Applied Mathematical Sciences, Vol. 20, 2026, no. 4, 273 - 280
HIKARI Ltd, www.m-hikari.com
https://doi.org/10.12988/ams.2026.919365

A Perception-Gap Theorem for Generative AI Adoption

Francesca Aquilino

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

Enterprise surveys of generative AI (GenAI) show that leaders underestimate employees' use of the technology, while most firms report returns at or above expectations and limited scaling. We model this as a two-stage game in which management chooses complementary investment and a continuum of employees chooses engagement under peer effects and displacement concerns, management holding a misspecified belief that scales the true engagement response by a perception factor $\mu \in (0, 1]$. The main theorem characterizes the unique equilibrium in closed form: investment is proportional to μ; below a critical perception $\mu_c = 2\theta$, with θ a composite displacement burden, the organization stays dormant; the profit loss equals $(1 - \mu)^2/(1 - 2\theta)$ of attainable profit; active organizations are always positively surprised by returns, whereas dormancy is a self-confirming equilibrium. Calibrated on McKinsey and Deloitte surveys, investment sits near one third of its efficient level.

Mathematics Subject Classification: 91A65, 91A27, 91B06, 90B70

Keywords: generative AI adoption, Stackelberg games, misspecified beliefs, self-confirming equilibrium, peer effects, organizational underinvestment

1 Introduction

At the task level, field evidence shows an average productivity gain of about 15% from a generative assistant in customer support [3]; at the organizational

level the picture differs. McKinsey reports that 92% of companies plan to increase their AI investments over the next three years, yet only 1% of leaders describe their company as mature in deployment [8]. Deloitte, surveying 2,773 leaders of AI-savvy organizations across fourteen countries between July and September 2024, finds that GenAI access still covers under 40% of the workforce and that over two thirds of respondents expect to fully scale 30% or fewer of their current experiments within six months [4].

A recent line of work explains the gap between capability and organizational value as a coordination problem among senior management, departments and employees [6]. In that multi-level Bayesian game, value co-creation [9] requires strategic complementarity, investment efficiency and effective management of displacement concerns to hold jointly; a companion note derives in closed form how the variance of departmental beliefs erodes the welfare generated by coordination [1]. Both treat beliefs as noisy but *unbiased*, whereas surveys point to a directional bias. C-suite leaders estimate that 4% of employees use GenAI for at least 30% of their daily work, while 13% of employees report doing so; only 20% of leaders see employees at that intensity now or within a year, against 47% of employees [8].

This paper studies the consequences of that downward bias. Beliefs are misspecified rather than noisy, so the solution concept combines Stackelberg leadership with a misperceived follower response, in the spirit of misspecified learning [5, 7]. The displacement and investment-efficiency thresholds of [6] collapse into a single composite burden θ that interacts with the bias, and the model yields a testable prediction: biased organizations that invest are positively surprised by returns, as Deloitte reports, while those below a critical perception never observe the evidence that would correct them.

2 The model

The organization consists of a management M and a unit mass of employees $j \in [0, 1]$; departmental adoption intensity is held at its equilibrium level and absorbed into the parameters, isolating the vertical interaction to which the McKinsey evidence speaks.

Stage 1. Management chooses complementary investment $I \geq 0$, interpreted as training, workflow integration and governance capacity, at cost $\frac{c}{2}I^2$ with $c > 0$.

Stage 2. After observing I, employees simultaneously choose engagement $e_j \in [0, 1]$. Employee j receives

$$u_j(e_j, \bar{e}) = (\rho I - \tau)\, e_j + \lambda\, e_j\, \bar{e} - \tfrac{1}{2}e_j^2, \tag{1}$$

where $\bar{e} = \int_0^1 e_k\, dk$ is average engagement, $\rho > 0$ measures how strongly investment raises the private return of engaging (skill complementarity), $\tau \geq 0$

is the net displacement concern associated with the current adoption intensity, and $\lambda \in [0, 1)$ captures peer effects in the linear-quadratic form of [2].

Management payoff. The organization earns $A\,\bar{e}$ from engagement, with $A > 0$ the productivity of engaged use at the prevailing adoption level, so that

$$\Pi(I) = A\,\bar{e}(I) - \tfrac{c}{2}I^2. \tag{2}$$

Perception. Management does not observe engagement directly. It evaluates investment through a perceived response $\hat{e}(I) = \mu\,\bar{e}(I)$ with $\mu \in (0, 1]$, and maximizes the perceived payoff $\hat{\Pi}_\mu(I) = A\,\mu\,\bar{e}(I) - \frac{c}{2}I^2$. The case $\mu = 1$ is the correctly specified benchmark; $\mu < 1$ is the underestimation documented in [8].

It is convenient to define the efficient investment scale and the composite displacement burden

$$K = \frac{A\rho}{c(1-\lambda)}, \qquad \theta = \frac{\tau\,c\,(1-\lambda)}{A\rho^2} = \frac{\tau}{\rho K}. \tag{3}$$

The burden θ is the displacement concern divided by the engagement incentive that efficient investment creates; it rises with τ/ρ and with c/A, the two ratios whose thresholds are emphasized in [6], and falls with peer effects.

Assumption 2.1 *$\theta < \frac{1}{2}$ and $\rho K - \tau \le 1 - \lambda$.*

The first inequality makes adoption worthwhile under correct beliefs; the second keeps efficient engagement interior. A *perception equilibrium* of $\Gamma(\mu)$ is a pair $(I_\mu, e^*(\cdot))$ in which $e^*(I)$ is the stage-2 Nash equilibrium for every $I \ge 0$ and I_μ maximizes $\hat{\Pi}_\mu$, ties resolved toward the smaller investment.

3 Main results

Lemma 3.1 *For every $I \ge 0$ the stage-2 game has a unique Nash equilibrium. It is symmetric and given by*

$$e^*(I) = \max\left\{0,\ \frac{\rho I - \tau}{1-\lambda}\right\} \tag{4}$$

whenever $\rho I - \tau \le 1 - \lambda$.

Proof. By strict concavity, every employee's best response is $e_j = P_{[0,1]}(\rho I - \tau + \lambda\bar{e})$, with $P_{[0,1]}$ the projection on $[0, 1]$. Hence $\bar{e}$ solves $\bar{e} = P_{[0,1]}(\rho I - \tau + \lambda\bar{e})$, whose right-hand side is a contraction with modulus $\lambda < 1$ because the projection is non-expansive; Banach's theorem gives a unique fixed point, which fixes every e_j. Direct substitution shows that (4) solves the fixed-point equation in both cases $\rho I \le \tau$ and $0 < \rho I - \tau \le 1 - \lambda$. ■

Theorem 3.2 (Perception gap) *Under Assumption 2.1, the game $\Gamma(\mu)$ has a unique perception equilibrium, with the following properties.*

(i) Dormancy threshold. Let $\mu_c = 2\theta$. If $\mu \le \mu_c$, then $I_\mu = 0$ and $e^* = 0$. If $\mu > \mu_c$, then

$$I_\mu = \mu K, \qquad e^*_\mu = \frac{\rho K\,(\mu - \theta)}{1 - \lambda} > 0. \tag{5}$$

(ii) Profit loss. For $\mu > \mu_c$, realized profit is $\Pi(\mu) = cK^2\big(\mu - \frac{\mu^2}{2} - \theta\big) > 0$, and

$$\frac{\Pi(1) - \Pi(\mu)}{\Pi(1)} = \frac{(1-\mu)^2}{1 - 2\theta}. \tag{6}$$

(iii) ROI surprise. For $\mu_c < \mu < 1$, realized profit strictly exceeds the profit management expected:

$$\Pi(\mu) - \hat{\Pi}_\mu(I_\mu) = cK^2\,(1-\mu)(\mu - \theta) > 0. \tag{7}$$

(iv) Self-confirming dormancy. For $\mu \le \mu_c$ the observable outcome $(I, \bar{e}) = (0, 0)$ is the same for every value of the true skill complementarity $\rho > 0$; the data generated in equilibrium therefore cannot reject management's belief, and the dormant equilibrium is self-confirming in the sense of [7].

Proof. By Lemma 3.1, investments with $\rho I \le \tau$ produce $\bar{e} = 0$ and perceived payoff $-\frac{c}{2}I^2$, so within that range $I = 0$ is optimal with value 0. For $\rho I > \tau$ the perceived payoff is $\hat{\Pi}_\mu(I) = \frac{A\mu}{1-\lambda}(\rho I - \tau) - \frac{c}{2}I^2$. Assumption 2.1 guarantees that the interior formula of Lemma 3.1 applies at every $I \le K$, which contains the relevant range because $\mu \le 1$. Strict concavity gives the unconstrained maximizer $I = \frac{A\mu\rho}{c(1-\lambda)} = \mu K$. Using $\frac{A\rho}{1-\lambda} = cK$ and $\frac{A\tau}{1-\lambda} = cK^2\theta$, both immediate from (3), its perceived value is

$$\hat{\Pi}_\mu(\mu K) = cK\mu \cdot \mu K - \mu\, cK^2\theta - \tfrac{c}{2}\mu^2K^2 = cK^2\,\mu\Big(\frac{\mu}{2} - \theta\Big).$$

This interior candidate is feasible only if $\rho\mu K > \tau$, that is $\mu > \theta$; if $\mu \le \theta$ the supremum over $\rho I > \tau$ is approached at the boundary $\rho I = \tau$ with negative value, so $I_\mu = 0$. If $\theta < \mu$, the candidate beats $I = 0$ if and only if $\mu/2 > \theta$. Hence $I_\mu = \mu K$ when $\mu > 2\theta$ and $I_\mu = 0$ when $\mu \le 2\theta$, the tie at $\mu = 2\theta$ being resolved toward zero. Because $2\theta > \theta$, the dormancy region contains the whole infeasible range. Engagement in (5) follows from (4) and $\tau = \rho K\theta$. Uniqueness follows from uniqueness of the stage-2 equilibrium and strict concavity of $\hat{\Pi}_\mu$ on the active range. This proves (i).

For (ii), substituting $I_\mu = \mu K$ into (2) with the true response gives

$$\Pi(\mu) = cK\mu K - cK^2\theta - \tfrac{c}{2}\mu^2K^2 = cK^2\Big(\mu - \frac{\mu^2}{2} - \theta\Big).$$

At $\mu = 1$ this equals $cK^2(\frac{1}{2} - \theta) > 0$ under Assumption 2.1, and $\Pi(1) - \Pi(\mu) = \frac{cK^2}{2}(1-\mu)^2$, which yields (6). Positivity for $\mu > 2\theta$ follows from $\mu - \frac{\mu^2}{2} - \theta > \mu - \frac{\mu^2}{2} - \frac{\mu}{2} = \frac{\mu(1-\mu)}{2} \geq 0$.

For (iii), subtracting the perceived value from the realized value,

$$\Pi(\mu) - \hat{\Pi}_\mu(\mu K) = cK^2\big(\mu - \mu^2 - \theta + \mu\theta\big) = cK^2(1-\mu)(\mu - \theta),$$

which is strictly positive for $\theta < 2\theta < \mu < 1$.

For (iv), when $I = 0$ Lemma 3.1 gives $\bar{e} = \max\{0, -\tau/(1-\lambda)\} = 0$ independently of ρ, so the outcome distribution is identical across all values of the complementarity parameter. Management's conjecture is correct at the chosen action $I = 0$ and wrong only off path. This is the defining property of a self-confirming equilibrium [7]. ■

Part (iii) inverts the usual reading of positive return surprises: a firm beating its expectations is revealed to have underestimated its workforce, and the surprise is largest at the intermediate bias $\mu = (1+\theta)/2$.

Corollary 3.3 (Learning) *Suppose that an active organization revises its perception after each period according to* $\mu_{t+1} = \mu_t + \eta\,(1 - \mu_t)$ *with* $\eta \in (0, 1]$ *whenever it observes a positive surprise (7). If* $\mu_0 > \mu_c$*, then* $\mu_t \to 1$ *geometrically,* $1 - \mu_t = (1-\eta)^t(1-\mu_0)$*, and the profit loss (6) decays at rate* $(1-\eta)^2$ *per period. If* $\mu_0 \leq \mu_c$*, then* $\mu_t = \mu_0$ *for all* t*. The critical perception* $\mu_c = 2\tau c(1-\lambda)/(A\rho^2)$ *is increasing in* τ *and* c *and decreasing in* A*,* ρ *and* λ*.*

Proof. Active organizations stay active as μ_t rises, and each period yields a positive surprise by Theorem 3.2(iii). The recursion $1 - \mu_{t+1} = (1-\eta)(1-\mu_t)$ gives the closed form, and (6) is proportional to $(1-\mu_t)^2$. Dormant organizations observe no surprise by Theorem 3.2(iv), so the update is never triggered. The comparative statics follow from differentiating 2θ in (3). ■

Organizations above the threshold thus self-correct through return evidence, whereas those below it need an external shock: direct measurement of employee usage, which moves μ, or training (ρ), credible job-security commitments (τ) and peer programmes (λ), which move μ_c below μ without changing beliefs at all.

4 Empirical analysis

The evidence consists of aggregate survey statistics, so the exercise is a calibration with sign checks rather than an econometric estimation. *Data sources.* The McKinsey data come from the report [8] (published 28 January 2025), based on surveys fielded in October and November 2024 of 3,613 employees and 238 C-level executives, 81% from the United States and the rest from Australia,

India, New Zealand, Singapore and the United Kingdom. The leader-employee usage comparisons (Exhibit 2) use the US subsamples ($n = 118$ executives, $n = 3{,}002$ employees). The Deloitte data come from the fourth-quarter report of the *State of Generative AI in the Enterprise* series [4] (January 2025): 2,773 director to C-suite respondents from six industries in fourteen countries (1,203 in the US), surveyed July to September 2024, from organizations screened for AI already in daily use. Table 1 lists every figure used and its location in the source.

Statistic	Value	Model	Source
McKinsey, Superagency in the Workplace [8]			
Leaders' estimate of users at ≥30% of work	4%	$\hat{e}$	Exh. 2
Employees self-reporting that intensity	13%	$\bar{e}$	Exh. 2
Leaders: that intensity now or within a year	20%	$\hat{e}$, 1 yr	Exh. 2
Employees: that intensity now or within a year	47%	$\bar{e}$, 1 yr	Exh. 2
Employees citing formal training as first driver	48%	ρ	Exh. 3
Companies raising AI investment, 3 years	92%	$\mu_t \uparrow$	text
Leaders calling their rollout mature	1%	$\mu \approx 1$	Exh. 11
Deloitte, State of GenAI in the Enterprise, Q4 [4]			
Top initiative meeting / beating ROI expectations	43 / 31%	$\Pi \geq \hat{\Pi}$	p. 22
Workforce with GenAI access	<40%	$I_\mu < K$	p. 17
Scaling ≤30% of experiments in 3 to 6 months	>2/3	$I_\mu < K$	p. 18

Table 1: Survey figures used in the analysis, with exhibit or page of the source report.

Calibrating the perception factor. Since $\hat{e} = \mu\bar{e}$, the ratio of the leaders' estimate to the employees' self-report identifies μ directly. The current-usage pair gives $\hat{\mu} = 4/13 \approx 0.31$; the now-or-within-a-year pair gives $\hat{\mu} = 20/47 \approx 0.43$. By Theorem 3.2(i), investment lies at the same fraction of K.

Calibrating the burden. As θ is not identified, we use $\theta \in \{0.05, 0.10, 0.15\}$, for which $\mu_c = 2\theta$ lies below both estimates. Table 2 evaluates the closed forms of Section 3: the investment ratio $I_\mu/K = \mu$, relative engagement $e^*_\mu/e^*_1 = (\mu - \theta)/(1 - \theta)$, the profit loss (6), and the realized-to-expected profit ratio implied by (7).

Underinvestment is large: at the current-usage estimate, organizations invest about 31% of the efficient amount, forgo 53% to 69% of attainable profit, and reach a fifth to a quarter of efficient engagement. This agrees with limited access and scaling [4] and the 1% maturity figure [8]. The direction of return surprises matches Theorem 3.2(iii): 31% of Deloitte respondents report that their most advanced initiative exceeds return expectations and 43% that it meets them. In the model every active organization with $\mu < 1$ is positively surprised, so the 74% meeting-or-exceeding share is the expected signature of underestimating management rather than of accurate planning. The extreme

$\hat{\mu}$	θ	μ_c	I_μ/K	e^*_μ/e^*_1	Profit loss	$\Pi/\hat{\Pi}$
0.31	0.05	0.10	0.31	0.27	0.53	6.6
0.31	0.10	0.20	0.31	0.23	0.60	9.7
0.31	0.15	0.30	0.31	0.19	0.69	93
0.43	0.05	0.10	0.43	0.40	0.37	4.1
0.43	0.10	0.20	0.43	0.36	0.41	4.9
0.43	0.15	0.30	0.43	0.32	0.47	6.9

Table 2: Calibrated equilibrium outcomes. Profit loss is expressed as a share of attainable profit $\Pi(1)$.

ratio in the third row reflects proximity to $\mu_c = 0.30$, where small perception differences decide between activity and dormancy.

Learning dynamics. The 92% of companies planning higher AI investment [8], and Deloitte's majority expecting to need at least a year to resolve challenges of training and trust, match the correction path of Corollary 3.3. With a moderate revision speed $\eta = 0.5$, an organization starting at $\hat{\mu} = 0.31$ reaches $\mu \geq 0.9$ after $\lceil \ln(0.1/0.69)/\ln 0.5 \rceil = 3$ periods. Dormant firms are absent from screened samples such as Deloitte's, which therefore overstate the share of firms on a correcting path.

Limitations. Leaders and employees in [8] belong to different firms, so $\hat{e}/\bar{e}$ is an average rather than a firm-level perception factor, and the exercise speaks to orders of magnitude and signs only.

5 Concluding remarks

A single documented distortion, management's underestimation of the engagement response, generates three facts that dominate GenAI surveys: widespread underinvestment in complements, a predominance of return outcomes at or above expectation, and a long tail of organizations that stay inactive without learning that they are wrong. The theorem gives each fact a closed form tied to a composite burden that unifies two thresholds of [6]. Measuring actual usage moves μ, while training, job-security commitments and peer programmes move μ_c; either route breaks self-confirming dormancy. Natural extensions are firm-level estimation of μ from matched leader and employee responses and a heterogeneous workforce, connecting the analysis to the uneven gains documented in [3].

Acknowledgements. The author is grateful to Professor Massimiliano Ferrara (Decisions Lab, DiGiES, Università degli Studi Mediterranea di Reggio Calabria) for his supervision and for many helpful discussions on the multilevel framework from which this work originates.

References

[1] F. Aquilino and M. Ferrara, A belief-variance threshold for value co-destruction in multi-level generative AI adoption, *Applied Mathematical Sciences,* **20** (2026), no. 4, 195 - 202. https://doi.org/10.12988/ams.2026.919339

[2] C. Ballester, A. Calvó-Armengol and Y. Zenou, Who's who in networks. Wanted: the key player, *Econometrica,* **74** (2006), no. 5, 1403 - 1417. https://doi.org/10.1111/j.1468-0262.2006.00709.x

[3] E. Brynjolfsson, D. Li and L. Raymond, Generative AI at work, *The Quarterly Journal of Economics,* **140** (2025), no. 2, 889 - 942. https://doi.org/10.1093/qje/qjae044

[4] Deloitte AI Institute, *Now Decides Next: Generating a New Future. Deloitte's State of Generative AI in the Enterprise, Quarter Four Report,* Deloitte, January 2025. https://www.deloitte.com/content/dam/assets-shared/docs/about/2025/quarter-4.pdf

[5] I. Esponda and D. Pouzo, Berk-Nash equilibrium: a framework for modeling agents with misspecified models, *Econometrica,* **84** (2016), no. 3, 1093 - 1130. https://doi.org/10.3982/ECTA12609

[6] 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

[7] D. Fudenberg and D.K. Levine, Self-confirming equilibrium, *Econometrica,* **61** (1993), no. 3, 523 - 545. https://doi.org/10.2307/2951716

[8] H. Mayer, L. Yee, M. Chui and R. Roberts, *Superagency in the Workplace: Empowering People to Unlock AI's Full Potential at Work,* McKinsey and Company, January 2025. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/superagency-in-the-workplace-empowering-people-to-unlock-ais-full-potential-at-work

[9] 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

Received: October 5, 2026; Published: November 6, 2026