Noname manuscript No.
(will be inserted by the editor)

Information elicitation for aggregate demand prediction with costly forecasting

Péter Egri

the date of receipt and acceptance should be inserted later

Abstract This paper presents a multiple newsvendor-type purchasing problem where demand forecasts of a number of individual consumer agents can be generated at a price. Firstly, we derive the optimal solution for the model. Next, an information elicitation mechanism is presented that results in the optimal solution despite the autonomous, self-interested participants and the information asymmetry in between consumers and the supplier. Specifically, the incentive compatibility, efficiency, individual rationality and budget balance properties of the mechanism are proved and also illustrated by several numerical experiments.

Keywords Information elicitation · Mechanism design · Scoring rule

1 Introduction

In many domains the satisfaction of *aggregate, uncertain demand* may incur lower costs (or higher profits) than meeting individual demands apiece. Apart from exploiting economies of scale, in such a setting the risk incurred by demand uncertainty may be decreased and shared by the parties; both factors result in lower overall costs. For instance, in *markets of electricity* it is now quite a common practice that a service provider elicits demand forecasts of individual customers, aggregates the forecasts and then plans and contracts for the satisfaction of aggregated demand [16]. A recent study about smart grid related trends also suggests, that consumption related data should be managed on the network-customer frontier independently from the Distribution System Operator (DSO) [13]. The study emphasise the role of aggregators in representing small and medium sized consumers (and on the other side, producers, too) in order to better integrate local supply and load. Furthermore, it suggests changing the simple payment scheme with more complex ones, such as *two-part tariff*.

P. Egri
Fraunhofer Project Center for Production Management and Informatics, Institute for Computer Science and Control, Hungarian Academy of Sciences, Kende u. 13-17, 1111 Budapest, Hungary E-mail: egri@sztaki.mta.hu

Similar situation arises in *supply networks* where a single supplier is responsible for serving a number of different retailers applying *Vendor Managed Inventory* contracts [6]. Here, again, it is better to plan on the basis of aggregated demand than going for the satisfaction of the individual demands of the retailers. Of course, in any case at the time of realization the actual needs of customers may differ from their forecasts. Unless aggregate supply meets aggregate demand this incurs extra cost, and the cost of mismatch—both of surplus and shortage—can be decreased only by improving the precision of forecasts.

Further on, similar information elicitation problems occur also at distributed *sensor networks*, such as aggregating information from autonomous weather sensors, traffic sensors, or users of a social network site [11]. All three coordination problems mentioned above have similar characteristics: there are several autonomous decision makers who possess private information unknown for the others, the future is uncertain, and the emergent equilibrium is not *Pareto-efficient*, i.e., it can be improved for some of the agents while not harming the others. Since we present a general model, we use the word *commodity* throughout the paper, which in specific cases can either be electricity, newspapers, computational capacity in grid computing and so on.

In cases in which forecasting is costly, the consumers want to save money by creating less precise forecasts. However, they would also like to hide the fact from their suppliers that they provide imprecise information for them, therefore they explain the difference between the forecast and realized demand only with external uncertainty. The goal of this paper is to provide a mechanism that inspire agents to increase forecast precision and remove the effect of the information asymmetry.

In what follows we consider the model presented in [16] with the following assumptions: the demands are forecasted by normal distributions and are private information of the agents; the prices—including the forecasting prices—are common knowledge; the agents decide rationally; the demands are independent from each other, cannot be influenced by the agents and are observable by the mechanism. For this setting we present a novel information elicitation mechanism that is incentive compatible, efficient, individually rational and budget balanced. Therefore applying the suggested mechanism results in the optimal cooperative solution in a non-cooperative case.

2 Literature review

The aggregate demand prediction problem investigated in this paper is introduced in [16] for an electricity network, where an aggregator agent elicits consumption forecasts from the home agents and purchases electricity for them. Purchasing can be done on the *forward market* based on the forecasts, and on the *balancing market* based on the realised consumptions. Since the former option is much cheaper, the aggregator wants to collect as precise forecasts as possible in order to avoid excessive trade on the balancing market. However, increasing the precision of the forecasts is not free, it incurs some costs for the home agents. Therefore an optimal balance should be sought between the forecasting and the balancing costs, minimising the total cost of the electricity. The difficulty is that the home agents are concerned only with their forecasting costs, thus without coordination the system gets into a suboptimal equilibrium.

The *information elicitation* or *prediction mechanism design* problem, which has recently come into the focus also of the multiagent research, consists of several agents with some private information about the probability of some stochastic future event, and a centre, whose goal is to obtain and aggregate the dispersed information [8, 17]. A subfield of the information elicitation problem deals with *prediction markets*, where there is always a clear, objective outcome [3]. Such problems can usually be handled by applying so-called *strictly proper scoring rules* [9] that we briefly define here. Let us assume a set D of possible events, and $\mathcal{P}$, a class of probability measures over them. A scoring rule $S : \mathcal{P} \times D \to \Re$ is called strictly proper, if whenever an event $\xi \in D$ is drawn from the distribution $\theta \in \mathcal{P}$, then for any other $\hat{\theta} \neq \theta : \mathbb{E}_\theta[S(\theta, \xi)] < \mathbb{E}_\theta[S(\hat{\theta}, \xi)]$. With other words, the score can be minimised (in expectation), if it is parametrised with the real distribution of the stochastic event.[1] The application of strictly proper scoring rules for information elicitation is straightforward: if the agent with the private information is penalized proportionally, it becomes interested in creating and providing as good a forecast as possible in order to minimise the penalty.

However, it is implicitly assumed in these models that the forecasts can be generated free of charge. On the other hand, when generating or improving the forecast involves some cost, the scoring rule or the mechanism has to be modified accordingly. An example for such a case can be found in [12], where several agents can provide forecasts for the same stochastic variable at different costs, and the authors present a two-stage mechanism including a reverse second-price auction for solving this information elicitation problem.

In [7] a similar model is studied, where the agents get signals about the same stochastic variable. The authors assume that both the forecast precision and the forecasting cost are binary. The goal of their mechanism is to select the experts from a group of agents and use “the wisdom of crowds” in order to create an aggregate forecast. They apply a betting framework, exploiting that people with more accurate information would bet more money on their forecasts, and combine these forecast weighted with the bets.

Apart from the costly forecast generation, there are several other extensions of the information elicitation problem. In some models, the agents have interests in the decision of the centre, therefore they might disclose false information in order to manipulate the decision maker [2]. Such situation is considered e.g., in [6], where the logistic decisions of a supplier can cause shortages at the retailers, which affects their profits. Further difficulties occur, when the objective function of the agents are not known exactly, or when the agents can manipulate the outcome and therefore influence the evaluation of their reported forecasts (the score).

The application of scoring rules in electricity networks (smart grids) has recently become widespread both for aggregating production [15] and consumption forecasts [1] in the scientific literature. Chakraborty et al. investigate a smart grid consisting of consumers, a provider and an electricity generator. The consumers forecast their consumptions based on the fluctuations in the price, while the provider aggregates these forecasts and forwards the resulted predictions to the generator. A multi-layered scoring rule is applied to ensure truthful disclosure of information. The authors then show that the proposed mechanism can decrease both the electricity consumption and the payment of the consumers.

[1] Note that in contrast to the usual notation, for convenience, we minimize the score.

In [16] the authors present a scoring rule based mechanism that fairly distributes the savings among the agents in an electricity network, and prove its individual rationality, incentive compatibility and *ex ante* weak budget balance properties. However, incentive compatibility in that paper relates only to the truthful reporting of forecasts, the achieved forecast precisions are usually not globally optimal, therefore the mechanism is not efficient. In addition, their mechanism has to artificially limit the accepted precision from the agents, thus even the optimal solution could be excluded by the mechanism itself.

In this paper we adapt the same model as [16], but contrarily, the mechanism presented here is efficient, does not bind the accepted precision, moreover it inspires the home agents to generate optimally precise forecasts. The mechanism is based on the scoring rule we applied in [6], but it is extended for considering the costly forecast generation. Furthermore, we prove additional properties of the mechanism in the discussed model like individual rationality and *ex ante* strong budget balance. There is a difference between the model of [16] and our model: they assume that the centre (aggregator) has prior beliefs about the agents' demand, and if it is more precise than the agents' estimations, they will be considered instead of the agents' forecasts. We however assume that the centre cannot have more precise estimation free of charge than the agents, therefore we omitted this aspect from our model. It can be straightforwardly added without modifying the mechanism, but we decided to keep it as simple as possible.

3 The basic model

We assume that the forecast of agent $i \in [1, n]$ $(n > 1)$ is represented by a *normal distribution* with mean m_i and standard deviation $\sigma_i > 0$. This forecast can be generated at price $p_i(\sigma_i) = \alpha_i/\sigma_i^2$, where $1/\sigma_i^2$ is the so-called *precision* of the forecast that can be chosen arbitrarily, and $\alpha_i > 0$ is a constant unit price. This formulation models the possibility of improving the forecast precision by applying more sophisticated and more costly methods, e.g., market research. Note that the $\alpha_i > 0$ parameters are considered to be public knowledge in this paper.

It is also assumed that the consumptions of the agents are independent from each other, therefore the aggregate forecast of the total consumption will also be normally distributed with mean and standard deviation $m = \sum_{i=1}^n m_i$ and $\sigma = \sqrt{\sum_{i=1}^n \sigma_i^2}$.

The centre decides about the purchase quantity q that is bought on the forward market at price c. During the demand realisation period, agent i consumes ξ_i commodity, and if the total realised consumption $\xi = \sum_{i=1}^n \xi_i$ is less or equal than the quantity q bought on the forward market, the centre can provide the necessary commodity. However, if $\xi > q$, the centre has to buy additional commodity on the balancing market at *buy price* b. On the other hand, if $\xi < q$, the surplus can be sold there at *sell price* s. It is natural to assume the following relation between the different prices: $b > c > s$.

The decision variables in the basic model are therefore the σ_i standard deviations and the q purchase quantity. In the decentralised setting however, while the agents still have to decide about the desired precision of their forecasts, they either buy directly from the market—therefore they have to determine q_i purchase quantities—or they delegate the purchasing decision to a centre—which requires a decision about what forecast to tell to the centre $(\hat{m}_i, \hat{\sigma}_i)$—who in turn is responsible for q.

4 Cooperative solution

In the following two subsections we study the *social welfare maximising* solution in the centralised model considering cooperative agents and complete information. The resulted decision problem has two stages: in the first stage the optimal forecast precisions, while in the second stage the optimal purchase quantity is computed. We solve the problem in a *backward induction* manner: firstly, we determine the optimal q assuming a given forecast with a special form of the newsvendor model [14], then we derive the optimal forecast precisions.

4.1 Optimal purchase quantity

In this subsection we consider the normally distributed total consumption $\xi \sim \mathcal{N}(m, \sigma)$, and denote its probability and cumulative distribution functions with ϕ and Φ, respectively.

If the centre purchases quantity q, its resulted *valuation* will be

$$v(q) = -cq - b\max(\xi - q, 0) + s\max(q - \xi, 0) , \tag{1}$$

i.e., the payment for the commodity on the forward market and the (negative or positive) payment of matching supply and demand on the balancing market.

The optimal purchase quantity in the newsvendor model is well known, but for the sake of completeness I repeat it here. (The analysis of the valuation (1) can also be done applying Theorem 9 of [10] with $x = 1$, $y = \xi$, $g(x) = (b - s)x$ and $\alpha = (b - c)/(b - s)$.

The expected value of the valuation, can be expressed in the following form:

$$\begin{aligned}
\mathbb{E}[v(q)] &= -cq - b\mathbb{E}[\max(\xi - q, 0)] + s\mathbb{E}[\max(q - \xi, 0)] \\
&= -cq - b\int_q^{\infty}(x - q)\phi(x)dx + s\int_{-\infty}^{q}(q - x)\phi(x)dx \\
&= -cq - b\left(\int_q^{\infty} x\phi(x)dx - q\int_q^{\infty}\phi(x)dx\right) \\
&\quad + s\left(q\int_{-\infty}^{q}\phi(x)dx - \int_{-\infty}^{q}x\phi(x)dx\right) \\
&= -cq - b\left(m - \int_{-\infty}^{q}x\phi(x)dx - q\left(1 - \Phi(x)\right)\right) \\
&\quad + s\left(q\Phi(x) - \int_{-\infty}^{q}x\phi(x)dx\right) \\
&= -cq + b(q - m) + (b - s)\left(\int_{-\infty}^{q}x\phi(x)dx - q\Phi(q)\right) .
\end{aligned} \tag{2}$$

Since this function is concave $((\mathbb{E}[v(q)])'' = -(b-s)\phi(q) < 0)$, the optimal q^* can be determined by the first derivative test

$$\mathbb{E}[v(q^*)]' = b - c - (b-s)\Phi(q^*) = 0 \ , \tag{3}$$

which results in

$$q^* = \Phi^{-1}\left(\frac{b-c}{b-s}\right) = m + \sigma\sqrt{2}\,\mathrm{erf}^{-1}\left(\frac{b-2c+s}{b-s}\right) \ , \tag{4}$$

where

$$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}}\int_0^x e^{-t^2}\,dt \ , \tag{5}$$

is the *Gauss error function* and $\mathrm{erf}^{-1}(x)$ denotes its inverse. The expected valuation using the *optimal purchase quantity* can then be expressed by substituting (4) into (2) as

$$\mathbb{E}[v(q^*)] = -cm - K_{cbs}\sigma \ , \tag{6}$$

where

$$K_{cbs} = \frac{(b-s)e^{-\left(\mathrm{erf}^{-1}\left(\frac{b-2c+s}{b-s}\right)\right)^2}}{\sqrt{2\pi}} \ . \tag{7}$$

Note that K_{cbs} depends only on the cost parameters c, b and s, furthermore $K_{cbs} > 0$.

4.2 Optimal forecast precisions

We now examine the optimal forecast precisions, or equivalently, the optimal standard deviations. The *utility* will be the valuation minus the forecasting price of the agents:

$$U(\sigma_1, \ldots, \sigma_n) = v(q^*) - \sum_{i=1}^{n} p_i(\sigma_i) \ . \tag{8}$$

Theorem 1 *The expected utility function* $\mathbb{E}[U(\sigma_1, \ldots, \sigma_n)]$ *is concave.*

Proof We prove that

$$-\mathbb{E}[U(\sigma_1, \ldots, \sigma_n)] = cm + K_{cbs}\sqrt{\sum_{i=1}^{n}\sigma_i^2} + \sum_{i=1}^{n}\frac{\alpha_i}{\sigma_i^2} \tag{9}$$

is convex, since it is a sum of convex functions. The first term cm is a constant. The second term is the multiplication of a positive number and $\sqrt{\sum_{i=1}^n \sigma_i^2}$. This latter is the L^2 (Euclidean) norm of the vector $(\sigma_1, \ldots, \sigma_n)$, whose convexity follows from the Minkowski (triangle) inequality. The convexity of the rest $(\sum_{i=1}^n \alpha_i/\sigma_i^2)$ can be proven by the second partial derivative test: its Hessian matrix is diagonal $(\mathrm{diag}(6\alpha_i/\sigma_i^4))$, which is positive definite. $\square$

Theorem 2 *The maximal expected utility is assumed at the stationary point* $(\sigma_1^*, \ldots, \sigma_n^*)$, *where*

$$\sigma_i^* = \sqrt[4]{\alpha_i} \sqrt[3]{\frac{2}{K_{cbs}} \sqrt{\sum_{j=1}^{n} \sqrt{\alpha_j}}} \, . \tag{10}$$

Proof The first partial derivatives of the expected utility function are

$$\frac{\partial \mathbb{E}[U(\sigma_1, \ldots, \sigma_n)]}{\partial \sigma_i} = -K_{cbs} \frac{\sigma_i}{\sqrt{\sum_{j=1}^{n} \sigma_j^2}} + \frac{2\alpha_i}{\sigma_i^3} \qquad (i = 1, \ldots, n) \, , \tag{11}$$

therefore we get the following system of equations

$$K_{cbs} \frac{\sigma_i^*}{\sqrt{\sum_{j=1}^{n} (\sigma_j^*)^2}} = \frac{2\alpha_i}{(\sigma_i^*)^3} \qquad (i = 1, \ldots, n) \, , \tag{12}$$

or equivalently

$$K_{cbs} \frac{(\sigma_i^*)^4}{2\alpha_i} = \sqrt{\sum_{j=1}^{n} (\sigma_j^*)^2} \qquad (i = 1, \ldots, n) \, . \tag{13}$$

Note that the right hand side of the equations are the same, therefore for arbitrary i and j we have

$$K_{cbs} \frac{(\sigma_i^*)^4}{2\alpha_i} = K_{cbs} \frac{(\sigma_j^*)^4}{2\alpha_j} \, , \tag{14}$$

hence we have the following relationship among the variables

$$\sigma_j^* = \sqrt[4]{\frac{\alpha_j}{\alpha_i}} \sigma_i^* \, . \tag{15}$$

Then σ_i^* can be computed from (13) considering (15)

$$\begin{aligned} (\sigma_i^*)^4 &= \frac{2\alpha_i}{K_{cbs}} \sqrt{\sum_{j=1}^{n} (\sigma_j^*)^2} = \frac{2\alpha_i}{K_{cbs}} \sqrt{\sum_{j=1}^{n} (\sqrt[4]{\frac{\alpha_j}{\alpha_i}} \sigma_i^*)^2} \\ &= \frac{2\alpha_i}{K_{cbs}} \sigma_i^* \sqrt{\sum_{j=1}^{n} \sqrt{\frac{\alpha_j}{\alpha_i}}} \qquad (i = 1, \ldots, n) \, . \end{aligned} \tag{16}$$

From this

$$\sigma_i^* = \sqrt[3]{\frac{2\alpha_i}{K_{cbs}} \sqrt{\sum_{j=1}^{n} \sqrt{\frac{\alpha_j}{\alpha_i}}}} \, . \tag{17}$$

and the statement of the theorem follows. □

5 Non-cooperative (baseline) solution

In the previous section we have seen that in an optimal situation, any agent i generates forecast with precision defined by Theorem 2, while the centre orders the quantity defined by (4) on the forward market. However, in reality the agents are non-cooperative in the sense that they do not intend to maximise the utility of the whole system, but rather their own utility. In this section we investigate the decisions of agent i, if it does not participate in the cooperation, but purchases directly from the market.

Let q_i denote the quantity bought on the forward market, then analogously to the derivation of Section 4.1, its valuation becomes

$$v_i(q_i) = -cq_i - b\max(\xi_i - q_i, 0) + s\max(q_i - \xi_i, 0) \ , \tag{18}$$

and its expected value

$$\mathbb{E}[v_i(q_i)] = -cq_i + b(q_i - m_i) + (b - s)\left(\int_{-\infty}^{q_i} x\phi_i(x)dx - q_i\Phi_i(q_i)\right) \ , \tag{19}$$

where ϕ_i and Φ_i are the probability and cumulative distribution functions of ξ_i, respectively. This yields an optimal purchase quantity

$$q_i^* = m_i + \sigma_i\sqrt{2}\,\mathrm{erf}^{-1}\left(\frac{b - 2c + s}{b - s}\right) \ , \tag{20}$$

and

$$\mathbb{E}[v_i(q_i^*)] = -cm_i - \sigma_i K_{cbs} \ . \tag{21}$$

Then the optimal non-cooperative forecast precision can be derived analogously to Section 4.2 using the non-cooperative utility function

$$\bar{U}_i(\sigma_i) = v_i(q_i^*) - p_i(\sigma_i) \ : \tag{22}$$

$$\bar{\sigma}_i^* = \sqrt[3]{\frac{2\alpha_i}{K_{cbs}}} \ . \tag{23}$$

6 Non-cooperative mechanism

In this section we continue to consider agents maximising their own utility, but instead of purchasing directly from the market, they interact with the centre. The process is as follows.

1. Agent i decides about the forecast precision (σ_i) and generates normally distributed forecast (m_i, σ_i). The precision may differ from the optimal one of Theorem 2, because the agent might want to spare some money on forecasting.
2. Agent i discloses its forecast to the centre. Since the disclosure can be dishonest, we denote the reported forecast with $(\hat{m}_i, \hat{\sigma}_i)$.
3. The centre decides about the purchase quantity q.
4. The demands $\xi_1, \ldots, \xi_n$ realise. Simultaneously with the supply, the centre trades the surplus or the deficit on the balancing market.
5. The agents pay $t_i(\hat{m}_i, \hat{\sigma}_i, \xi_i)$ based on their reported forecasts and realised demand to the centre.

We are looking for a mechanism—defined by the payment or *transfer function*—that fulfils the following four key properties:

- It is *incentive compatible* that implies in this case two things: (i) the agents create optimally precise forecasts, and (ii) they report these forecasts to the centre truthfully, if they want to maximise their expected utility.
- The mechanism is *efficient*, i.e., the centre purchases the optimal quantity of commodity from the market.
- The mechanism is *individually rational*, i.e., the expected utility of the agents are not less than in the baseline solution.
- Finally, it is *budget balanced* meaning that the mechanism does not require external sources for financing deficit or surplus on the long term, or in other words, the expected utility of the centre is zero.

We assume that the market prices c, b and s are common knowledge, since the agents can buy directly from the market if they do not join the mechanism. We suggest the following transfer function that consists of two terms: (i) the payment for the purchased commodity, and (ii) a penalty for the forecast error, based on a scoring rule:

$$t_i(\hat{m}_i, \hat{\sigma}_i, \xi_i) = c\xi_i + \gamma_i \left(\frac{(\xi_i - \hat{m}_i)^2}{\hat{\sigma}_i} + \hat{\sigma}_i \right) , \tag{24}$$

where $\hat{m}_i$ and $\hat{\sigma}_i$ are the communicated forecast of agent i, ξ_i is its realised consumption, while γ_i is a positive constant. As we shall soon see, in order to achieve the four required properties mentioned above, γ_i cannot be arbitrary, but it should assume a specific value consisting only of publicly known parameters c, b, s and α_j $(j \in [1, n])$:

$$\gamma_i = \frac{K_{cbs} \sqrt[4]{\alpha_i}}{2\sqrt{\sum_{j=1}^{n} \sqrt{\alpha_j}}} . \tag{25}$$

In what follows we examine one by one whether and how the non-cooperative mechanism exhibits the four basic properties required above.

6.1 Incentive compatibility

In the first two phases the agents generate forecasts m_i and σ_i, next, they report the forecasts to the centre. Let us first examine the latter phase with generated forecast (m_i, σ_i), reported forecast $(\hat{m}_i, \hat{\sigma}_i)$, and realized demand ξ_i. Note that at this point—since $p_i(\sigma_i)$ has already been invested in the forecast—we consider only the expected payment here.

Theorem 3 *The unique optimal solution for minimising the expected payment is $\hat{m}_i = m_i$ and $\hat{\sigma}_i = \sigma_i$, therefore the agents are inspired to report the forecasts truthfully.*

Proof The statement of the theorem follows from Theorem 7 of [10] with $\phi(x) = x^2$, but for the sake of completeness I include a more detailed derivation below.

The expected payment is

$$\begin{aligned}\mathbb{E}[t_i(\hat{m}_i,\hat{\sigma}_i,\xi_i)] &= c\mathbb{E}\left[\xi_i\right] + \gamma_i \mathbb{E}\left[\frac{(\xi_i-\hat{m}_i)^2}{\hat{\sigma}_i}+\hat{\sigma}_i\right] \\ &= cm_i + \gamma_i \mathbb{E}\left[\frac{\xi_i^2+\hat{m}_i^2-2\hat{m}_i\xi_i}{\hat{\sigma}_i}+\hat{\sigma}_i\right] \\ &= cm_i + \gamma_i \frac{\mathbb{E}\left[\xi_i^2\right]+\hat{m}_i^2-2\hat{m}_i\mathbb{E}[\xi_i]}{\hat{\sigma}_i}+\hat{\sigma}_i \\ &= cm_i + \gamma_i \left(\frac{m_i^2+\sigma_i^2+\hat{m}_i^2-2\hat{m}_i m_i}{\hat{\sigma}_i}+\hat{\sigma}_i\right) , \end{aligned} \tag{26}$$

where we have applied the identity $\mathbb{E}\left[\xi_i^2\right] = m_i^2+\sigma_i^2$. The partial derivative of the expected payment by $\hat{m}_i$ is

$$\frac{\partial \mathbb{E}[t_i(\hat{m}_i,\hat{\sigma}_i,\xi_i)]}{\partial \hat{m}_i} = \gamma_i \left(\frac{2\hat{m}_i - 2m_i}{\hat{\sigma}_i}\right) , \tag{27}$$

which equals zero iff $\hat{m}_i = m_i$, independently from the value of $\hat{\sigma}_i$. This yields the minimum, since the expected payment is convex in $\hat{m}_i$:

$$\frac{\partial^2 \mathbb{E}[t_i(\hat{m}_i,\hat{\sigma}_i,\xi_i)]}{\partial \hat{m}_i^2} = \gamma_i \left(\frac{2}{\hat{\sigma}_i}\right) \geq 0 . \tag{28}$$

For calculating the other partial derivative, we already exploit that $\hat{m}_i = m_i$ in an optimal solution:

$$\frac{\partial \mathbb{E}[t_i(m_i,\hat{\sigma}_i,\xi_i)]}{\partial \hat{\sigma}_i} = \gamma_i \left(-\frac{\sigma_i^2}{\hat{\sigma}_i^2}+1\right) , \tag{29}$$

that equals zero iff $\hat{\sigma}_i = \sigma_i$, which is the minimum, since the expected payment is convex also in $\hat{\sigma}_i$:

$$\frac{\partial^2 \mathbb{E}[t_i(m_i,\hat{\sigma}_i,\xi_i)]}{\partial \hat{\sigma}_i^2} = \gamma_i \frac{\sigma_i^2}{\hat{\sigma}_i^3} \geq 0 . \tag{30}$$

□

Let us now examine the first phase knowing that later the forecasts will be reported truthfully. In this case, the expected payment can be derived from (26): $\mathbb{E}[t_i(m_i,\sigma_i,\xi_i)] = cm_i + 2\gamma_i\sigma_i$, thus the expected utility of agent i becomes:

$$\mathbb{E}[U_i(\sigma_i)] = -\mathbb{E}[t_i(m_i,\sigma_i,\xi_i)] - p_i(\sigma_i) = -cm_i - 2\gamma_i\sigma_i - \frac{\alpha_i}{\sigma_i^2} . \tag{31}$$

Since this is concave in σ_i ($\mathbb{E}[U_i(\sigma_i)]'' = -6\alpha_i/(\sigma_i)^4 < 0$), the optimal standard deviation can be determined by $\mathbb{E}[U_i(\sigma_i^*)]' = -2\gamma_i + 2\alpha_i/(\sigma_i^*)^3 = 0$, which yields

$$\sigma_i^* = \sqrt[3]{\frac{\alpha_i}{\gamma_i}} = \sqrt[3]{\frac{\alpha_i 2\sqrt{\sum_{j=1}^n \sqrt{\alpha_j}}}{K_{cbs}\sqrt[4]{\alpha_i}}} = \sqrt[4]{\alpha_i}\sqrt[3]{\frac{2}{K_{cbs}}\sqrt{\sum_{j=1}^n \sqrt{\alpha_j}}} , \tag{32}$$

where we applied (25). The result is the optimal forecast precision of Theorem 2.

All in all, applying payment function (24) with parameter (25) inspires the utility maximising agents to (i) achieve the optimal forecast precision, and (ii) report the forecast truthfully.

6.2 Efficiency

The utility of the centre will be the collected transfer payments minus the price of the commodity bought, plus the price of the commodity eventually sold—this last two items are summed up by the valuation (1), i.e.,

$$U_a(q) = \sum_{i=1}^{n} t_i(m_i, \sigma_i, \xi_i) + v(q) \ . \tag{33}$$

Since the first term is independent from the decision variable q, thus the centre intends to maximise the expected valuation that results in the optimal quantity derived in (4).

6.3 Individual rationality

In this subsection we compare the agents' utility applying the mechanism and the baseline solution. The expected utility of agent i in the baseline solution follows from (22) and (23):

$$\begin{aligned}\mathbb{E}[\bar{U}_i(\bar{\sigma}_i^*)] = \mathbb{E}[v_i(q_i^*)] - p_i(\bar{\sigma}_i^*) &= -cm_i - \bar{\sigma}_i^* K_{cbs} - \frac{\alpha_i}{(\bar{\sigma}_i^*)^2} \\ &= -cm_i - \frac{3}{\sqrt[3]{4}} \sqrt[3]{\alpha_i K_{cbs}^2} \ .\end{aligned} \tag{34}$$

However, by using the mechanism (10), (25) and (31) are valid, thus

$$\begin{aligned}\mathbb{E}[U_i(\sigma_i^*)] &= -\mathbb{E}[t_i(m_i, \sigma_i^*, \xi_i)] - p_i(\sigma_i^*) = -cm_i - 2\gamma_i \sigma_i^* - \frac{\alpha_i}{(\sigma_i^*)^2} \\ &= -cm_i - \frac{3}{\sqrt[3]{4 \sum_{j=1}^{n} \sqrt{\frac{\alpha_j}{\alpha_i}}}} \sqrt[3]{\alpha_i K_{cbs}^2} \ .\end{aligned} \tag{35}$$

Comparing (34) and (35) shows that $\mathbb{E}[U_i(\sigma_i^*)] > \mathbb{E}[\bar{U}_i(\bar{\sigma}_i^*)]$ (since $\sum_{j=1}^{n} \sqrt{\alpha_j/\alpha_i} > 1$ if $n > 1$), i.e., the expected utility using the mechanism is always greater than without the mechanism. Hence, the agents have an incentive to use the service of the centre mechanism when meeting their individual demand for the commodity.

6.4 Budget balance

The expected utility of the centre can be determined by substituting (6) and (10) into (33):

$$\begin{aligned}\mathbb{E}[U_a(q^*)] &= \mathbb{E}\left[\sum_{i=1}^{n} t_i(m_i, \sigma_i^*, \xi_i)\right] + \mathbb{E}[v(q^*)] \\ &= \sum_{i=1}^{n} \left(cm_i + 2\gamma_i \sigma_i^*\right) - cm - K_{cbs} \sqrt{\sum_{i=1}^{n} (\sigma_i^*)^2}\end{aligned}$$

$$
\begin{aligned}
&= \sum_{i=1}^{n} \frac{K_{cbs}\sqrt{\alpha_i}}{\sqrt{\sum_{j=1}^{n}\sqrt{\alpha_j}}} \sqrt[3]{\frac{2\sqrt{\sum_{j=1}^{n}\sqrt{\alpha_j}}}{K_{cbs}}} - K_{cbs}\sqrt{\sum_{i=1}^{n}\sqrt{\alpha_i}\sqrt[3]{\frac{4\sum_{j=1}^{n}\sqrt{\alpha_j}}{K_{cbs}^2}}} \\
&= \frac{K_{cbs}}{\sqrt{\sum_{j=1}^{n}\sqrt{\alpha_j}}} \sqrt[3]{\frac{2\sqrt{\sum_{j=1}^{n}\sqrt{\alpha_j}}}{K_{cbs}}} \sum_{i=1}^{n}\sqrt{\alpha_i} - K_{cbs}\sqrt{\sqrt[3]{\frac{4\sum_{j=1}^{n}\sqrt{\alpha_j}}{K_{cbs}^2}}\sum_{i=1}^{n}\sqrt{\alpha_i}} \\
&= K_{cbs}\sqrt[3]{\frac{2\sqrt{\sum_{j=1}^{n}\sqrt{\alpha_j}}}{K_{cbs}}}\sqrt{\sum_{i=1}^{n}\sqrt{\alpha_i}} - K_{cbs}\sqrt[3]{\frac{2\sqrt{\sum_{j=1}^{n}\sqrt{\alpha_j}}}{K_{cbs}}}\sqrt{\sum_{i=1}^{n}\sqrt{\alpha_i}} \\
&= 0\,, \qquad (36)
\end{aligned}
$$

thus the mechanism is *ex ante* (in expectation) budget balanced. In other words, no payments or debts are accumulated at the centre on the long term.

7 Benefit sharing with the centre

In Section 6 we have proved the budget balance property of the mechanism, which is desired from the theoretical point of view, however, this means that the centre does not profit from providing the service for the agents. This could make the whole model practically inapplicable, but with a small modification the positive utility of the centre can be assured.

Let $B_i = \mathbb{E}[U_i(\sigma_i^*) - \bar{U}_i(\bar{\sigma}_i^*)]$ denote the expected *benefit* of agent i, which is strictly positive, as we have proven in Section 6.3. Also note that B_i is independent from the decision variables; it characterises the achievable gain applying a mechanism. If the agents are willing to share their benefits with the centre, its expected utility will be positive, while maintaining incentive compatibility, efficiency and individual rationality.

Let us introduce a sharing ratio $\delta_i \in (0, 1)$ and modify the payment function to the following:

$$
t_i(\hat{m}_i, \hat{\sigma}_i, \xi_i) = c\xi_i + \gamma_i\left(\frac{(\xi_i - \hat{m}_i)^2}{\hat{\sigma}_i} + \hat{\sigma}_i\right) + \delta_i B_i\,. \qquad (37)
$$

The additional third term is independent from the decision variables of the model, causes only a constant shift in the utilities, therefore serves as an instrument for *benefit balancing* [5].

8 Numerical illustrations

In this section we illustrate the properties of the mechanism (without benefit sharing) on a numerical example. We set the parameters similarly to [16], i.e., $b = 170$, $c = 100$ and $s = 50$. We consider uniformly distributed α_i in $[0.01, 0.1]$ and m_i in $[30, 50]$. We also considered different number of agents $n \in \{1, \ldots, 100\}$, and for each case we run 10000 independent simulations and computed the average values.

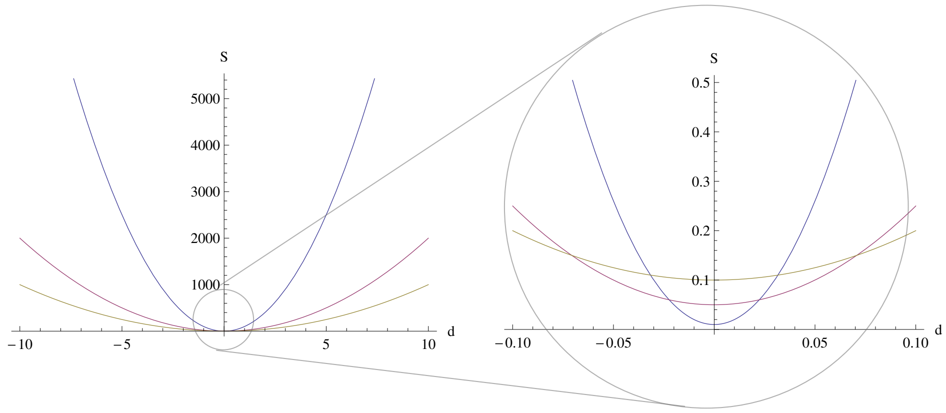

Fig. 1 The score in function of the difference of expected and realised demand.

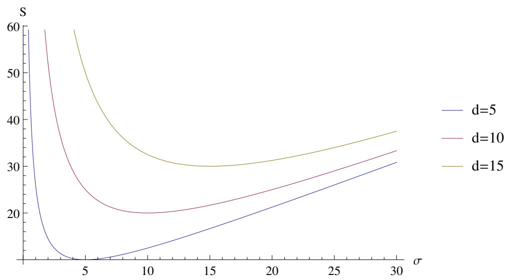

Fig. 2 The score in function of the standard deviation.

8.1 Scoring rule

The next figures illustrate the function $S(d, \sigma) = d^2/\sigma + \sigma$ that we applied in (24), where $d = \xi - m$ is the difference between the expected and the realised demand, and σ is the communicated standard deviation. Fig. 1 shows the score when we let d range from -10 to 10 for different σ values. It can be seen that the larger the difference between the forecasted and the realised demand, the larger the penalty the agent should pay to the centre for the inappropriate forecast.

Fig. 2 illustrates the score when σ changes, considering three different d values. In accordance with our theoretical results, for each d there is a unique σ minimising the score; any lower value suggests an underestimated, any higher an overestimated standard deviation. This property is used in the mechanism to inspire the agents achieving the optimal precision derived in Theorem 2.

Fig. 3 shows the score in a three dimensional plot in order to better observe the relation between d, σ and S.

8.2 Utility of the agents

Fig. 4 shows the difference between the average utility of an agent in a mechanism of n agents and the average utility of an agent without the mechanism resulted during the simulation. It can be observed that the larger n the larger the difference

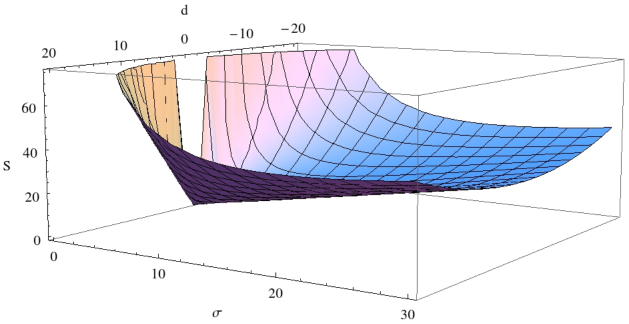

Fig. 3 Illustration of the score. The penalty for the difference of expected and realised demand is proportional to the declared precision.

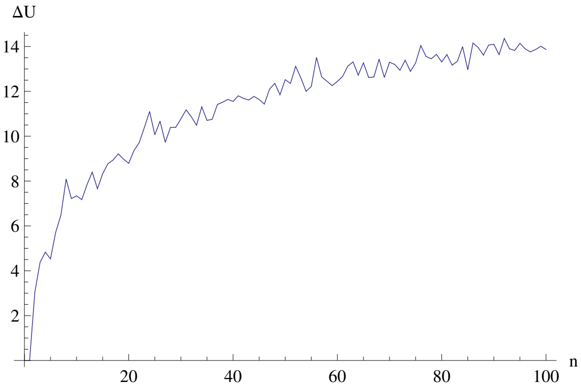

Fig. 4 Difference between $\bar{U}_i$ and U_i.

is in the utilities, but the curve becomes less steeper as n increases. This means that an individual agent can achieve relatively high improvement by joining with other agents in order to apply the purchase mechanism. However, the more agents participate in the mechanism, the less improvement can be achieved by increasing the number of agents. This can be interpreted in such a way, that although increasing the size of a mechanism is always beneficial, some efficiency may be sacrificed by forming several medium-sized mechanism instead of the *grand coalition*, one with all of the agents, if other considerations (administration cost, robustness, feasibility) are also present.

Another illustration of this phenomenon can be seen on Fig. 5, where the contribution of an additional agent to the average utility is presented.

8.3 Utility of the centre

The cumulated utility of the centre is shown on Fig. 6 through 10000 simulation steps considering 100 agents. Although it oscillates around zero as the *ex ante* budget balance property predicted, it can go far from it. Since this is usually undesirable, we are interested not only in its expected value, but also in its standard deviation, which is presented on Fig. 7.

It can be seen that by increasing the number of agents, the standard deviation of the centre's utility also increases. Therefore finding an appropriate mechanism

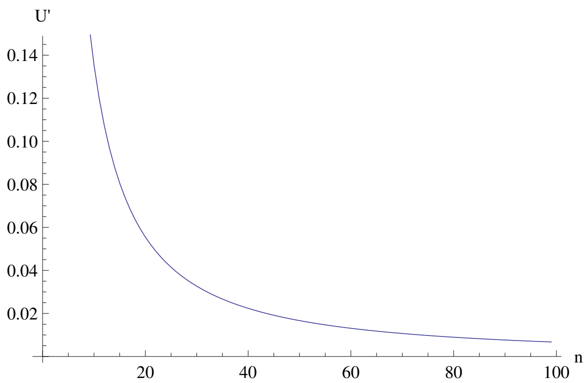

Fig. 5 Marginal utility of increasing the number of agents.

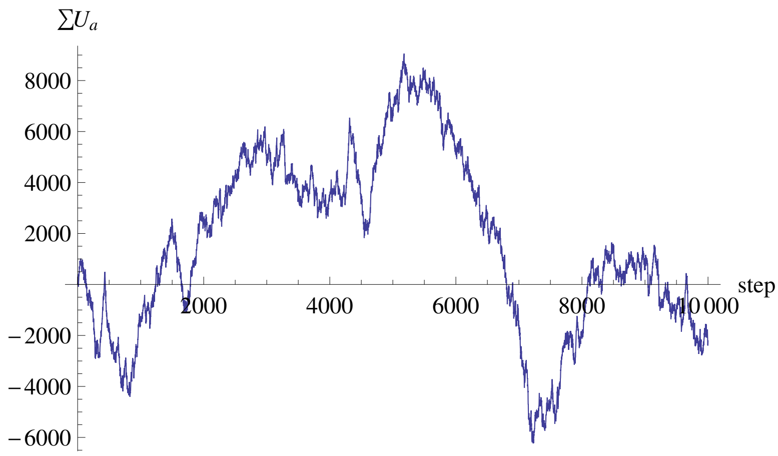

Fig. 6 Cumulated utility of the centre ($n = 100$).

size includes finding a trade-off between the agents' expected utility and the variance of the centre's utility.

9 Conclusions and future work

In this paper we have investigated the purchasing problem for satisfying uncertain demand with costly forecasting. We have examined the optimal cooperative solution—which is unrealistic, since the agents disregard their interests and minimise the overall cost—and the non-cooperative baseline solution—which is suboptimal, since agents disregard the others' utility. To resolve this problem, we have presented an information elicitation mechanism including a centre, where the decisions of rational, self-interested agents result in the optimal solution of the cooperative case. Since the previous works in this field could not provide an efficient mechanism for this model, our result is a novel contribution of this paper.

We have presented both formal proofs and numerical illustrations for the key properties like incentive compatibility, individual rationality, efficiency and budget balance. We have shown that the agents benefit from participating in the proposed

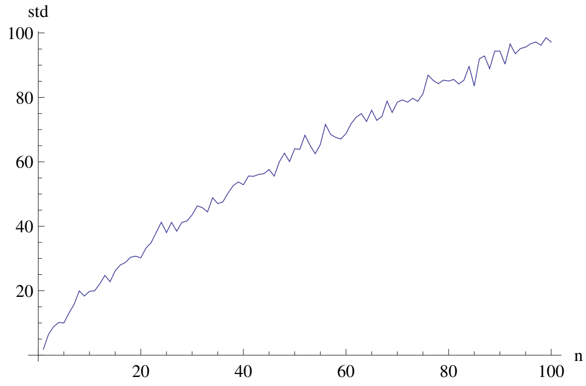

Fig. 7 Sample standard deviation of the centre's utility.

mechanism compared to managing their purchase individually, and the total benefit is distributed among the participants, since the expected utility of the centre is zero. The payment an agent should pay is *fair* in the sense that it is independent from other agents' forecasts or realised demands, therefore these parameters do not have to be common knowledge, which is desirable from the privacy point of view. Numerical experiments have shown that a mechanism with moderately many agents is almost as efficient as a larger one, which facilitates finding trade-offs when considering additional optimisation criteria.

Finally, the presented work can be extended into several directions. One can consider more complex planning problems than the newsvendor model, where an optimal solution cannot be guaranteed for problems of realistic size [4]. By applying approximations for such cases, not only the efficiency, but also the incentive compatibility can be lost, which necessitates further research in this direction. An other extension can be considering less common knowledge. A straightforward possibility is to assume privately known forecasting cost parameters (α_i), which rules out the current payment function (24). An other potential issue is that the realised demand (ξ_i) may not be directly observable for the centre, or the agents can influence its value, which introduces a new challenge into the model. In addition, the numerical illustrations suggest further theoretical investigations of the relationship between the number of agents in the system and the benefit of an individual agent.

Acknowledgements This work has been supported by the Hungarian Scientific Research Fund (OTKA), Grant No. 113038, by the European Union 7th Framework Programme Project No. NMP 2013-609087, Shock-robust Design of Plants and their Supply Chain Networks (RobustPlaNet), and by the János Bolyai scholarship No. BO/00659/11/6. The author would like to thank József Váncza for his help and support during the research and to the anonymous reviewers for their valuable remarks.

References

1. Chakraborty, S., Ito, T. & Senjyu, T. (2014). Smart pricing scheme: A multi-layered scoring rule application. *Expert Systems with Applications*, 41, 8, 3726–3735.
2. Chen, Y., Gao, X. A., Goldstein, R., & Kash, I. A. (2014). Market manipulation with outside incentives. *Autonomous Agents and Multi-Agent Systems*, in print. DOI 10.1007/s10458-014-9249-1.

3. Conitzer, V. (2009). Prediction markets, mechanism design, and cooperative game theory. *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, 101–108 .
4. Duan, L., Doğru, M. K., Özen, U., Beck, J. C. (2012). A negotiation framework for linked combinatorial optimization problems. *Autonomous Agents and Multi-Agent Systems*, 25, 158–182.
5. Egri, P., Döring, A., Timm, T., & Váncza, J. (2011). Collaborative planning with benefit balancing in Dynamic Supply Loops. *CIRP Journal of Manufacturing Science and Technology*, 4, 3, 226–233.
6. Egri, P., & Váncza, J. (2013). A distributed coordination mechanism for supply networks with asymmetric information. *European Journal of Operational Research*, 226, 3, 452–460.
7. Fang, F., Stinchcombe, M., Whinston, A. (2007). ,,Putting your money where your mouth is" – A betting platform for better prediction. *Review of Network Economics*, 6, 2, 214–238.
8. Farfel, J., & Conitzer, V. (2011). Aggregating value ranges: preference elicitation and truthfulness. *Autonomous Agents and Multi-Agent Systems*, 22, 127–150.
9. Gneiting, T. & Raftery, A.E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102, 477, 359–378.
10. Gneiting, T. (2011). Making and evaluating point forecasts. *Journal of the American Statistical Association*, 106, 494, 746–762.
11. Papakonstantinou, A. (2010). *Mechanism design for eliciting costly observations in next generation citizen sensor networks.* PhD thesis, University of Southampton.
12. Papakonstantinou, A., Rogers, A., Gerding, E. H., & Jennings, N. R. (2011). Mechanism design for the truthful elicitation of costly probabilistic estimates in distributed information systems. *Artificial Intelligence*, 175, 2, 648–672.
13. Mallet, P., Granström, P.-O., Hallberg, P., Lorenz, G. & Mandatova, P. (2014). Power to the People! European Perspectives on the Future of Electric Distribution. *IEEE power & energy magazine*, 12, 2, 51–64.
14. Qin, Y., Wang, R., Vakharia, A.J., Chen, Y. & Seref, M.M.H. (2011). The newsvendor problem: Review and directions for future research. Invited review. *European Journal of Operational Research*, 213, 2, 361–374.
15. Robu, V., Kota, R., Chalkiadakis, G., Rogers, A., & Jennings, N. R. (2012). Cooperative virtual power plant formation using scoring rules. *Proceedings of the AAAI-12: Twenty-Sixth Conference on Artificial Intelligence*, 370–376.
16. Rose, H., Rogers, A., & Gerding, E. H. (2012). A scoring rule-based mechanism for aggregate demand prediction in the smart grid. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, 661–668.
17. Zohar, A., & Rosenschein, J. S. (2008). Mechanisms for information elicitation. *Artificial Intelligence*, 172, 16–17, 1917–1939.