum_{j=1}^i r(j)$ and $O(i) = \sum_{j=1}^i \omega(j)$.

Appendix B. Proof of Theorem 1

Let us now provide the proof of the properties mentioned in Theorem 1.

1. *Budget balance:* For any vector of reports \mathbf{r} , wagers \mathbf{m} and an outcome ω ,

$$\begin{aligned} \sum_i \hat{\Pi}_i &= \sum_i \Pi_i(\mathbf{r}, \mathbf{m}, \omega) + \sum_i \frac{\tilde{s}(r_i, \omega) m_i}{\sum_j \tilde{s}(r_j, \omega) m_j} U \\ &= \sum_i m_i + \sum_i s(r_i, \omega) m_i - \left(\sum_i m_i \right) \left(\frac{\sum_j s(r_j, \omega) m_j}{\sum_j m_j} \right) \\ &\quad + \sum_i \frac{\tilde{s}(r_i, \omega) m_i}{\sum_j \tilde{s}(r_j, \omega) m_j} U \\ &= \sum_i m_i + U. \end{aligned}$$

2. *Anonymous*: Let σ be any permutation of \mathcal{I} . For any $\mathbf{r}, \mathbf{m}, \omega$, and i ,

$$\begin{aligned}
\hat{\Pi}_{\sigma(i)} \left((r_{\sigma^{-1}(j)})_{j \in \mathcal{I}}, (m_{\sigma^{-1}(j)})_{j \in \mathcal{I}}, \omega, U \right) &= m_{\sigma^{-1}(\sigma(i))} \left(1 + s(r_{\sigma^{-1}(\sigma(i))}, \omega) \right. \\
&\quad \left. - \frac{\sum_j s(r_{\sigma^{-1}(j)}, \omega) m_{\sigma^{-1}(j)}}{\sum_j m_{\sigma^{-1}(j)}} \right. \\
&\quad \left. + \frac{s(r_{\sigma^{-1}(\sigma(i))}, \omega)}{\sum_j s(r_{\sigma^{-1}(j)}, \omega) m_{\sigma^{-1}(j)}} U \right) \\
&= m_i \left(1 + s(r_i, \omega) - \frac{\sum_j s(r_j, \omega) m_j}{\sum_j m_j} \right. \\
&\quad \left. + \frac{s(r_i, \omega)}{\sum_j s(r_j, \omega) m_j} U \right) \\
&= \hat{\Pi}_i \left((r_j)_{j \in \mathcal{I}}, (m_j)_{j \in \mathcal{I}}, \omega, U \right).
\end{aligned}$$

3. *Individually rational*: The skill factor Π_i of the payoff function in (3) is individually rational by Theorem 1 in Lambert et al. (2008) and the utility factor is always non-negative. Thus, the payoff $\hat{\Pi}_i$ is individually rational, i.e., $\mathbb{E}[\hat{\Pi}_i - m_i] \geq 0$.

4. *Sybilproofness*: Let a vector of reports \mathbf{r} and vectors of wagers \mathbf{m} and \mathbf{m}' such that for a subset of players $S \subset \mathcal{I}$ the reports $r_i = r_j$ for $i, j \in S$, the wagers $m_i = m'_i$ for $i \notin S$ and that $\sum_{i \in S} m_i = \sum_{i \in S} m'_i$. Let players $i \in S$ post a common forecast report r then, for any $i \notin S$,

$$\begin{aligned}
\hat{\Pi}_i(\mathbf{r}, \mathbf{m}, \omega, U) &= m_i \left(1 + s(r_i, \omega) - \frac{\sum_{j \notin S} s(r_j, \omega) m_j + s(r, \omega) \sum_{j \in S} m_j}{\sum_{j \notin S} m_j + \sum_{j \in S} m_j} \right. \\
&\quad \left. + \frac{\tilde{s}(r_i, \omega)}{\sum_{j \notin S} s(r_j, \omega) m_j + s(r, \omega) \sum_{j \in S} m_j} U \right) \\
&= m'_i \left(1 + s(r_i, \omega) - \frac{\sum_{j \notin S} s(r_j, \omega) m'_j + s(r, \omega) \sum_{j \in S} m'_j}{\sum_{j \notin S} m'_j + \sum_{j \in S} m'_j} \right. \\
&\quad \left. + \frac{\tilde{s}(r_i, \omega)}{\sum_{j \notin S} s(r_j, \omega) m'_j + s(r, \omega) \sum_{j \in S} m'_j} U \right) \\
&= \hat{\Pi}_i(\mathbf{r}, \mathbf{m}', \omega, U).
\end{aligned}$$

Additionally, for all $i \in S$

$$\begin{aligned}
\sum_{i \in S} \hat{\Pi}_i(\mathbf{r}, \mathbf{m}, \omega, U) &= \sum_{i \in S} m_i \left(1 + s(r, \omega) - \frac{\sum_{j \notin S} s(r_j, \omega) m_j + s(r, \omega) \sum_{j \in S} m_j}{\sum_{j \notin S} m_j + \sum_{j \in S} m_j} \right) \\
&\quad + \sum_{i \in S} \frac{\tilde{s}(r, \omega) m_i}{\sum_{j \notin S} s(r_j, \omega) m_j + s(r, \omega) \sum_{j \in S} m_j} U \\
&= \left(\sum_{i \in S} m_i \right) \left(1 + s(r, \omega) - \frac{\sum_{j \notin S} s(r_j, \omega) m_j + s(r, \omega) \sum_{j \in S} m_j}{\sum_{j \notin S} m_j + \sum_{j \in S} m_j} \right) \\
&\quad + \frac{\tilde{s}(r, \omega) \sum_{i \in S} m_i}{\sum_{j \notin S} s(r_j, \omega) m_j + s(r, \omega) \sum_{j \in S} m_j} U \\
&= \left(\sum_{i \in S} m'_i \right) \left(1 + s(r, \omega) - \frac{\sum_{j \notin S} s(r_j, \omega) m'_j + s(r, \omega) \sum_{j \in S} m'_j}{\sum_{j \notin S} m'_j + \sum_{j \in S} m'_j} \right) \\
&\quad + \frac{\tilde{s}(r, \omega) \sum_{i \in S} m'_i}{\sum_{j \notin S} s(r_j, \omega) m'_j + s(r, \omega) \sum_{j \in S} m'_j} U \\
&= \sum_{i \in S} \hat{\Pi}_i(\mathbf{r}, \mathbf{m}', \omega, U).
\end{aligned}$$

5. *Conditionally truthful for players:* The skill factor Π_i of the payoff function in (3) is truthful by Theorem 1 in Lambert et al. (2008). Furthermore, for $U > 0$ utility becomes proportional to the strictly proper score function given that $U \propto \phi(s(\hat{r}, \omega) - s(r_c, \omega))$. Hence, players can maximize utility by reporting their true belief \mathbf{p} , i.e., $\mathbb{E}_p [U(s(\mathcal{A}(\mathbf{p}, \mathbf{m}), \omega), s(r_c, \omega), \phi)] > \mathbb{E}_p [U(s(\mathcal{A}(\mathbf{r}, \mathbf{m}), \omega), s(p, \omega), \phi)]$ is satisfied for all $\mathbf{r} \neq \mathbf{p}$. Finally, a player does not have enough information and influence on term $\left(\frac{\tilde{s}(r_i, \omega) m_i}{\sum_j \tilde{s}(r_j, \omega) m_j} \right)$ in (3) to create a beneficial arbitrage between skill and utility factors. Thus, we conclude that the payoff $\hat{\Pi}_i$ is conditionally truthful in practical situations.
6. *Truthful for client:* From the design of utility, i.e., $U \propto \phi(s(\hat{r}, \omega) - s(r_c, \omega))$, it is proportional to the strictly proper score function. Furthermore, the predictions of forecasters r_i , for all $i \in \mathcal{I}$ are independent of original forecast report of the client. Writing r_c the client forecast report, we consequently have that the expected payment of the client (allocated utility U) is minimized when the client posts their true belief p , i.e.,

$$\mathbb{E}_p [U(s(\hat{r}, \omega), s(r_c, \omega), \phi)] > \mathbb{E}_p [U(s(\hat{r}, \omega), s(p, \omega), \phi)], \quad \forall r_c \neq p,$$

and

$$p = r_c \quad \implies \quad \mathbb{E}_p [U(s(\hat{r}, \omega), s(r_c, \omega), \phi)] = \mathbb{E}_p [U(s(\hat{r}, \omega), s(p, \omega), \phi)] .$$

7. *Stimulant*: For a player $i \in \mathcal{I}$, the skill factor Π_i of the payoff function in (3) is monotone by Theorem 1 in Lambert et al. (2008) and the utility factor is proportional to wager m_i thus, the payoff $\hat{\Pi}_i$ is stimulant. ■

References

- Acemoglu, D., Makhdoumi, A., Malekian, A., & Ozdaglar, A. (2019). Too much data: Prices and inefficiencies in data markets. URL: <https://economics.mit.edu/files/21108>.
- Agarwal, A., Dahleh, M., & Sarkar, T. (2019). A marketplace for data: An algorithmic solution. In *Proceedings of the 2019 ACM Conference on Economics and Computation* (pp. 701–726).
- Agueh, M., & Carlier, G. (2011). Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43, 904–924.
- Andrade, J. R., & Bessa, R. J. (2017). Improving renewable energy forecasting with a grid of numerical weather predictions. *IEEE Transactions on Sustainable Energy*, 8, 1571–1580.
- Berg, J. E., Nelson, F. D., & Rietz, T. A. (2008). Prediction market accuracy in the long run. *International Journal of Forecasting*, 24, 285–300.
- Bergemann, D., & Bonatti, A. (2019). Markets for information: An introduction. *Annual Review of Economics*, 11, 85–107.
- Browell, J., & Gilbert, C. (2020). Probcast: Open-source production, evaluation and visualisation of probabilistic forecasts. In *2020 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)* (pp. 1–6). IEEE.
- Ghorbani, A., & Zou, J. (2019). Data shapley: Equitable valuation of data for machine learning. In *International Conference on Machine Learning* (pp. 2242–2251).
- Ghosh, A., & Roth, A. (2011). Selling privacy at auction. In *Proceedings of the 12th ACM Conference on Electronic Commerce* (pp. 199–208).
- Gneiting, T. (2011). Quantiles as optimal point forecasts. *International Journal of forecasting*, 27, 197–207.
- Gneiting, T., & Katzfuss, M. (2014). Probabilistic forecasting. *Annual Review of Statistics and Its Application*, 1, 125–151.
- Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102, 359–378.

- Hong, T., Pinson, P., Fan, S., Zareipour, H., Troccoli, A., & Hyndman, R. J. (2016). Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond. *International Journal of Forecasting*, *32*, 896–913.
- Jose, V. R. R., Nau, R. F., & Winkler, R. L. (2009). Sensitivity to distance and baseline distributions in forecast evaluation. *Management Science*, *55*, 582–590.
- Kilgour, D. M., & Gerchak, Y. (2004). Elicitation of probabilities using competitive scoring rules. *Decision Analysis*, *1*, 108–113.
- Knüppel, M., & Krüger, F. (2022). Forecast uncertainty, disagreement, and the linear pool. *Journal of Applied Econometrics*, *37*, 23–41.
- Lambert, N. S., Langford, J., Wortman, J., Chen, Y., Reeves, D., Shoham, Y., & Penno k, D. M. (2008). Self-financed wagering mechanisms for forecasting. In *Proceedings of the 9th ACM Conference on Electronic Commerce* (pp. 170–179).
- Lichtendahl Jr., K. C., Grushka-Cockayne, Y., & Winkler, R. L. (2013). Is it better to average probabilities or quantiles? *Management Science*, *59*, 1594–1611.
- Linde, F., & Stock, W. G. (2011). *Information Markets: A Strategic Guideline for the I-Commerce*. Association for Information Science & Technology.
- Messner, J. W., & Pinson, P. (2019). Online adaptive lasso estimation in vector autoregressive models for high dimensional wind power forecasting. *International Journal of Forecasting*, *35*, 1485–1498.
- Morales, J. M., Conejo, A. J., Madsen, H., Pinson, P., & Zugno, M. (2014). Renewable energy sources—modeling and forecasting. In *Integrating Renewables in Electricity Markets* (pp. 15–56). Springer.
- NUMERAI (2022). NUMERAI, the hardest data science tournament on the planet. <https://numer.ai/>. Accessed: 2022-01-07.
- Ospina, R., & Ferrari, S. L. (2010). Inflated beta distributions. *Statistical Papers*, *51*, 111–126.
- Petropoulos, F., Apiletti, D., Assimakopoulos, V., Babai, M. Z., Barrow, D. K., Taieb, S. B., Bergmeir, C., Bessa, R. J., Bijak, J. et al. (2022). Forecasting: theory and practice. *International Journal of Forecasting*, *38*, 705–871.
- Pinson, P. (2012). Very-short-term probabilistic forecasting of wind power with generalized logit–normal distributions. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, *61*, 555–576.
- Ranjan, R., & Gneiting, T. (2010). Combining probability forecasts. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *72*, 71–91.
- Spiekermann, S., Acquisti, A., Böhme, R., & Hui, K.-L. (2015). The challenges of personal data markets and privacy. *Electronic Markets*, *25*, 161–167.
- Winkler, R. L., Grushka-Cockayne, Y., Lichtendahl Jr., K. C., & Jose, V. R. R. (2019). Probability forecasts and their combination: A research perspective. *Decision Analysis*, *16*, 239–260.
- Winkler, R. L., Muñoz, J., Cervera, J. L., Bernardo, J. M., Blattenberger, G., Kadane, J. B., Lindley, D. V., Murphy, A. H., Oliver, R. M., & Ríos-Insua, D. (1996). Scoring rules and the evaluation of probabilities. *Test*, *5*, 1–60.
- Witkowski, J., Freeman, R., Vaughan, J., Pennock, D., & Krause, A. (2018). Incentive-compatible forecasting competitions. In *Proceedings of the AAAI Conference on Artificial Intelligence*. volume 32.
- Wolfers, J., & Zitzewitz, E. (2006). Prediction markets in theory and practice. URL: <https://www.nber.org/papers/w12083>.

Zhou, Z., Botterud, A., Wang, J., Bessa, R. J., Keko, H., Sumaili, J., & Miranda, V. (2013). Application of probabilistic wind power forecasting in electricity markets. *Wind Energy*, *16*, 321–338.