

A Belief-Variance Threshold for Value Co-destruction in Multi-Level Generative AI Adoption

Francesca Aquilino

Decisions Lab, DiGiES
Università degli Studi Mediterranea di Reggio Calabria
Via dell'Università 25, 89124 Reggio Calabria, Italy

Massimiliano Ferrara

Decisions Lab, DiGiES
Università degli Studi Mediterranea di Reggio Calabria
Via dell'Università 25, 89124 Reggio Calabria, Italy

This article is distributed under the Creative Commons by-nc-nd Attribution License.
Copyright © 2026 Hikari Ltd.

Abstract

We study the intra-organizational adoption of generative artificial intelligence (GenAI) as a multi-level Bayesian game played by senior management, departmental units and individual employees. Existing treatments characterize value co-creation through three parameter thresholds (complementarity, investment efficiency, displacement management) but treat the role of information asymmetry only informally, asserting that high belief variance pushes the equilibrium toward co-destruction without deriving the mechanism. The present note fills this gap. Working with an aggregative reduction of the departmental stage game, we prove that expected organizational welfare is strictly decreasing and concave in the variance of beliefs about rivals' adoption, and we exhibit a closed-form critical variance beyond which equilibrium welfare drops below the no-coordination benchmark. This makes information transparency a fourth, analytically explicit lever alongside the three known parameter thresholds. A Monte Carlo experiment with 10,000 replications confirms the monotone decay and locates the empirical critical variance in close agreement with the theoretical value.

Mathematics Subject Classification: 91A10, 91A27, 91B06, 90C15

Keywords: Bayesian games, aggregative games, generative AI adoption, value co-destruction, comparative statics, information asymmetry

1 Introduction

The diffusion of generative artificial intelligence inside firms has proved far more uneven than the raw capabilities of the underlying models would suggest. Identical tools deployed for identical use cases produce transformative gains in one organization and quiet failure in another, and the decisive variable is repeatedly found to be the architecture of strategic interaction surrounding the technology rather than the technology itself [3]. A productive way of formalizing this observation is to model adoption not as a single decision but as a nested game in which senior management sets investment and governance, departmental units choose adoption intensity, and employees choose engagement effort, each acting on beliefs about the others under incomplete information [3], [5].

Within this program, the conditions separating *value co-creation* from *value co-destruction* [4], [6] have been expressed through three parameter thresholds: a strategic complementarity threshold across departments, an investment efficiency threshold at the management level, and a displacement-management threshold at the employee level. All three must hold simultaneously for the coordinated regime to obtain. What the existing analysis leaves implicit is the role of *information*: it is asserted that, when the variance of beliefs about rivals' strategies grows large enough, defensive behaviour emerges and aggregate welfare falls below the cooperative benchmark, but the claim is stated as a qualitative failure mode rather than derived as a comparative-statics result.

The contribution of this note is to make that mechanism precise. We isolate the departmental stage game, reduce it to an aggregative game in the sense of Corchón [1] and Cornes and Hartley [2], and study how the equilibrium and the realized welfare respond to the belief variance σ^2 . We prove that expected welfare is strictly decreasing and concave in σ^2 (Proposition 3.2), and we obtain a closed-form critical variance σ_c^2 at which welfare crosses the no-coordination benchmark. Transparency of internal information flows thereby becomes a fourth lever, on the same analytical footing as the three parameter thresholds, and not merely a heuristic caveat. A Monte Carlo experiment validates the prediction.

2 The departmental stage game

We work with the departmental layer of the multi-level model, taking the management investment $I > 0$ and the average employee engagement $\bar{e} \in (0, 1]$ as fixed at their equilibrium values; this is the standard reduction used to obtain tractable comparative statics in the adoption layer. There are $n \geq 2$ symmetric departments. Department i chooses an adoption intensity $d_i \in [0, 1]$ and receives the payoff

$$U_i(d_i, d_{-i}) = (a\bar{e})d_i - \frac{c}{2}d_i^2 + \delta d_i \sum_{k \neq i} d_k - \frac{\gamma}{2} (d_i - \mathbb{E}_i[\bar{d}_{-i}])^2, \quad (1)$$

where $a > 0$ is the productivity coefficient, $c > 0$ the marginal cost of adoption, $\delta \in [0, \delta_{\max}]$ the cross-departmental spillover, and $\gamma > 0$ the coordination (alignment) intensity. The term $\bar{d}_{-i} = \frac{1}{n-1} \sum_{k \neq i} d_k$ is the average rival adoption, and $\mathbb{E}_i[\cdot]$ is department i 's belief operator.

The information structure is the one inherited from the parent model. Each department holds a common-prior belief about the rivals' average adoption; the belief is unbiased but noisy,

$$\mathbb{E}_i[\bar{d}_{-i}] = \bar{d}_{-i}^* + \varepsilon_i, \quad \mathbb{E}[\varepsilon_i] = 0, \quad \text{Var}(\varepsilon_i) = \sigma^2, \quad (2)$$

where \bar{d}_{-i}^* is the true equilibrium average and $\sigma^2 \geq 0$ measures the opacity of internal information flows. The case $\sigma^2 = 0$ is the full-information (transparent) benchmark.

Assumption 2.1 $c > (n-1)\delta$ and $a\bar{e} \leq c + (n-1)\delta$. The first inequality is the contraction (diagonal-dominance) condition that keeps own-cost effects above aggregate spillovers; the second keeps the unconstrained optimizer inside $[0, 1]$.

3 Equilibrium and welfare

We first record the equilibrium of the perturbed game and then study how realized welfare depends on σ^2 .

Lemma 3.1 *Under Assumption 2.1, the best response of department i is the affine map*

$$d_i(d_{-i}) = \frac{a\bar{e} + \delta \sum_{k \neq i} d_k + \gamma \mathbb{E}_i[\bar{d}_{-i}]}{c + \gamma}. \quad (3)$$

The symmetric Bayesian Nash equilibrium of the noiseless game ($\sigma^2 = 0$) is

$$d^* = \frac{a\bar{e}}{c - (n-1)\delta}, \quad (4)$$

and it is the unique equilibrium.

Proof. Differentiating (1) in d_i gives the first-order condition $a\bar{e} - cd_i + \delta \sum_{k \neq i} d_k - \gamma(d_i - \mathbb{E}_i[\bar{d}_{-i}]) = 0$, which rearranges to (3); the objective is strictly concave in d_i since $-(c + \gamma) < 0$, so the stationary point is the global maximizer. At $\sigma^2 = 0$ beliefs are correct, $\mathbb{E}_i[\bar{d}_{-i}] = \bar{d}_{-i}$, and the alignment term vanishes at a symmetric profile. Imposing $d_i = d_k = d^*$ in (3) yields $d^*(c + \gamma) = a\bar{e} + (n - 1)\delta d^* + \gamma d^*$, i.e. $d^*(c - (n - 1)\delta) = a\bar{e}$, which is (4). Uniqueness follows from Assumption 2.1: the best-response map (3) is a contraction in the sup-norm with modulus $(n - 1)\delta/(c + \gamma) < 1$, so Banach's fixed-point theorem gives a unique fixed point. ■

Realized organizational welfare is the sum of departmental payoffs evaluated at the played profile. We compare it with the *no-coordination benchmark* W^0 , the welfare obtained when departments ignore one another (δ and the alignment term are switched off), which is the analogue of the autarky benchmark of the parent framework.

Proposition 3.2 (Belief-variance welfare decay) *Let $W(\sigma^2)$ denote expected realized welfare, i.e. $W(\sigma^2) = \mathbb{E}[\sum_{i=1}^n U_i]$, when beliefs satisfy (2) with variance σ^2 , the realized profile being the noisy best response (3). Under Assumption 2.1,*

$$W(\sigma^2) = W(0) - \frac{n\gamma c}{2(c + \gamma)} \sigma^2. \quad (5)$$

Consequently W is strictly decreasing and (weakly) concave in σ^2 , and there exists a unique critical variance

$$\sigma_c^2 = \frac{2(c + \gamma)}{n\gamma c} (W(0) - W^0) \quad (6)$$

such that $W(\sigma^2) > W^0$ for $\sigma^2 < \sigma_c^2$ and $W(\sigma^2) < W^0$ for $\sigma^2 > \sigma_c^2$.

Proof. Write the belief noise as ε_i with mean 0 and variance σ^2 , independent across i . By (3) the played action is $d_i = d^* + \kappa \varepsilon_i$ with $\kappa = \gamma/(c + \gamma)$, where d^* is the noiseless equilibrium of Lemma 3.1: the productivity, cost and spillover terms are unchanged in expectation by the zero-mean perturbation, while the belief enters (3) with weight κ .

Expand U_i around d^* . The linear-in- ε terms vanish in expectation because $\mathbb{E}[\varepsilon_i] = 0$ and, by symmetry, $\mathbb{E}[\varepsilon_i \sum_{k \neq i} \varepsilon_k] = 0$ under independence. The only surviving second-order contributions are the own quadratic cost $-\frac{c}{2} \mathbb{E}[(\kappa \varepsilon_i)^2] = -\frac{c}{2} \kappa^2 \sigma^2$ and the alignment penalty. For the latter, the realized misalignment is $d_i - \mathbb{E}_i[\bar{d}_{-i}] = (d^* + \kappa \varepsilon_i) - (d^* + \varepsilon_i) = -(1 - \kappa)\varepsilon_i$, so $-\frac{\gamma}{2} \mathbb{E}[(d_i - \mathbb{E}_i[\bar{d}_{-i}])^2] = -\frac{\gamma}{2}(1 - \kappa)^2 \sigma^2$. Summing the two surviving terms over the n departments,

$$W(\sigma^2) - W(0) = -\frac{n}{2} (c\kappa^2 + \gamma(1 - \kappa)^2) \sigma^2.$$

Substituting $\kappa = \gamma/(c + \gamma)$ and $1 - \kappa = c/(c + \gamma)$,

$$c\kappa^2 + \gamma(1 - \kappa)^2 = \frac{c\gamma^2 + \gamma c^2}{(c + \gamma)^2} = \frac{\gamma c(\gamma + c)}{(c + \gamma)^2} = \frac{\gamma c}{c + \gamma}.$$

Hence $W(\sigma^2) - W(0) = -\frac{n}{2} \cdot \frac{\gamma c}{c + \gamma} \sigma^2$, which is exactly (5). (The spillover term contributes nothing: $\mathbb{E}[\varepsilon_i \varepsilon_k] = 0$ for $i \neq k$, so the cross-product $\delta \sum_{i \neq k} \mathbb{E}[d_i d_k]$ is unchanged from its noiseless value.) The map $\sigma^2 \mapsto W(\sigma^2)$ is affine with strictly negative slope, hence strictly decreasing and concave. Setting $W(\sigma^2) = W^0$ and solving for σ^2 produces the unique crossing point (6), with the stated sign of $W - W^0$ on either side because the slope is negative. ■

Remark 3.3 *Proposition 3.2 promotes information transparency to a fourth equilibrium lever. The three known thresholds determine the gap $W(0) - W^0$, i.e. the welfare headroom created by coordination; the belief variance then erodes that headroom at the constant rate $n\gamma c/2(c + \gamma)$. A firm may clear all three parameter thresholds and still slide into co-destruction if its internal information flows are opaque enough that σ^2 exceeds σ_c^2 . This is the precise sense in which transparency and parameter quality are complementary.*

Corollary 3.4 *The marginal welfare loss $|\partial W/\partial \sigma^2| = \frac{n}{2} \frac{\gamma c}{c + \gamma}$ is strictly increasing in the coordination intensity γ and in the marginal cost c , and is bounded above by $\frac{n}{2} \min\{\gamma, c\}$. For fixed headroom $W(0) - W^0$, the critical variance σ_c^2 is therefore strictly decreasing in γ : more tightly coordinated organizations are more exposed to belief opacity, because the alignment penalty that drives coordination also amplifies the cost of acting on wrong beliefs.*

Proof. The loss coefficient $g(\gamma) = \gamma c/(c + \gamma)$ has derivative $g'(\gamma) = c^2/(c + \gamma)^2 > 0$, so it is strictly increasing in γ ; by symmetry of the expression it is strictly increasing in c as well. Since $\gamma c/(c + \gamma) \leq \min\{\gamma, c\}$, the stated bound follows. By (6), σ_c^2 is proportional to $1/g(\gamma)$ at fixed headroom, hence strictly decreasing in γ . ■

4 Monte Carlo validation

To test Proposition 3.2 we simulate the perturbed departmental game. We fix a representative configuration drawn from the calibration ranges of the parent framework: $n = 5$ departments, productivity $a = 0.6$, average engagement $\bar{e} = 0.6$, marginal cost $c = 1.0$, spillover $\delta = 0.10$, coordination intensity $\gamma = 0.5$. With these values Assumption 2.1 holds ($c = 1.0 > (n - 1)\delta = 0.40$ and $a\bar{e} = 0.36 \leq c + (n - 1)\delta = 1.40$), the noiseless equilibrium is interior, $d^* = a\bar{e}/(c - (n - 1)\delta) = 0.36/0.60 = 0.60$, and the theoretical loss coefficient

in (5) is $n\gamma c/2(c + \gamma) = 5 \cdot 0.5 \cdot 1/(2 \cdot 1.5) \approx 0.833$. The noiseless coordinated welfare is $W(0) = 0.900$ and the no-coordination benchmark is $W^0 = 0.324$, giving a headroom of 0.576 and a closed-form critical variance $\sigma_c^2 = 2(c + \gamma)(W(0) - W^0)/(n\gamma c) = 2 \cdot 1.5 \cdot 0.576/2.5 \approx 0.691$.

For each value of σ^2 on a grid in $[0, 0.85]$ we draw 10,000 replications of the belief noise vector $\varepsilon \sim \mathcal{N}(0, \sigma^2 I_n)$, compute the played profile via (3), evaluate realized welfare $\sum_i U_i$, and average across replications. The no-coordination benchmark W^0 is computed by switching off δ and the alignment term. Table 1 reports the results.

Table 1: Monte Carlo welfare decay (10,000 replications per row).

σ^2	Simulated W	Theory (5)	$W - W^0$	Regime
0.00	0.900	0.900	+0.576	co-creation
0.05	0.859	0.858	+0.535	co-creation
0.10	0.817	0.817	+0.493	co-creation
0.20	0.734	0.733	+0.410	co-creation
0.40	0.570	0.567	+0.246	co-creation
0.69	0.326	0.324	+0.002	boundary
0.85	0.193	0.192	-0.131	co-destruction

Two findings emerge. The simulated welfare tracks the affine prediction (5) to within Monte Carlo error across the whole grid, confirming the strictly decreasing, concave (here exactly linear) decay. And the empirical sign change occurs at $\sigma^2 \approx 0.69$, coinciding with the closed-form critical variance $\sigma_c^2 \approx 0.691$ obtained above from (6). The experiment thus corroborates both the qualitative monotonicity and the quantitative location of the co-destruction boundary.

Remark 4.1 *The elasticity ordering reported in the parent framework places the complementarity-type parameters first among the welfare drivers; the present result is consistent with that ordering and refines it, in that the belief variance acts as a multiplicative discount on whatever headroom the parameters generate, rather than as an independent additive driver. Reducing σ^2 is most valuable precisely in the organizations whose parameters already place them just above the co-creation boundary, since there the headroom is small and a modest increase in opacity suffices to cross σ_c^2 .*

5 Concluding remarks

We have turned an informal failure-mode statement of the multi-level GenAI adoption framework into an explicit comparative-statics result. Expected or-

ganizational welfare falls linearly with the variance of inter-departmental beliefs, at a rate fixed by the cost and coordination parameters, and crosses the no-coordination benchmark at a closed-form critical variance. Information transparency is therefore a fourth lever, analytically on par with the complementarity, investment efficiency and displacement-management thresholds, and a Monte Carlo experiment confirms the prediction both in shape and in the location of the co-destruction boundary. Two extensions are natural: endogenizing σ^2 as the outcome of costly information acquisition, so that the critical variance becomes a design target rather than a primitive; and allowing asymmetric departments, where the scalar σ^2 is replaced by a belief covariance matrix and the welfare loss by a quadratic form in its spectrum.

Acknowledgements. The authors thank the Decisions Lab of the Mediterranean University of Reggio Calabria for support, and gratefully acknowledge the framework of Ferrara, Viglia and Romero on which the present analysis builds.

References

- [1] L.C. Corchón, Comparative statics for aggregative games: the strong concavity case, *Mathematical Social Sciences*, **28** (1994), no. 3, 151 - 165.
[https://doi.org/10.1016/0165-4896\(94\)90001-9](https://doi.org/10.1016/0165-4896(94)90001-9)
- [2] R. Cornes and R. Hartley, Fully aggregative games, *Economics Letters*, **116** (2012), no. 3, 631 - 633.
<https://doi.org/10.1016/j.econlet.2012.06.024>
- [3] M. Ferrara, G. Viglia and J.C. Romero, Strategic tensions in organizational GenAI adoption: a game theory modeling of internal resource competition, workforce dynamics, and value management, *Technological Forecasting and Social Change*, **227** (2026), 124653.
<https://doi.org/10.1016/j.techfore.2026.124653>
- [4] L. Plé and R. Chumpitaz Cáceres, Not always co-creation: introducing interactional co-destruction of value in service-dominant logic, *Journal of Services Marketing*, **24** (2010), no. 6, 430 - 437.
<https://doi.org/10.1108/08876041011072546>
- [5] M.H. Seifdar and B. Amiri, Strategic adoption of generative AI in organizations: a game-theoretic and network-based approach, *International Journal of Information Management*, **84** (2025), 102932.
<https://doi.org/10.1016/j.ijinfomgt.2025.102932>

- [6] A. Susarla, R. Gopal, J.B. Thatcher and S. Sarker, The Janus effect of generative AI: charting the path for responsible conduct of scholarly activities in information systems, *Information Systems Research*, **34** (2023), no. 2, 399 - 408. <https://doi.org/10.1287/isre.2023.ed.v34.n2>

Received: June 21, 2026; Published: July 1, 2026