When Central Bank Vagueness Reveals Low Confidence

Lakshya Narula*

Abstract

Vagueness can reveal low confidence, depending on equilibrium sorting. A central bank privately knows the reliability of its assessment and commits, before observing its direction, to an unobserved disclosure probability. I characterize all pure-strategy perfect Bayesian equilibria. In the monotone separating equilibrium, a directional announcement raises perceived reliability and a non-directional announcement lowers it. Persistence can sustain this sorting although the low-reliability bank prefers greater precision in a one-shot game. A closed-form example produces separation for discount factors at least 1/3, alongside precise pooling. All messages occur under both protocols, so Bayes' rule determines every belief.

Keywords: central bank communication; reputation; confidence; disclosure; signaling
JEL classification: D82; E52; E58

1 Introduction

Non-directional central bank communication can reveal low confidence rather than conceal it. If a low-reliability bank adopts a protocol that produces non-directional statements more often, such a statement lowers the market's estimate of reliability and the weight placed on the bank's next forecast. A forward-looking bank may accept a less accurate response today to obtain a better-calibrated response tomorrow.

I isolate this intertemporal force in a two-period signaling game with a common objective. The bank privately knows whether its directional assessment has low or high reliability. Before observing the first direction, it commits to a hidden protocol that discloses the direction with either a low or high probability. In the monotone separating equilibrium, the low type selects the coarse protocol and the high type the fine protocol. A directional announcement then raises perceived reliability, while a non-directional announcement lowers it. Persistence can sustain this sorting even when the low type would choose the fine protocol in a one-shot game.

The analysis characterizes all pure-strategy equilibria, including pooling and reverse sorting. In a closed-form example, precise pooling is the sole pure-strategy equilibrium for discount factors below 1/3. At and above 1/3, monotone separation appears alongside pooling; reverse sorting remains impossible. When the discount factor exceeds 31/35, separation also lowers ex ante common loss relative to pooling. Every realized message remains on path, so the result uses neither off-path beliefs nor a direct preference for prestige.

The closest paper is Gáti and Handlan (2026), in which private confidence and its public reputation shape static communication precision. Their central bank deliberately does not internalize how an announcement changes future reputation. Rather than solve the full dynamic version of their continuous-message model, this note gives a closed-form benchmark in which Bayesian inference from

*Department of Economics, University of Minnesota, 4-101 Hanson Hall, 1925 Fourth Street South, Minneapolis, MN 55455, USA. Corresponding author: narul014@umn.edu.

a current message creates an intertemporal incentive to sort over protocols. Dynamic communication can also reduce optimal precision because current disclosure constrains future persuasion (Gáti, 2023); here, instead, current messages signal persistent private reliability and thereby change the market's response to an independent future assessment. The underlying structure relates to privately informed senders choosing information structures (Hedlund, 2017) and to voluntary disclosure when information quality or information endowment is uncertain (Jung and Kwon, 1988; Kim and Pae, 2025). It differs from models in which competence changes credibility (Moscarini, 2007; Frankel and Kartik, 2018), experts value reputation directly (Ottaviani and Sørensen, 2006), central-bank actions build reputations for policy preferences (Hansen and McMahon, 2016; Amador and Phelan, 2026), or reputation disciplines communication (Camous and Matveev, 2025).

2 Reporting game

There are two periods and a common quadratic objective. Nature draws persistent reliability $h \in \{h_L, h_H\}$, where $0 < h_L < h_H < 1$ and $\Pr(h = h_H) = \mu \in (0,1)$, and only the bank observes it. Directions $\theta_t \in \{-1, 1\}$ are independent across periods and independent of h, with equal probabilities. The underlying state $y_t \in \{-1, 1\}$ satisfies

$$\Pr(y_t = \theta_t \mid h, \theta_t) = \frac{1+h}{2}.$$

Thus $E[y_t \mid h, \theta_t] = h\theta_t$. If the market chooses x_t, expected squared error is $(x_t - h\theta_t)^2 + 1 - h^2$. The final term is independent of communication, so the relevant loss is $(x_t - h\theta_t)^2$.

The first-period timing is as follows. After observing h but before observing θ_1, the bank irreversibly selects a hidden protocol $q \in \{c, f\}$, with $0 < c < f < 1$. An independent random device releases $m_1 = \theta_1$ with probability q and otherwise releases $m_1 = 0$; the bank cannot override its realization. The public observes m_1 and chooses x_1. The message $m_1 = 0$ represents a scheduled, publicly observed non-directional statement—for example, "the outlook remains uncertain"—rather than an unobserved failure to communicate. The fine protocol Blackwell-dominates the coarse protocol: erasing a directional message under f with probability $1 - c/f$ produces c.

Neither y_1 nor another signal of h becomes public before period two. A fresh direction θ_2 is then drawn and disclosed with certainty; the bank makes no second protocol choice. The market observes θ_2 and chooses x_2. Commitment to the first-period randomization is a primitive, which can represent an automated release rule or delegation to staff under a pre-approved procedure.

For a candidate profile (q_L, q_H), let p_1 and p_0 be the posterior probabilities of high reliability after a directional and a non-directional message. Bayes' rule gives

$$p_1 = \frac{\mu q_H}{\mu q_H + (1-\mu) q_L}, \qquad p_0 = \frac{\mu(1-q_H)}{\mu(1-q_H) + (1-\mu)(1-q_L)}. \tag{1}$$

Let $r_k = h_L + (h_H - h_L)p_k$ denote posterior mean reliability. By symmetry, either directional message induces p_1. The market chooses $x_1 = r_1 m_1$ after a directional message and $x_1 = 0$ after the non-directional message.

In period two, the market chooses $x_2 = r_k \theta_2$ after first-period message category $k \in \{0, 1\}$. The bank discounts this loss by $\delta \in [0, 1]$ and receives no direct utility from a high reputation. Holding the market's equilibrium belief mapping fixed, type j's fine-minus-coarse loss is $(f - c)\Psi_j$, where

$$\Psi_j = (r_1 - h_j)^2 - h_j^2 + \delta\big[(r_1 - h_j)^2 - (r_0 - h_j)^2\big]. \tag{2}$$

A type prescribed f requires $\Psi_j \le 0$; a type prescribed c requires $\Psi_j \ge 0$.

3 Equilibria

Write $\bar{h} = \mu h_H + (1-\mu)h_L$. Superscripts M and R denote posterior means from monotone separation (c, f) and reverse separation (f, c), respectively.

Proposition 1 (Pure-strategy equilibria)**.** *For every $\delta \in [0,1]$, the pure-strategy perfect Bayesian equilibria are characterized as follows.*

1. *Pooling on c is never an equilibrium.*
2. *Pooling on f is an equilibrium if and only if $\bar{h} \leq 2h_L$.*
3. *Monotone separation, $(q_L, q_H) = (c, f)$, is an equilibrium if and only if $\Psi_L^M \geq 0$. The high type's incentive constraint always holds. If $r_1^M < 2h_L$, this condition is equivalent to*

$$\delta \geq \delta_M \equiv \frac{h_L^2 - (r_1^M - h_L)^2}{(r_1^M - h_L)^2 - (r_0^M - h_L)^2}. \tag{3}$$

4. *Reverse separation, $(q_L, q_H) = (f, c)$, is an equilibrium if and only if $\Psi_L^R \leq 0 \leq \Psi_H^R$.*

An increase in δ relaxes both incentive constraints for monotone separation: it makes c more attractive to the low type and f more attractive to the high type.

Proof. Under pooling, all messages induce $\bar{h}$. Continuation loss is constant in the protocol, and type j's fine-minus-coarse current loss, divided by $f - c$, is $(\bar{h} - h_j)^2 - h_j^2$. This expression is negative for the high type, so pooling on c is impossible. It is nonpositive for the low type exactly when $\bar{h} \leq 2h_L$, which proves the pooling claims.

Under monotone separation, Bayes' rule implies $h_L < r_0^M < \bar{h} < r_1^M < h_H$. For the high type, both terms in Ψ_H^M are negative. For the low type, the coefficient on δ is positive. Its prescribed choice c is optimal exactly when $\Psi_L^M \geq 0$. At $\delta = 0$, this requires $r_1^M \geq 2h_L$; when that inequality fails, rearrangement gives (3). Under reverse separation, $h_L < r_1^R < \bar{h} < r_0^R < h_H$. The prescribed choices require, respectively, $\Psi_L^R \leq 0$ and $\Psi_H^R \geq 0$. These four profiles exhaust the pure strategies. Every message has positive probability under either protocol, so Bayes' rule determines all beliefs used in the argument. □

Corollary 1 (Closed-form example)**.** *Let*

$$(h_L, h_H, \mu, c, f) = \left(\frac{2}{5}, \frac{4}{5}, \frac{7}{10}, \frac{3}{10}, \frac{9}{10}\right).$$

For $0 \leq \delta < 1/3$, the only pure-strategy equilibrium is pooling on f. For $1/3 \leq \delta \leq 1$, pooling on f and monotone separation are the only pure-strategy equilibria. When separation exists, it has lower ex ante relevant loss than pooling if and only if $\delta > 31/35$.

Proof. Here $\bar{h} = 17/25$. Under monotone separation, $(r_0^M, r_1^M) = (1/2, 3/4)$ and

$$\Psi_L^M = -\frac{3}{80} + \delta\frac{9}{80},$$

which is nonnegative exactly when $\delta \geq 1/3$. Under reverse separation, $(r_0^R, r_1^R) = (101/130, 23/40)$. The high type's current term is $-943/1600$, while its continuation term is less than $81/1600$. Hence $\Psi_H^R < 0$ for every $\delta \in [0,1]$, so the high type deviates to f. Finally, $\bar{h} < 2h_L$, so pooling on f exists, while Proposition 1 rules out pooling on c.

For $e \in \{\text{sep}, \text{pool}\}$, define ex ante relevant loss by

$$L^e \equiv \mathbb{E}\big[(x_1 - h\theta_1)^2 + \delta(x_2 - h\theta_2)^2 \,\big|\, e\big].$$

Under pooling, current and continuation relevant losses are 499/6250 and 21/625. Under separation, they are 91/1000 and 21/1000. Therefore

$$L^{\text{sep}} - L^{\text{pool}} = \frac{279 - 315\delta}{25000},$$

which is negative exactly when $\delta > 31/35$. □

The example isolates the dynamic result. Below the threshold, the low type's current gain from disclosure prevents informative sorting. Above it, a non-directional message moves future reputation toward h_L enough to offset that gain. Precise pooling survives because messages do not change reputation under pooling. Dynamics creates an informative equilibrium; it does not select it uniquely.

The direction of sorting is parameter-dependent. Under reverse separation, $r_1^R < r_0^R$, so a larger δ lowers Ψ_L^R and raises Ψ_H^R. It therefore relaxes both reverse-sorting constraints as well. Conditional on monotone separation, the empirical prediction is that a non-directional statement lowers the market's sensitivity to the bank's next forecast; reverse separation reverses that prediction.

Two compact extensions show that the result is not tied to the binary menu or perfect non-verification. If q can take any value in $[c, f]$, expected loss is affine in q for a fixed equilibrium belief mapping. Hence every strict endpoint equilibrium characterized above remains an equilibrium of the continuous-menu game, and every strict pure equilibrium uses an endpoint. If, independently of type, protocol, and message, an intervening signal reveals h perfectly with probability s, the continuation term in (2) is multiplied by $1 - s$. The mechanism survives whenever the effective discount factor $\delta(1-s)$ crosses the relevant threshold. At strict inequalities, continuity also preserves the equilibrium in an open neighborhood of the parameters.

4 Conclusion

Persistent reliability can make the form of a central bank's announcement informative about confidence. In the monotone equilibrium, a low-reliability bank chooses vagueness to correct an inflated future reputation even at a current cost. The mechanism requires neither lying nor a taste for prestige, but it does require commitment to the disclosure protocol and sufficiently limited interim verification. Multiplicity remains: the model identifies when low-confidence vagueness can arise, not a universal direction of sorting.

References

Amador, M., Phelan, C., 2026. Central bank reputation with noise. Journal of Monetary Economics 158, 103906. doi:10.1016/j.jmoneco.2026.103906.

Camous, A., Matveev, D., 2025. Monetary Communication and Credibility in a Multi-Sector Economy. Working Paper Series 1013. Banque de France.

Frankel, A., Kartik, N., 2018. What kind of central bank competence? Theoretical Economics 13, 697–727. doi:10.3982/TE2290.

Gáti, L., 2023. Talking over time—dynamic central bank communication. Journal of Money, Credit and Banking 55, 1147–1176. doi:10.1111/jmcb.12934.

Gáti, L., Handlan, A., 2026. Reputation for Confidence. Working Paper Series 3141. European Central Bank. doi:10.2866/8342498.

Hansen, S., McMahon, M.F., 2016. First impressions matter: Signalling as a source of policy dynamics. Review of Economic Studies 83, 1645–1672. doi:10.1093/restud/rdw007.

Hedlund, J., 2017. Bayesian persuasion by a privately informed sender. Journal of Economic Theory 167, 229–268. doi:10.1016/j.jet.2016.11.003.

Jung, W.O., Kwon, Y.K., 1988. Disclosure when the market is unsure of information endowment of managers. Journal of Accounting Research 26, 146–153. doi:10.2307/2491118.

Kim, E., Pae, S., 2025. Voluntary disclosure when information quality is unknown. The Accounting Review 100, 269–297. doi:10.2308/TAR-2023-0540.

Moscarini, G., 2007. Competence implies credibility. American Economic Review 97, 37–63. doi:10.1257/aer.97.1.37.

Ottaviani, M., Sørensen, P.N., 2006. Reputational cheap talk. RAND Journal of Economics 37, 155–175. doi:10.1111/j.1756-2171.2006.tb00010.x.