Forecast Elicitation and Frequency Control

Nawaaz Khalfan   Guillaume Roger

November 18, 2024

Frequency control

On an electricity grid there are generators who supply power and consumers who draw it.

At any point in time, the amount supplied must equal to that drawn. If not, the system *frequency* adjusts to equate power in with power out.

A change in *frequency* is costly to the grid:

- degrades generator and consumer use as well as the grid itself
- at worst disrupts service i.e. blackouts
- necessitates dispatch on and offload from the grid

As a result, modern electricity grids are typically run by a market operator or scheduler.

Causer pay rules

Of recent concern, is how the scheduler is incentivising participation on the grid.

Among other features, the market operator in Australia offers contracts to generators with **causer pay rules**. Generally these:

1. pay rewards to generators who over-produce when net supply is negative and under-produce when net supply is positive
2. collect fines from generators who over-produce when net supply is positive and under-produce when net supply is negative

This is puzzling to us:

- if the contracts are designed to elicit a generator response, why not just run a normal market and let price dictate supply?
- if the contracts are designed to elicit information, why pay generators conditional on the types of their competitors?

Information elicitation

For a market to necessitate a scheduler, there must be some uncertainty regarding the ultimate supply (and demand), that price cannot resolve, and there must be an additional cost than simply the foregone trade.

Let a market that has these features be a **stochastic** and **scheduled** market, and it seems a natural fit for electricity.

In such a market, the must be a premium for any private information that market participants have concerning this uncertainty as it would improve the schedulers ability to match supply and demand.

This is of particular importance for the on-boarding of renewable generators who's generation costs are very low but who's supply is inherently stochastic and thus costly for the grid.

A simple scheduling problem

Consider a market with,

- consumers described by a total benefit function, $B(x)$, where x the total amount consumed,
- n producers, each of whom at the point of trade produce a stochastic amount X_i according to F_i at zero marginal cost, and
- and a scheduler who selects a subset of producers to participate, $S \subseteq N$, so that $x(S) = \sum_{i \in S} x_i$ is the realised production level.

Additionally suppose the scheduler must set a target production level, $\bar{x} \in \mathbb{R}$, so that their incentives are described by the benefit of servicing $\bar{x}$ consumers net of the quadratic cost of meeting this target:

$$v(S, \bar{x} \mid x) = B(\bar{x}) - c \cdot (x - \bar{x})^2$$

Private information

Going forward we'd like to think of F_i as being the private information of the producers, and subsequently the scheduler being able to incentivise their reporting - a *mechanism design* problem.

The timing of their decisions are then given as:

1. The scheduler commits to a mechanism, and F_i is drawn for all $i \in N$
2. Reports are made and collected, and the scheduler selects $S \subseteq N$ and $\bar{x}$
3. x_i is realised for all $i \in S$, the residual is $|x - \bar{x}|$ is amended at quadratic cost $c > 0$, and $\bar{x}$ is consumed

The scheduler selects S and $\bar{x}$ at the interim stage, so they select the mechanism that maximises the interim payoff

$$V(S, \bar{x}) = B(\bar{x}) - c \cdot \mathbb{E}[(x - \bar{x})^2]$$

Efficient scheduling

Consider the first best or *efficient* schedule.

Observation 1

$\bar{x}(S) \geq \mathbb{E}[x(S)]$ and thus $\mathbb{E}[(x - \bar{x}(S))^2] \geq \mathbb{V}[x(S)]$

Given the efficient target, let the final value function be given by:

$$V = \max_{S \subseteq N} B(\bar{x}(S)) - c \cdot \mathbb{E}[(x - \bar{x}(S))^2]$$

Going forward, suppose F_i is the binary distribution:

$$X_i = \begin{cases} 1 & \text{w.p.} \quad p_i \\ 0 & \text{w.p.} \quad 1 - p_i \end{cases}$$

Observation 2

V is convex and increasing in p_i

Scoring rules

The private information that the producers have are *beliefs* about future states of the world which differ from the usual auction setup.

The study of incentivising the reporting of such beliefs is typically attributed to McCarthy (1956) and Savage (1971), and contracts designed to elicit truthful reporting of these beliefs are known in computer science as **scoring rules**.

Definition 1 (Gneiting and Raftery (2007))

A *scoring rule* is any real valued function $S : P \times \{0,1\} \to \mathbb{R}$ where P is the type space. A scoring rule is *proper* if $S(p|p) \geq S(q|p)$ for all $p, q \in P$.

Theorem 1 (McCarthy (1956), Savage (1971))

The scoring rule S is proper if and only if $S(p|p)$ is convex for all $p \in P$.

A candidate contract menu

Given observation 2, the efficient scheduling rule is characterized by single threshold, $\bar{p}$, such that, i is scheduled if and only if $p_i \geq \bar{p}$.

Consider the (obviously) convex scoring rule, S, such that:

- $S(p_i|x_i = 0) = \beta$ and $S(p_i|x_i = 1) = \alpha$ if $p_i \geq \bar{p}$, and
- $S(p_i|x_i = 0) = S(p_i|x_i = 1) = 0$ if $p_i \leq \bar{p}$.

where (α, β) are any real values such that $\alpha\bar{p} + \beta(1 - \bar{p}) = 0$ and $\alpha \geq 0$.

Result #1

Proposition 1

The scoring rule S and scheduling rule $\bar{p}$ achieves the efficient schedule in dominant strategies.

This means that "buying the producers information" is essentially costless and given we've assumed the producers have zero marginal cost, the following is easily achievable:

Corollary 1

The scoring rule S and scheduling rule $\bar{p}$ can be achieved without a budget deficit.

Forecast acquisition?

Now consider adding a stage before types are reported where producers can privately invest in a forecast update. Can we use the flexibility in S to incentivize private investment when those forecasts are socially valuable?

Think of this update as a costly experiment.

To keep things simple, suppose there is only one experiment available, characterized by the pair (ϵ, c), where for all $p \in [\epsilon, 1-\epsilon]$ the producer can chose to forgo $c > 0$ and receive a stochastic update to a new type $\hat{p}$ such that:

$$\hat{p} = \begin{cases} p + \epsilon & \text{w.p.} \quad 0.5 \\ p - \epsilon & \text{w.p.} \quad 0.5 \end{cases}$$

The candidate menu doesn't incentivise efficient acquisition ...

Fixing the draws of other producers it's clear that for some (ϵ, c) our candidate does not incentivise the efficient acquisition of forecasts!

For example, consider a producer with an endowed type of $p + \epsilon$. Whether they invest or not they will still be scheduled and investing in the update will has no net effect on their prior valuation of the contract (Bayes plausibility!). As such, there can be no positive cost worth paying.

But the scheduler benefits from re-optimising who they schedule and which target they set (observation 2: the value function is convex!). As such, for sufficiently low costs, they would like the producer to invest.

but there are other candidates!

But we should not forget the first theorem: we can implement any scoring rule that is convex. So we have a lot of freedom.

This problem resembles that of public good provision where the collectives information determines the resulting allocation, and so we should be thinking of a VCG mechanism - Vickrey (1961), Clarke (1971), and Groves (1973).

The wrinkle is here we have what is essentially type investment, a *moral hazard* problem, before types are reported. This happens to be precisely the study of Bergemann and Välimäki (2002) who confirmed the following:

Theorem 2 (Bergemann and Välimäki (2002))

With independent private values, every local social optimum can be achieved by the VCG mechanism.

Result #2

Proposition 2

There exists a scoring rule S^+ and scheduling rule $\bar{p}$ achieves the efficient schedule in dominant strategies for any forecast product (ϵ, c).

Summary

Returning to the electricity grid, the core observation is that appropriately chosen, state-contingent contracts can be used to incentivize forecast reporting, as well forecast acquisition, without the sacrifice of efficient scheduling.

This seems particularly relevant for encouraging and accommodating renewable generators who suffer acutely from stochastic supply.

We're unsure right now what about the budget implications of the second proposition and whether there is a *simple* contract menu that implements the efficient schedule in dominant strategies akin to the first proposition.

But we're confident in the generality of the arguments concerning the toy model, in particular,

- the production distributions or private information, and
- the menu of forecasts or experiments available to the producers.

References

Bergemann, Dirk and Juuso Välimäki. 2002. "Information Acquisition and Efficient Mechanism Design." *Econometrica* 70:1007–1033.

Clarke, Edward. 1971. "Multipart pricing of public goods." *Public Choice* 11(1):17–33.

Gneiting, Tilmann and Adrian E. Raftery. 2007. "Strictly Proper Scoring Rules, Prediction, and Estimation." *Journal of the American Statistical Association* 102(477):359–378.

Groves, Theodore. 1973. "Incentives in Teams." *Econometrica* 41(4):617–631.

McCarthy, John. 1956. "Measures of the Value of Information." *Proceedings of the National Academy of Sciences* 42(9):654–655.

Savage, Leonard J. 1971. "Elicitation of Personal Probabilities and Expectations." *Journal of the American Statistical Association* 66:783–801.

Vickrey, William. 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *The Journal of Finance* 16(1):8–37.