The Cassandra and the Commander: Strategic Divergence in Decision-Dependent Forecasting

Eugene Shcherbinin[1,2], Eugene Brotons Batista[2], and Gemini 3 Pro[3]

[1]London School of Economics (LSE)
[2]Bloomsbury Technology , {eugene.shcherbinin, eugenia.brotons}@bloomsburytech.com
[3]AI Co-author

January 28, 2026

Abstract

We refine the distinction between passive (GJP-style) and active (Strategic) forecasting by introducing a parameter for objective task difficulty, θ. We demonstrate that while passive observers minimize error by adhering to base rates, strategic actors maximize utility by diverging from these rates—specifically when the marginal elasticity of effort is high. We formally distinguish between *informational calibration* (accuracy) and *instrumental signaling* (impact), showing that "rational overconfidence" is not a cognitive bias but a solution to the coordination problem in probabilistic environments.

1 Introduction

Critiques of standard forecasting models often conflate the *observer's* goal (minimizing Brier scores) with the *actor's* goal (maximizing expected value). A key friction arises in high-stakes scenarios (e.g., geopolitical conflict) where the probability of success is endogenous to the actor's commitment. This paper addresses the "Quagmire Trap" critique: how do we distinguish rational strategic exertion from irrational refusal to accept inevitable defeat?

2 Literature Review

The tension between calibration and influence is well-documented across operations research and game theory.

Endogenous Uncertainty: Unlike classical stochastic programming where uncertainty is exogenous, recent work in robust optimization models "decision-dependent uncertainty," where the probability distribution of outcomes is a function of the agent's actions (Hellemo et al., 2023; Vayanos et al., 2021). This literature supports our core premise that the "true" probability π is not a static property of the world but a dynamic surface manipulated by effort.

Strategic Overconfidence: Game theoretic models have long established that rational Bayesian agents may exhibit overconfidence. Mobius et al. (2011) demonstrate that agents manage self-confidence to overcome risk aversion or motivation deficits. Similarly, Foerster et al. (2023) show in multi-agent reinforcement learning that agents must "shape" the learning dynamics of their opponents, a process that often requires non-equilibrium behavior akin to strategic over-commitment.

Forecasting in Conflict: Green and Armstrong (2002) empirically differentiate between "unaided judgment" (passive forecasting) and "role-playing" (simulation of active decision-making). They find that passive forecasting fails in conflict situations because it ignores the reflexive nature of the adversary's response to the forecaster's own signals.

3 The Enhanced Model

Let the probability of a successful outcome $E = 1$ depend on two variables:

1. $a \in [0, 1]$: The aggregate effort/resources committed by the actor.
2. $\theta \in [0, \infty)$: The objective difficulty of the task (where high θ implies high difficulty).

The probability production function is defined as:

$$P(E = 1|a, \theta) = \pi(a, \theta) \tag{1}$$

We assume decreasing returns to effort ($\frac{\partial^2 \pi}{\partial a^2} < 0$) and that difficulty negatively impacts success ($\frac{\partial \pi}{\partial \theta} < 0$).

3.1 The Observer's Problem (The Cassandra)

The passive observer (GJP) observes the difficulty θ and the actor's current state. They seek to minimize prediction error:

$$\min_f (f - \pi(a_{observed}, \theta))^2 \tag{2}$$

The observer correctly identifies "Quagmires"—situations where θ is so high that $\pi(1, \theta) \approx 0$.

3.2 The Actor's Problem (The Commander)

The actor does not just choose physical effort a_{self}, but broadcasts a *signal* $\sigma \in [0, 1]$ (projected confidence) to motivate subordinates or allies, whose effort a_{allies} depends on this signal. Total effort $a = a_{self} + a_{allies}(\sigma)$.

The actor maximizes:

$$\max_{\sigma, a_{self}} [\pi(a(\sigma), \theta) \cdot V - C(a_{self}, \sigma)] \tag{3}$$

Where $C(a, \sigma)$ captures the reputational cost of being wrong ("loss of face").

4 Strategic Implications

Proposition 1 (The Signal-Belief Wedge)**.** *In equilibrium, the optimal signal σ^* (public forecast) will strictly exceed the actor's private conditional probability estimate $\pi(a, \theta)$ whenever the participation of allies is elastic to confidence.*

Proof. Assume allies only contribute effort if they believe success is plausible ($\sigma > \sigma_{threshold}$). The actor calculates that acknowledging the "true" low probability (say, 20%) results in $a_{allies} \to 0$, causing collapse ($\pi \to 0$). By signaling $\sigma = 1$ ("We will win"), the actor secures $a_{allies} > 0$. Even if the resulting success probability is only 40%, the expected value calculation ($0.4V - C$) may exceed the certainty of defeat (0). Thus, $\sigma^* \gg P(Success)$. The actor is not "deluded"; they are *instrumentally calibrated.* □

Proposition 2 (The Quagmire Condition)**.** *Rational overreaction is bounded by the sensitivity of the outcome to effort, denoted by $\epsilon_{a,\pi} = \frac{\partial \pi}{\partial a}$.*

Proof. The critique suggests that actors might rationally pursue impossible goals. However, the first-order condition requires:

$$\frac{\partial \pi}{\partial a} \cdot \frac{\partial a}{\partial \sigma} \cdot V = \frac{\partial C}{\partial \sigma} \tag{4}$$

- **Case A (Pivotal Moment):** The outcome is highly sensitive to effort ($\frac{\partial \pi}{\partial a}$ is large). High signaling is rational.
- **Case B (The Quagmire):** The difficulty θ is sufficiently high such that $\frac{\partial \pi}{\partial a} \to 0$. In this limit, further signaling is irrational (pure cost).

The "Superforecaster" advantage lies in accurately estimating θ (Quagmire detection), whereas the "Strategic Actor" advantage lies in maximizing $\frac{\partial a}{\partial \sigma}$ (Mobilization). □

5 Conclusion

Forecasting accuracy and strategic efficacy are often orthogonal goals. We conclude that "overconfidence" is a rational feature of the executive function, provided it is calibrated to the *elasticity of the outcome*, not merely the static probability.

References

Jakob Foerster, Alon Pitan, Matei Mateescu, and Illia Polosukhin. Scaling opponent shaping to high dimensional games. *arXiv preprint arXiv:2312.12568*, 2023.

Kesten C Green and J Scott Armstrong. Forecasting decisions in conflict situations: a comparison of game theory, role-playing, and unaided judgement. *International Journal of Forecasting*, 18(3):321–344, 2002.

Lars Hellemo, Anders Hügnell, and Asgeir Tomasgard. Solving decision problems with endogenous uncertainty and conditional information revelation using influence diagrams. *arXiv preprint arXiv:2402.15486*, 2023.

Markus M Mobius, Muriel Niederle, Paul Niehaus, and Tanya S Rosenblat. Managing self-confidence: Theory and experimental evidence. Working Paper 17014, National Bureau of Economic Research, 2011.

Phebe Vayanos, Angelos Georghiou, and Han Yu. Robust optimization with decision-dependent information discovery. *Management Science*, 67(3):1363–1381, 2021.