Arianna Anniciello

DECISION MAKING & ARTIFICIAL INTELLIGENCE

Tutor: Prof. Elio Masciari

Cycle:XXXVIII

Year: 3

Summary

- My background
- Summary of Study Activities
- My Research Area
- Research Products
- Research Results
- PhD thesis

My background

- MSc degree: Management Engeneering
- Research group: Picus Lab
- PhD start date: 01/11/2022
- Scholarship type: none

Summary of study activities

- Ad hoc PhD courses / Other courses
  - I pilastri della Trasformazione Digitale
  - La Scienza moderna e il problema della disciplina giuridica dell'IA
  - PM Academy PMI PMP Certificate

- Conferences / events attended
  - 2022 IEEE International Conference on Bioinformatics and Biomedicine – IEEE BIBM 2022 – December 6-9, Las Vegas
  - 2023 31st Euromicro International Conference on Parallel, Distributed and Network-Based Processing – PDP 2023 – March 1-3, Naples, Italy – Published paper
  - 2023 European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases - ECML PKDD 2023 – September 18-22, Turin, Italy

Year	Course Title	Type	Credits	Lecturer	Organization
1st	On the challenges and impact of Artificial Intelligence in the Insurance domain	Ad-hoc Course	3	Ing. Lorenzo Ricciardi Celsi	DIETI-Unina
1st	Using Deep Learning properly	Ad-hoc Course	4	Dr. Raffaele Della Corte, DIETI	DIETI-Unina
1st	IoT Data Analysis	Ad-hoc Course	4	Dr. Andrea Apicella, DIETI	DIETI-Unina
1st	Academic Entreprenuership	Ad-hoc Course	4	Prof. P. Rippa, DII	DII-Unina
1st	I pilastri della trasformazione digitale	Ad-hoc Course	3	prof. Nicola Mazzocca - DIETI, Unina	DIETI-Unina
2nd	Percorso per il rafforzamento delle competenze sulla progettazione europea	Ad-hoc Course	3.4	Dr. Tommaso FOGLIA, Dr. Federico PORCEDDA. Dr. Veronica ROCCO	Ministero dell'Università e della Ricerca, Ateneo Federico II
3rd	La Scienza moderna e il problema della disciplina giuridica dell'IA	Ad-hoc Course	6	prof. Lucio Franzese, DIETI - Unina	DIETI-Unina
3rd	ESG per il futuro	B	3.2	Sole 24h	Sole 24h
3rd	PM Academy	B	3.5	Luiss Business School	Luiss Business School
		B			

Research Area

- **Artificial Intelligence & Decision Making**

Research areas

- Computational Social Choice and Multi Criteria Decision Making
  - Problem: Help decision-makers make rational, global, and collective choices
  - Objective: Distilling human expertise and enhancing it through a perpetual learning mechanism driven by feedback data from the actual performance of decisions made.
  - Methodologyes: Our first approach to enhance Decision Making Processes was the application of Clustering Algorithms to Computational Social Choice, then Hybrid Majority Judgment and Multicriteria Decision Making (MCDM) tools → an original Method for Multi Stakeholder Multi Criteria Investment Evaluation
- Artificial Intelligence Risk Management Framework
- Economic Impacts of Artificial Intelligence

Research Results

Multicriteria Majority Judgment

A rating scale in natural language to express a judgment for each criteria for each alternative.

Decision Makers' judgments are aggregated using MJ, finding a *majority grade* for each leaf element.

Judgment on the scale are converted into triangular fuzzy numbers.

Hierarchical recomposition method is applied to get to a colletive global evaluation for each alternative.

MJ rating scale	Min Value	Most Likely Value	Max Value	Fuzzy Number
Very Poor	0	1	2	(1,2,2)
Poor	1	2	3	(3,4,4)
Decent	3	4	5	(4,5,6)
Good	5	6	7	(7,8,8)
Very Good	7	8	9	(8,9,9)
Excellent	8	9	10	(9,10,10)

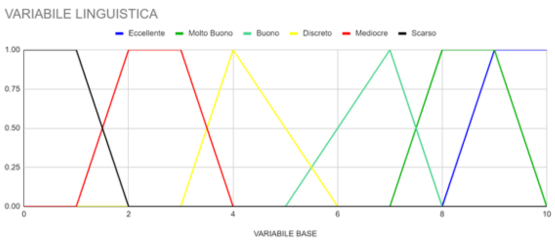

ARIAN

Algorithm 1

Require: $k \geq 0$
Ensure: $n\_winners = (n_1, \ldots, n_k), k > 1$

$k \leftarrow number\_winners$
$max\_cluster \leftarrow k$
$condition \leftarrow$ "ko"
while $condition$ = "ko" **do**
  $cluster\_list \leftarrow cluster(vote\_list)$
  for all list_cluster **do**
    $winners\_per\_cluster \leftarrow compute\_winners(cluster)$
    $all\_winners \leftarrow list\_of\_all\_winners(winners\_per\_cluster)$
  end for
  $list\_winner\_distinct = list\_of\_all\_distinct\_winners(all\_winners)$
  $option\_remaining \leftarrow number\_winners - len(list\_winner\_distinct)$
  if $option\_remaining = 0$ **then**
    $condition =' ok'$
  else
    $k \leftarrow option\_remaining$
    $condition \leftarrow' ko'$
  end if
end while

itee PhD
information technology
electrical engineering

Research Products

Cluster algorithm for social choice *A. Anniciello, E. d’Ajello, D. Formica, E. Masciari, G. Mattia, C. Moscariello, S. Quintarelli and D. Zaccarella,* *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases, ECML PKDD 2022 Workshops, published.*
Covid-19 impact on health information technology: the rapid rise of e-Health and Big Data driven innovation of healthcare processes. – A. Anniciello, S. Fioretto, E. Masciari, E. Napolitano, *2022 IEEE International Conference on Bioinformatics and Biomedicine – BIBM – published*
A Judgment Aggregation Method For Fuzzy Multi Criteria Decision Making *A. Anniciello, E. Masciari,* *B. 31st Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, PDP 2023*
Digital Twins for Traffic Congestion in Smart Cities: a novel solution using Data Mining techniques *A. Anniciello, S. Fioretto, E. Masciari, E. Napolitano,* *2023 15th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management, KMIS 2023*
Decision making with Clustered Majority Judgment. *Anniciello, A., d'Ajello, E., Formica, D., Masciari, E., Mattia, G., Quintarelli, S., & Zaccarella, D.* *In Proceedings of the 26th International Database Engineered Applications Symposium (pp. 156-160).*
How pandemic affected the adoption of e-health systems. *Napolitano, E. V., Fioretto, S., Masciari, E., & Anniciello, A. (2023, May).* *In Proceedings of the 27th International Database Engineered Applications Symposium (pp. 94-98).*
Human in The Loop Generative AI for Explainable Insurance Decision Support *A.Anniciello,S. Fioretto, E.Masciari, E.V.Napolitano,* *iiWAS 2025*

Problem Statement

A relevant question to be solved, of great interest and value for both academia and enterpises, is the evaluation of costs and benefits of Artificial intelligence initiatives in enterprises. Considering the hype on AI, companies need to evaluate carefully which projects to invest on, and there is no simple answer to these questions. In order to enable digital transformation, it is of outost importance to be able to assess costs and benefits of Artificial Intelligence initiatives and give effective tools for strategic portfolio management.

PhD thesis overview

- Problem Statement

  Thought AI investments are rising, only 26% of companies have advanced beyond the proof-of concept stage to generate value, and only 4% of them consistently generate substantial value.

- Objective

  Assessing Expected Benefits of AI Investment in order to enable Decision Maker to make informed decisions and identify Value and Failure Drivers

- Methodology

  Going from static to multi-criteria dynamic modeling of the socio-technological phenomenons that determines the Beneifts Trajectory over time.

Artificial Intelligence Initiative: Business Case process

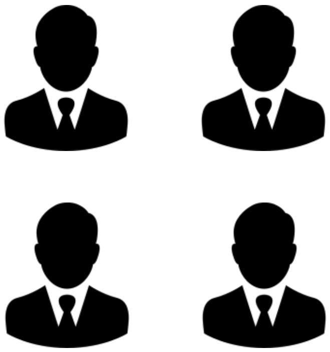

Stakeholders

1 Business Manager: We want AI to improve process X efficiency by 80%

2 CFO: Let's see B/C

3 CIO: We need to outsource skilled resources

4 Risk Manager: are we compliant to regulations?

5 End User: just joking, we're not in the loop

Value and Failure Drivers

- As remarked by many reports and paper, the value of AI depends only for a small percentage on technological maturity and algorithm capabilities. The main value drivers for AI projects in companies hides in people and processess

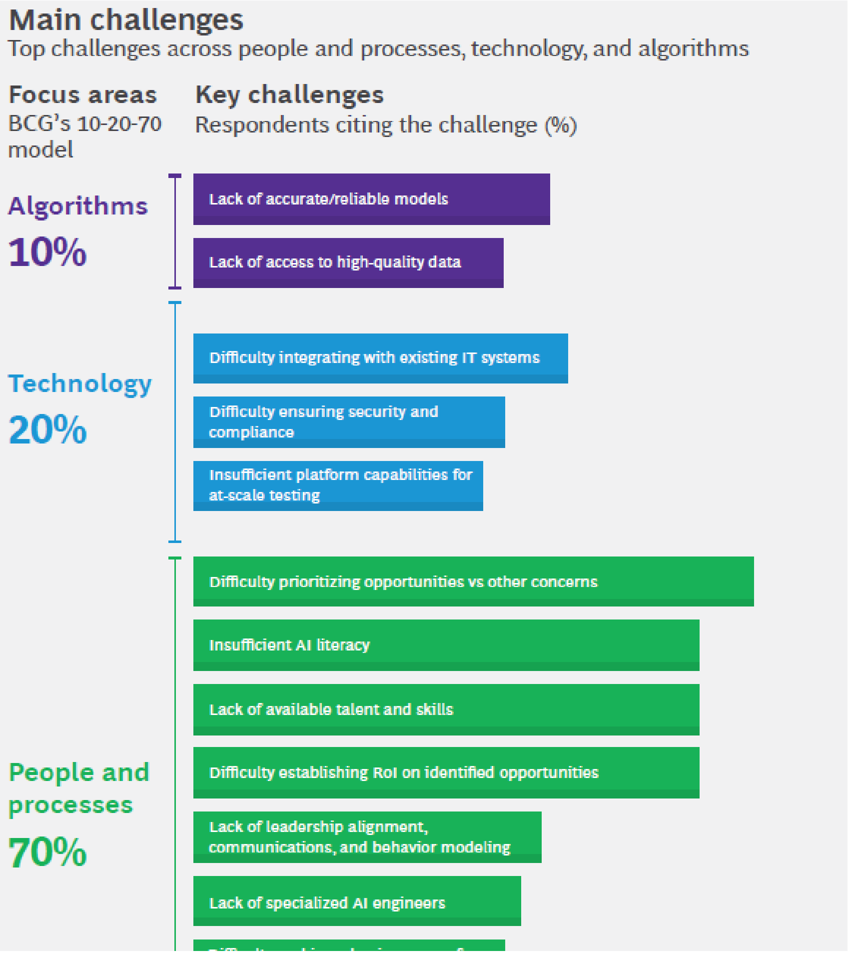

Multi Criteria nature of Artificial Intelligence Benefits

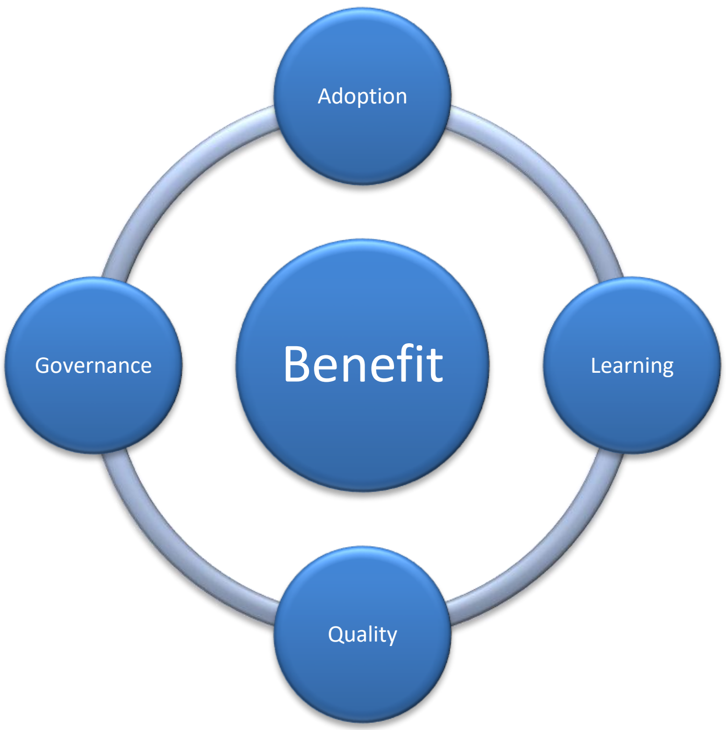

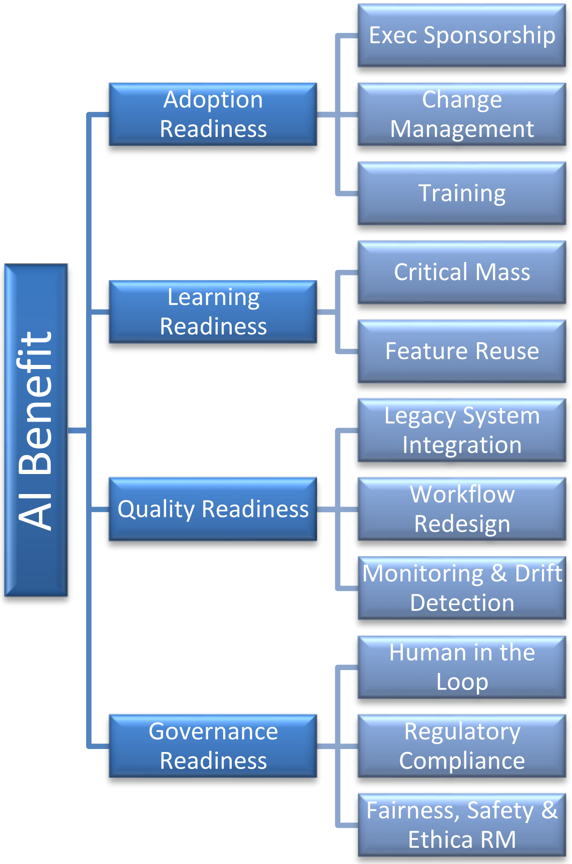

Static evaluation is not enough

- Our DMAI model core is built on a simple intuition: the success of an AI project depends not only on its initial accuracy or efficiency, but on its resilience: in other words it should learn from experience, resist decay, and align with evolving business and regulatory contexts.
- In these conditions, **value is no more an outcome but a trajectory; not a number, but a curve.**

Adoption, Learning, Quality and Governance are time dependent

The DBAI framework builds upon the confluence of three traditions of research:

– **diffusion and adoption theory**, which explains how technologies penetrate social systems through S-shaped patterns {Rogers2003,Moore1991};

– **learning-curve economics**, describing how cumulative experience translates into productivity gains that eventually saturate {Wright1936,Kaplan2020};

– **reliability and quality decay modeling**, which characterizes how systems degrade over time in the absence of maintenance {Sculley2015}.

Adoption Dynamic

Literature on innovation diffusion provides a rich tradition to draw from. Since Rogers' seminal work, adoption has been described as following an S-shaped pattern several mathematical formulations embody this intuition: logistic function, the Gompertz function, and the Bass model.

We choose tp model adoption with an asymmetric Gompertz curve, reflecting inertia, delayed take-off, and saturation in enterprise diffusion.

$$A_i(t; R_A) = \exp\left[-\, \eta_A \, e^{-k_A(R_A)(t-\tau_A)} \right],$$

where

η_A *denotes the initial inertia,*
τ_A *the time of inflection,*
k_A *the adoption velocity*
R_A *the organizational readiness, defined with our Fuzzy AHP MJ framework*

This functional form captures the slow early phase of experimentation, the acceleration driven by early-majority adoption, and the subsequent saturation as full integration is achieved.
Managerially, $A_i(t, R_A)$ *measures how effectively the enterprise transforms AI capabilities into widespread operational use.*

Adoption Dynamic

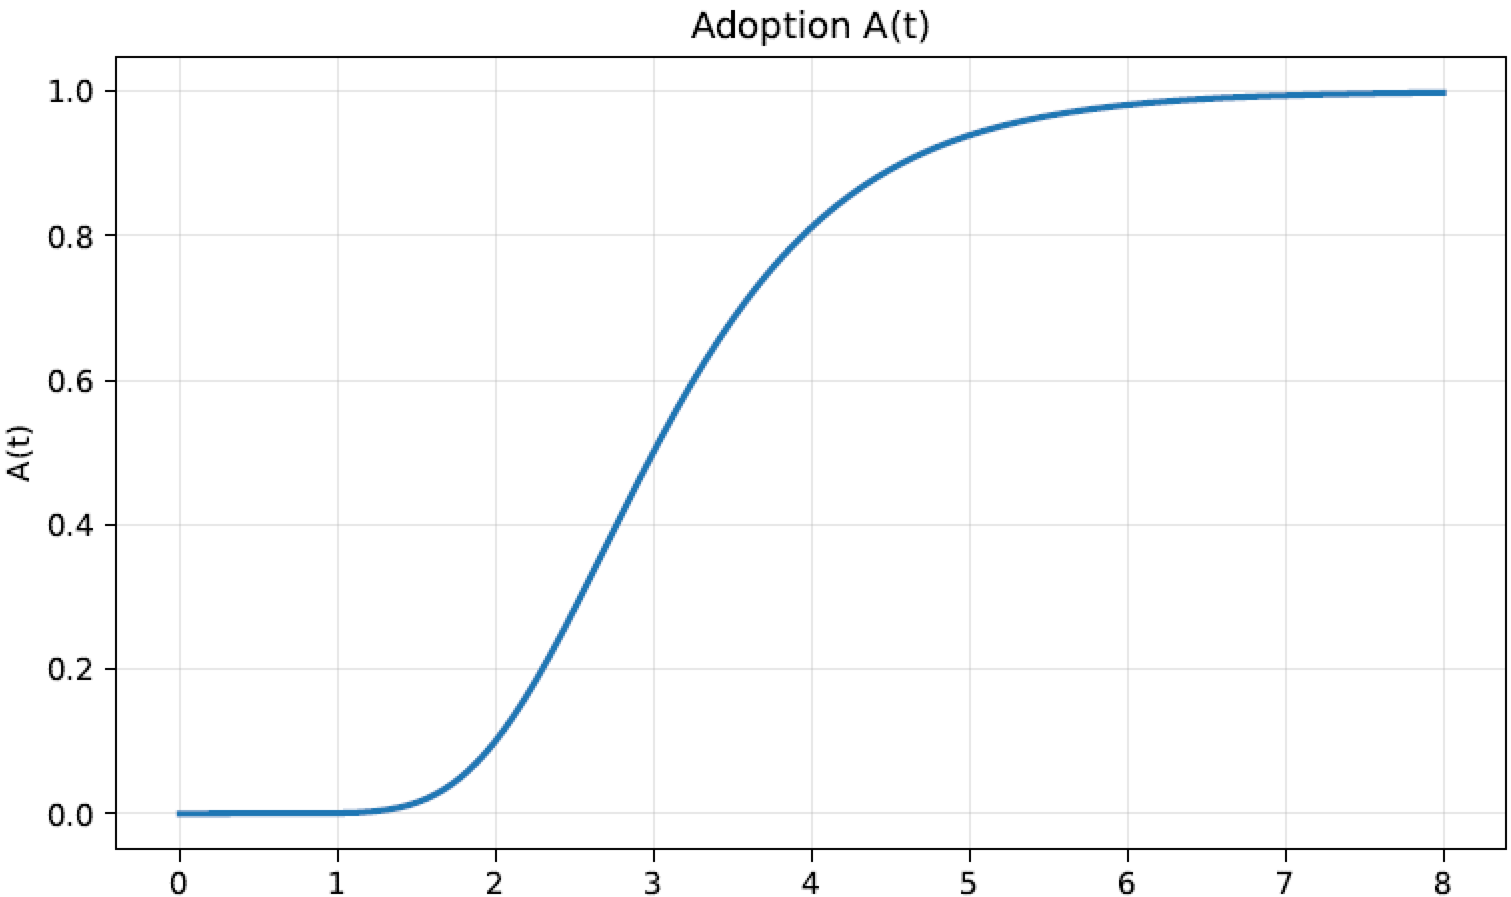

$A(t) = \exp(-\eta e$ [illegible] $)$;

Learning Dynamic

Large-scale studies on scaling laws {kaplan2020,hoffmann2022} demonstrate that model performance, measured for example by loss or accuracy, improves predictably as a power-law function of training data and compute.

The more users interact with an AI system, the more feedback is collected, and the more opportunities exist for retraining and refinement.

Yet the incremental value of each additional user or dataset decreases over time, because the system has already learned the most common patterns and only rarer or more subtle cases remain to be captured.

$$L_i(t; R_L) = \frac{S_{\mathrm{eff}}(t)^n}{S_{\mathrm{eff}}(t)^n + S_{50,i}(R_L)^n}$$

We modeled Learning as a Hill function, with a saturating response of an effective stock fed by adoption.

with $S_i^{\mathrm{eff}}(t) = w_i^{\mathrm{DD}}\,(1+\sigma_i)\,S_i^{\mathrm{base}}(t)$ $\qquad$ $S_i^{\mathrm{base}}(t) = \int_0^t \rho\, A(u)\, du.$

where

$w_i^{DD} \in (0,1]$ *penalizes projects affected by data debt,*

$\sigma_i > 0$ *accounts for micro-synergies,*

ρ represents the density of data produced per unit of adoption.

Learning Dynamic

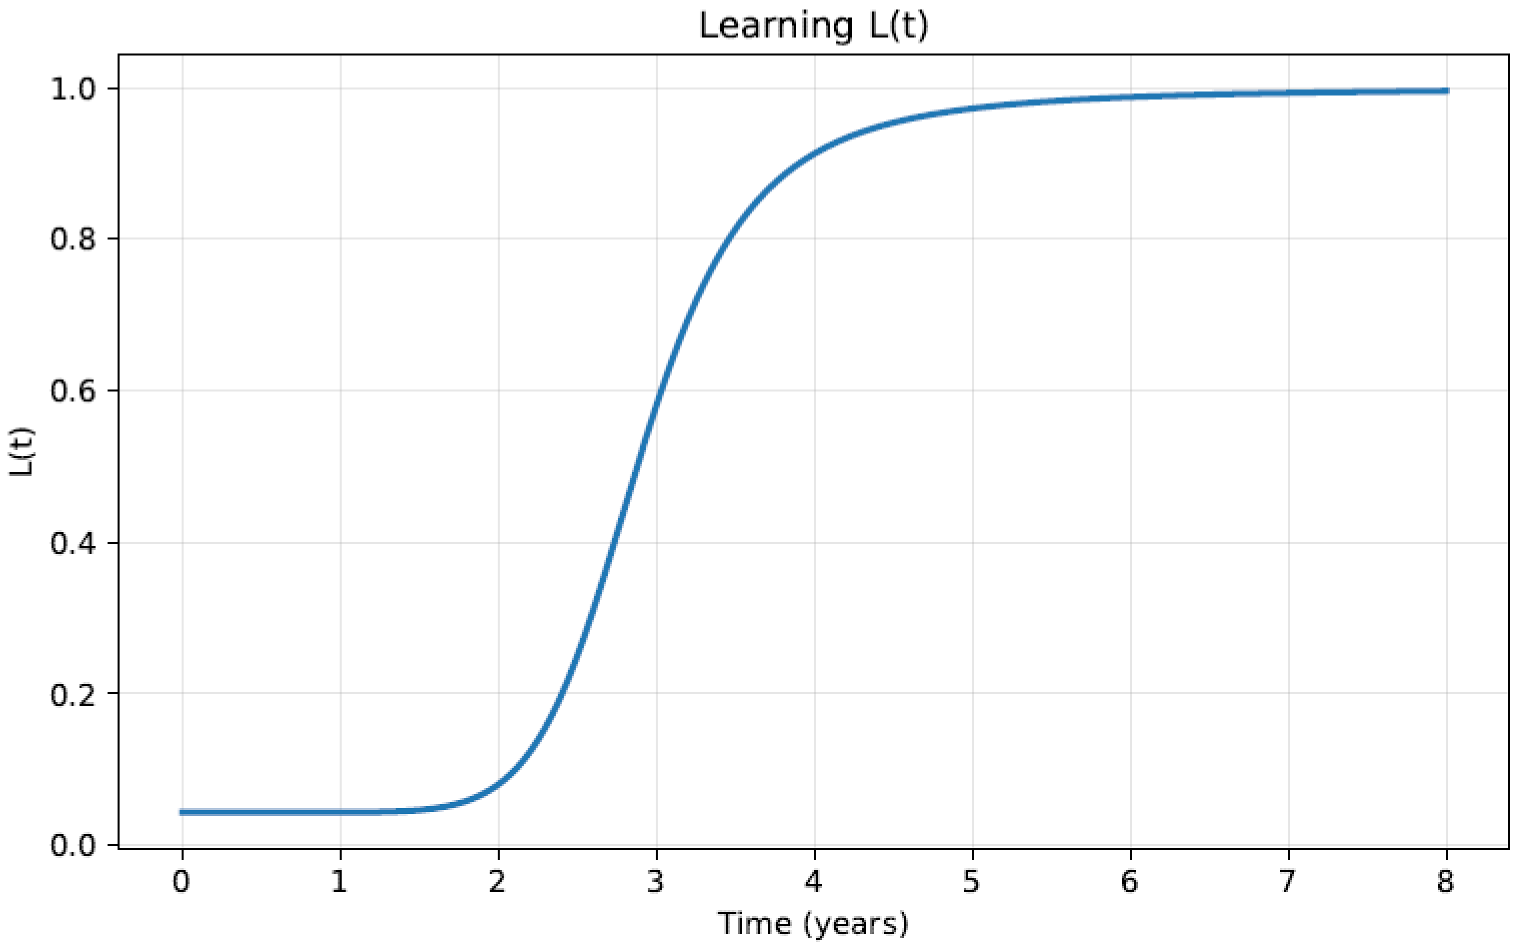

Quality Dynamic

Classical reliability growth models, such as those developed by Musa and Okumoto {musa1984}, show how defect rates decrease over time as failures are discovered and corrected, but also how systems never reach perfect reliability.

Similarly, in machine learning contexts, studies on concept drift demonstrate that model accuracy decays unless monitoring and retraining are systematically implemented {gama2014}.

We adopt exponential decay with a floor to capture temporal dataset shift and model aging without refresh; readiness R_Q moderates the effective decay rate $\Phi_{eff,i}$ meaning slower decay and longer periods of stable performance.

$$Q_i(t) = q_{\min} + \left(Q_0 - q_{\min}\right) e^{-\Phi_{\mathrm{eff},i}(R_Q)\,t},$$

Where

Q_0 is the initial operational quality at deployment, representing the baseline accuracy or fitness level achieved under controlled training conditions,

q_{min} denotes the minimum attainable quality, corresponding to the residual usefulness of the system once degradation stabilizes and periodic maintenance cannot fully recover performance,

$\Phi_{eff,i}$ is the effective decay rate, defined as $\Phi_{eff,i}$ R_Q = $\Phi_{max}(1 - R_Q)\gamma$ where Φ_{max} is the maximum deterioration rate observed in the absence of governance,

R_Q is the readiness index derived from multicriteria assessment of data quality, monitoring, and MLOps maturity,

$\gamma > 1$ controls the curvature of the relationship between readiness and quality decay.

Quality Dynamic

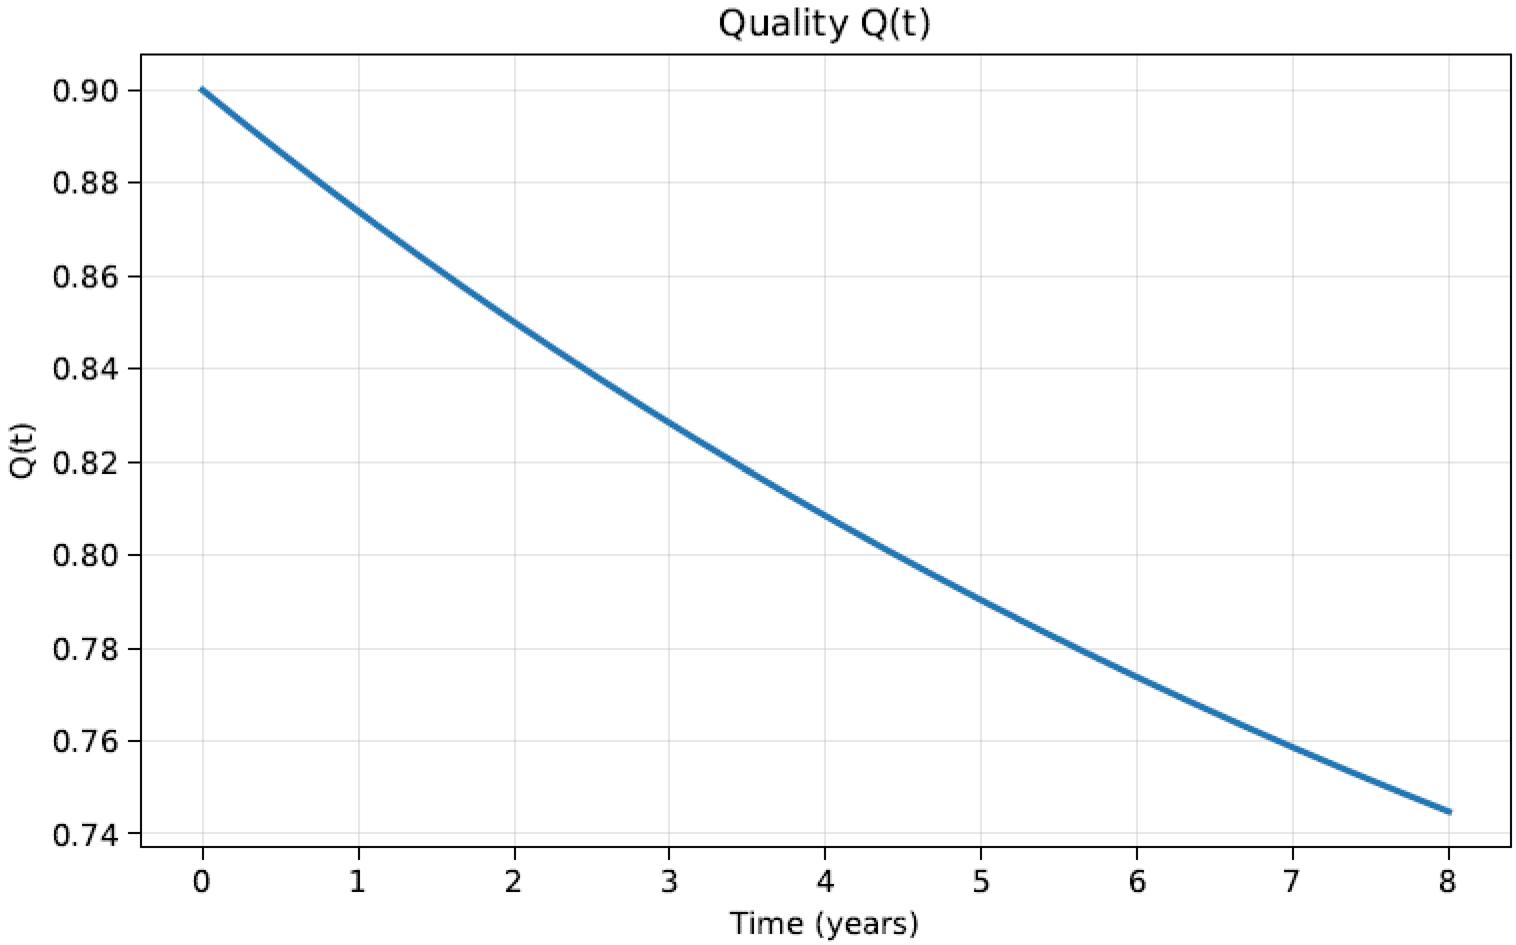

Governance Dynamic

The compliance component represents the organizational capacity to meet regulatory, ethical, and social requirements over time:

compliance dynamics follow a sigmoid pattern, starting from low levels of trust and progressively increasing as the organization implements governance controls, validation processes, and transparency measures, in line with continuous TEVV and governance-by-design recommended by NIST AI RMF and consistent with EU AI Act obligations.

$$G(t) = \frac{1}{1 + \exp[-k_G\,(t - \tau_G(R_G))]}$$

The readiness index canticipates the timing of these transitions: mature organizations reach compliance thresholds earlier, thereby reducing the risk of project interruption or reputational loss.

Governance Dynamic

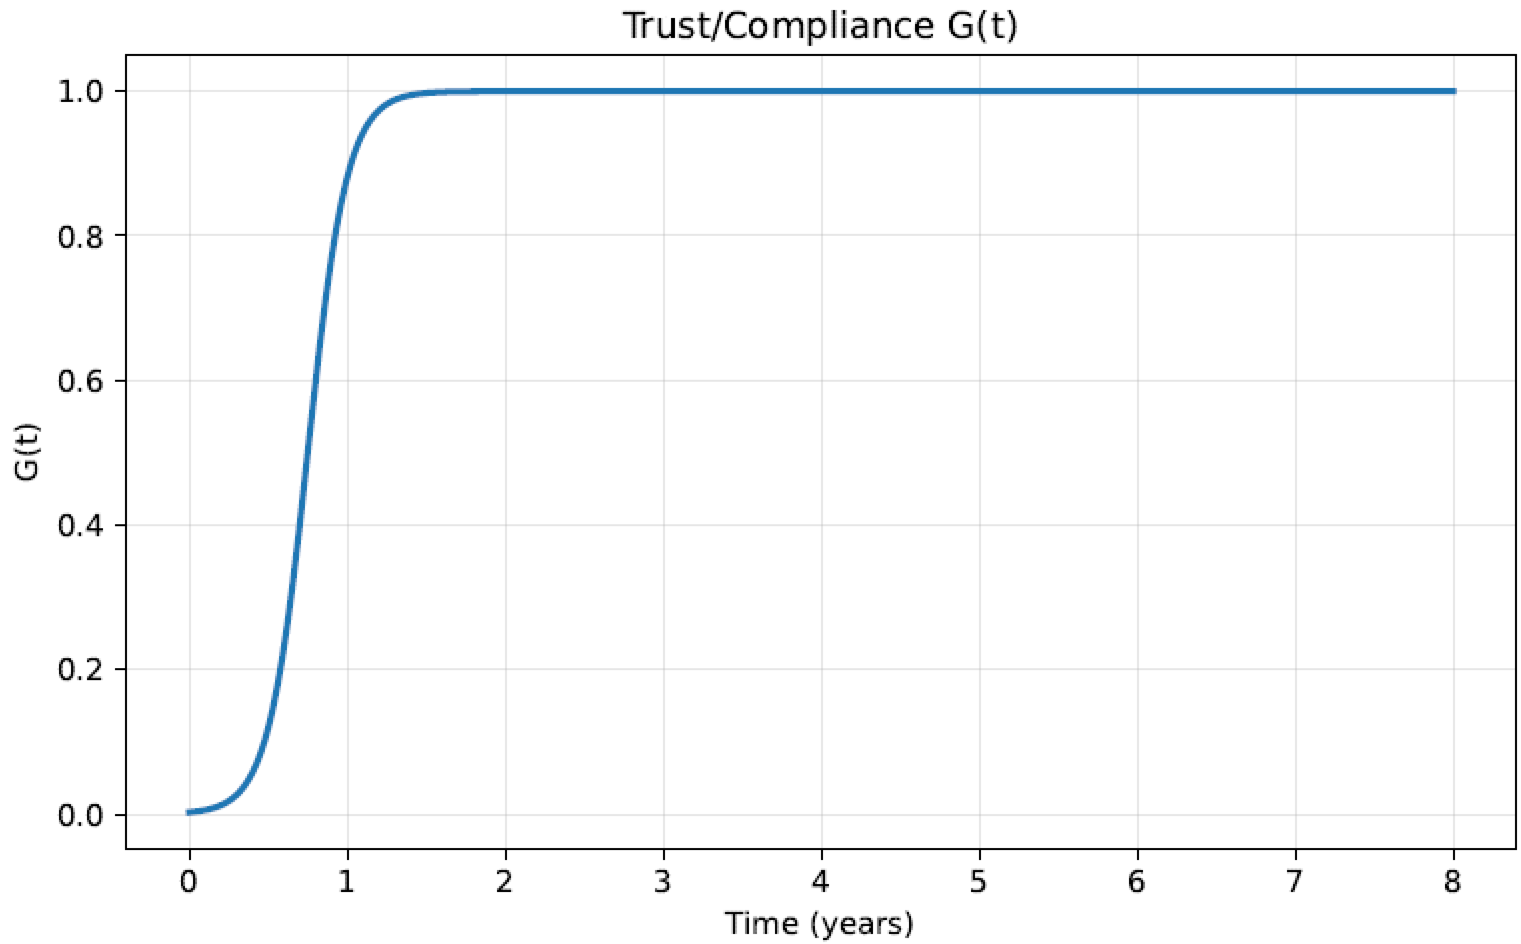

DBAI

The general form of the benefit function is expressed as:

$$B_i(t) = B_{\max,i}^{(\text{risk-adj})}\, A_i(t;R_A)^{\alpha_A}\, L_i(t;R_L)^{\alpha_L}\, Q_i(t;R_Q)^{\alpha_Q}\, C_i(t;R_C)^{\alpha_C}$$

The multiplicative structure in DBAI formulation is deliberate. Each dimension represents a necessary condition for benefit realization: high adoption without compliance readiness generates little usable value; strong learning effects cannot compensate for poor operational quality; and conversely, robust governance is irrelevant without user adoption.

The term $B_{max,i}^{(risk.adj)}$ represents the theoretical maximum benefit achievable under ideal conditions, corrected for managerial and cognitive biases that typically distort ex-ante estimations.

The four exponents α_A, α_L, α_Q, and α_G denote the relative importance of each component and are determined through the Analytic Hierarchy Process (AHP) using pairwise comparisons among criteria.

These exponents can be interpreted as elasticity coefficients: they quantify the sensitivity of total benefit to changes in each underlying process.

By combining these factors multiplicatively, the model mirrors the interdependence observed in real-world AI deployments, where the weakest link often determines overall success.

DBAI

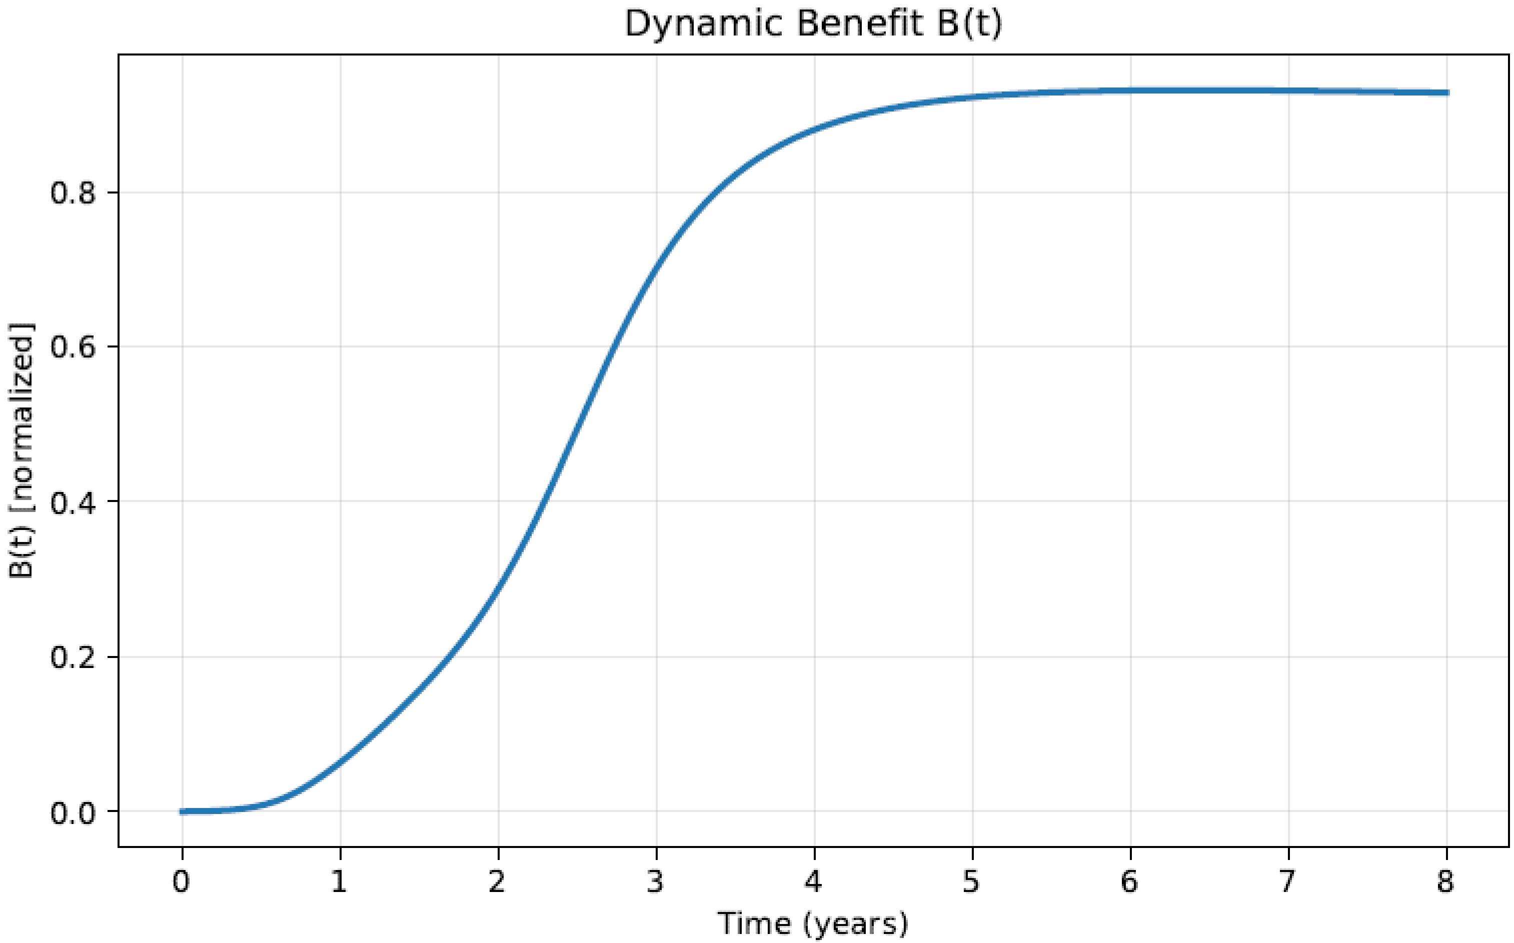

Dynamic Benefit B(t)
B(t) [normalized]
0.0
0.2
0.4
0.6
0.8
0
1
2
3
4
5
6
7
8
Time (years)

BIAS in Management Expectation

In traditional business case analysis, B_{max} is typically inferred from expected ROI or forecasted savings, under the implicit assumption that

technical performance translates directly into business performance.

However, empirical evidence from industrial AI deployments demonstrates that this translation is far from perfect: projects often overestimate potential benefits due to

- **misalignment with strategic objectives**, proxied by the proportion of project KPIs not directly traceable to organizational OKRs;
- **managerial literacy**, measured as the fraction of managers untrained or uncertified in AI;
- **technological feasibility**, quantified as the percentage of AI proof-of-concept initiatives that fail to scale within the evaluation horizon
- and **time-horizon expectations**, deviation between the expected ROI period of the project and the organization's average investment cycle, normalized by the latter.

To account for these deviations, the DBAI introduces a bias adjusted term defined as

$$R_b(t) = 1 - \theta_{\text{base}} - \left(\beta_1\, M_{\text{mis}} + \beta_2\, M_{\text{lit}} + \beta_3\, M_{\text{hype}} + \beta_4\, M_{\text{fit}}\right)$$

The multiplicative structure ensures that biases compound rather than cancel each other, reflecting the empirical observation that overconfidence, strategic drift, and hype exposure often co-occur in failed AI programs.

Capex and Opex

Rather than aggregating all expenditures into a static value, the formulation distinguishes between non-recurring and recurring components, allowing the cost trajectory to reflect the transition from an initial investment phase to steady-state operations.

$$C_i(t) = C_i^{\mathrm{nrec}}(t) + C_i^{\mathrm{rec}}(t)$$

$$C_i^{\mathrm{nrec}}(t) = I_i \, e^{-\delta_i t}$$

$$C_i^{\mathrm{rec}}(t) = \left(c_i^{\mathrm{use}}(t) + c_i^{\mathrm{obs}}(t)\right) \left(1 + \theta_i\right)$$

The first component, represents the non-recurring investment cost associated with the initial deployment of the AI system.
The parameter I_i denotes the magnitude of the upfront investment, while $\delta_i$controls the rate at which these initial costs decay over time, modeling the progressive completion of deployment activities and transition to steady-state operations.

The second component captures the variable recurring cost of usage. It scales proportionally with the adoption curve Ai(t), which represents the share of processes or users actively leveraging the AI system at time t.

This proportional structure ensures that recurring costs grow in step with actual utilization, reflecting the scalability of AI services: low costs during the early adoption phase and higher costs as adoption approaches saturation.

Managerially, this term highlights how usage-related expenditures act as a financial proxy for system load, compute demand, and data-processing intensity.

Capex and Opex

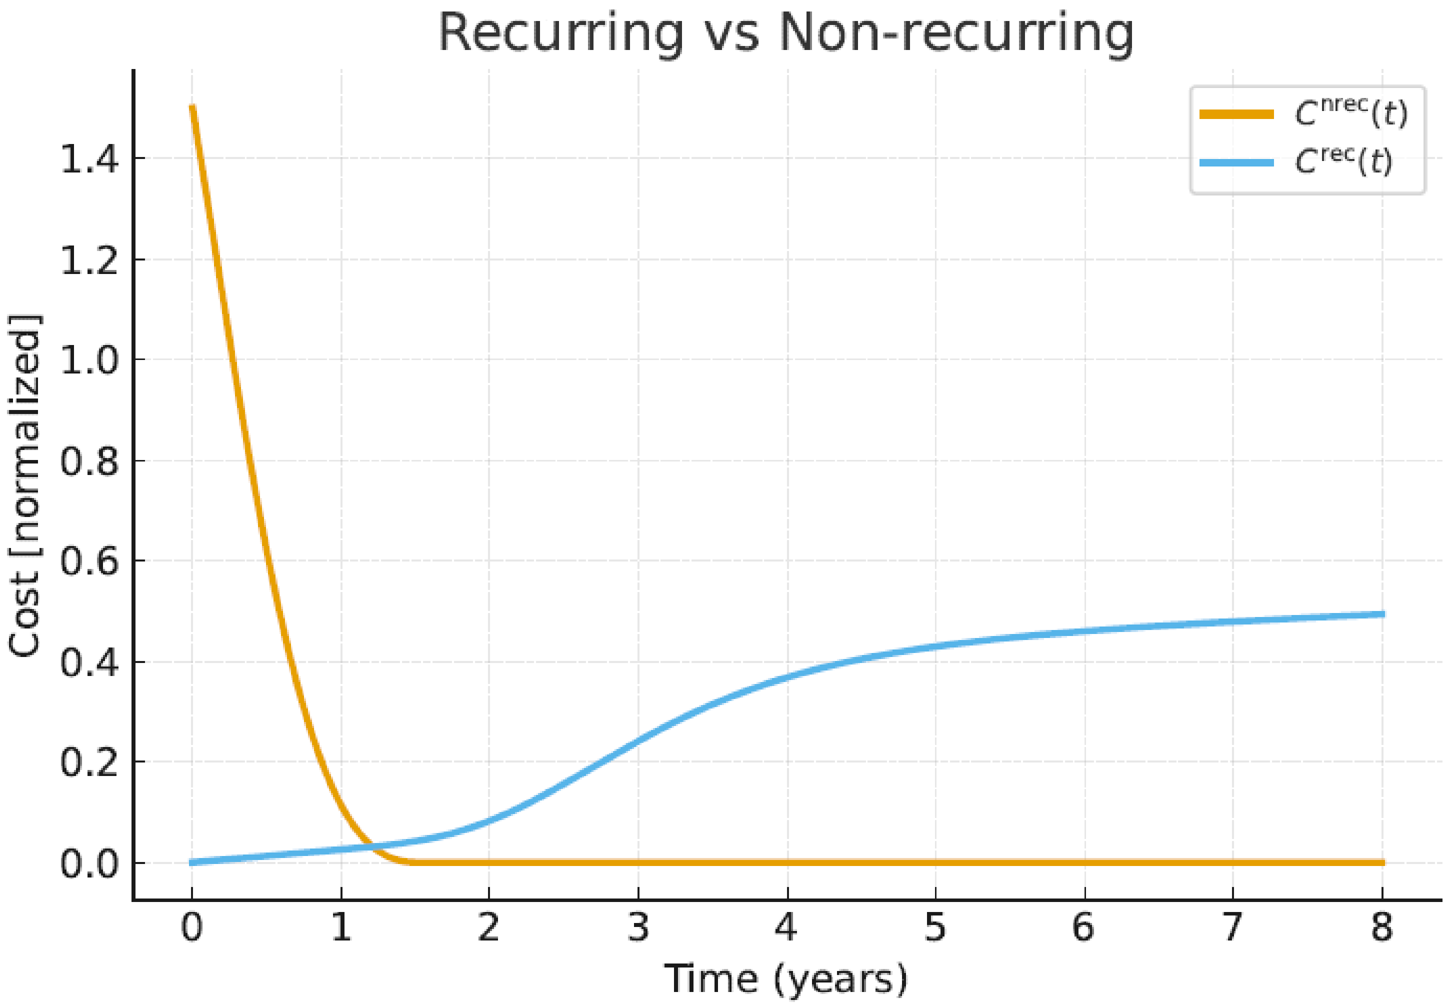

B/C Ratio and J curve

When combined, these four curves generate trajectories of benefit realization that are markedly different from the static projections of traditional financial methods.

They naturally reproduce J-shaped profiles, in which early costs outweigh benefits until adoption and learning accelerate, quality stabilizes, and governance unlocks deployment.

They also make explicit the levers available to management: investing in change management reduces inertia in A(t), strengthening data pipelines increases ρ and decreases $\mathrm{R_Q}$ *in L(t), maturing MLOps practices raises* $\mathrm{R_Q}$ *in Q(t), and advancing transparency and fairness improves* R_G *in G(t).*

In this way, the model is not only descriptive but prescriptive, highlighting how organizational choices translate into altered trajectories of value.

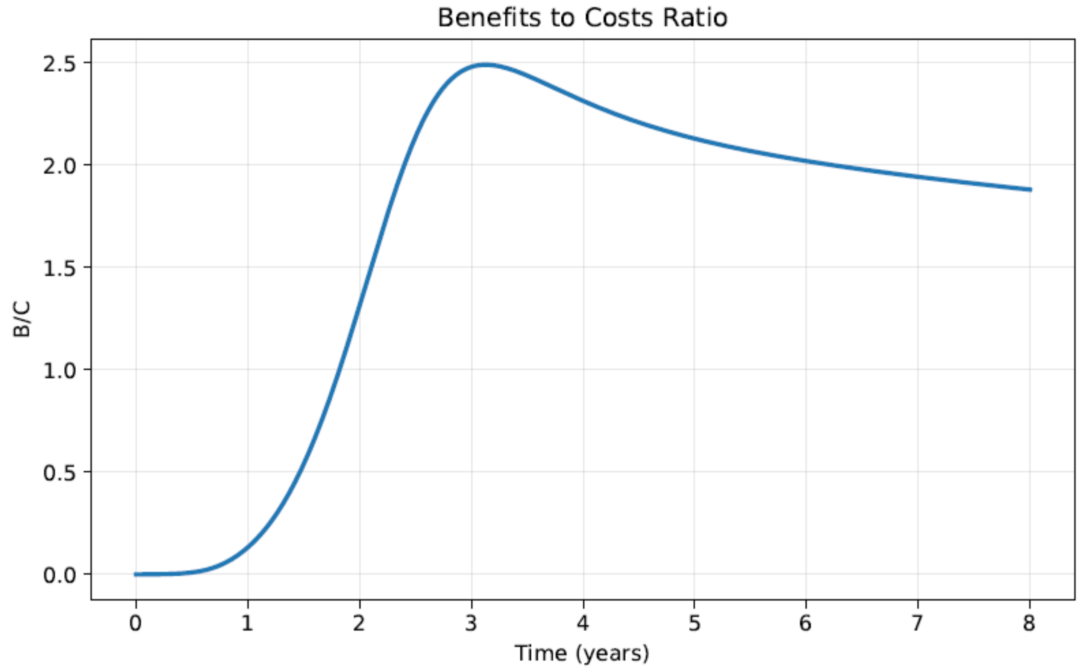

B/C Ration and J curve

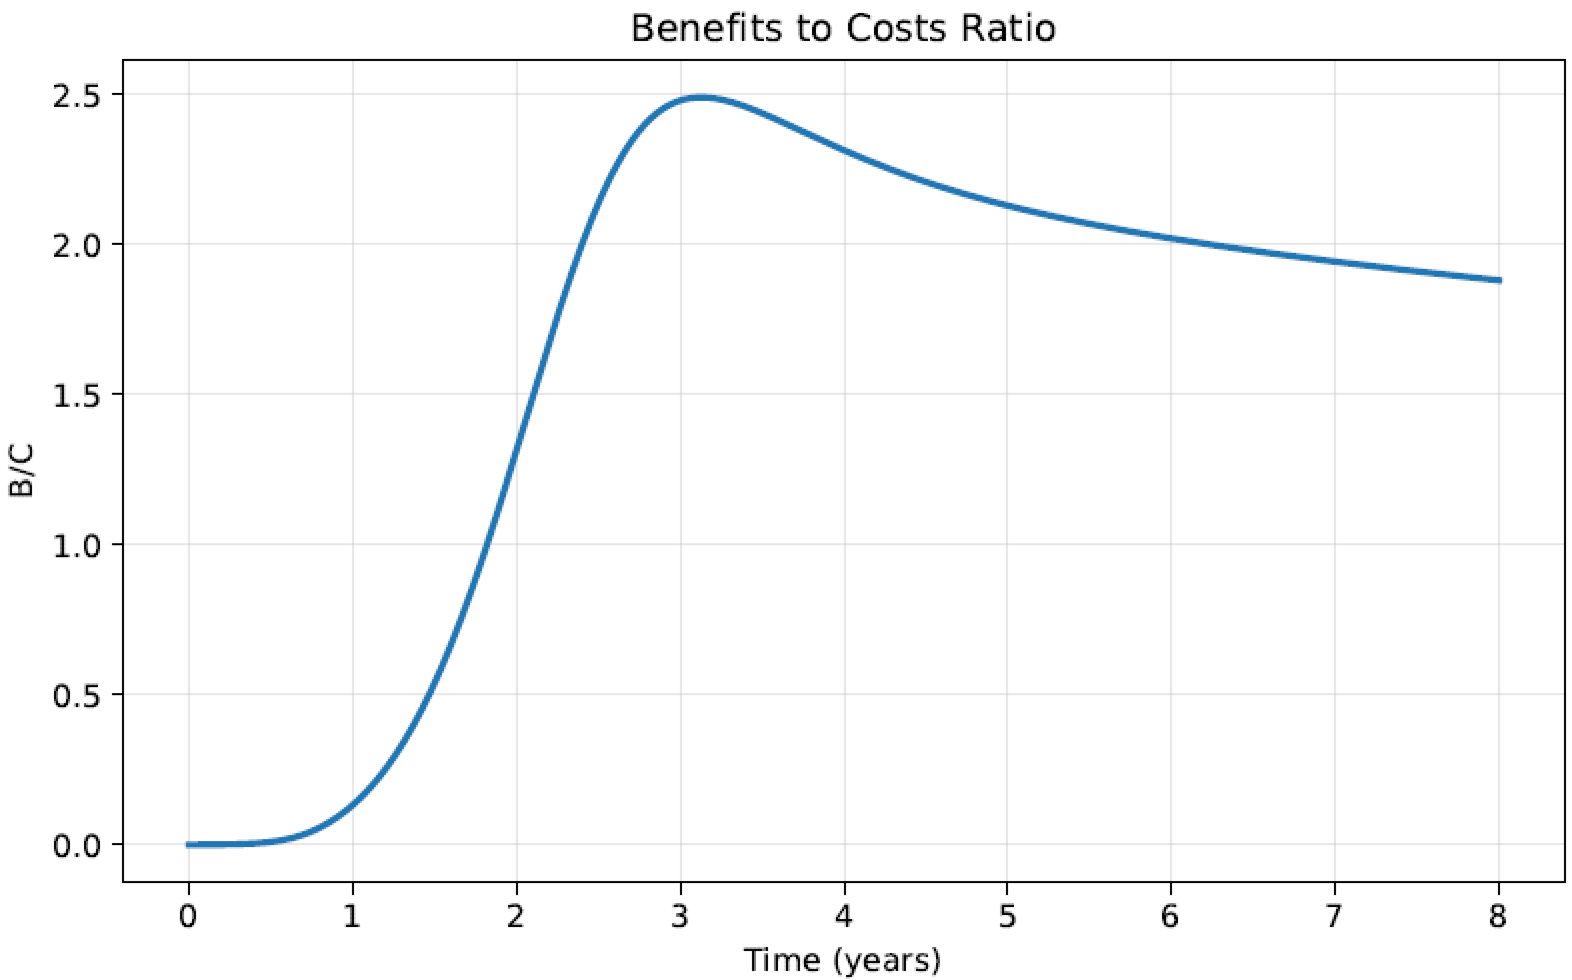

Validation and Application

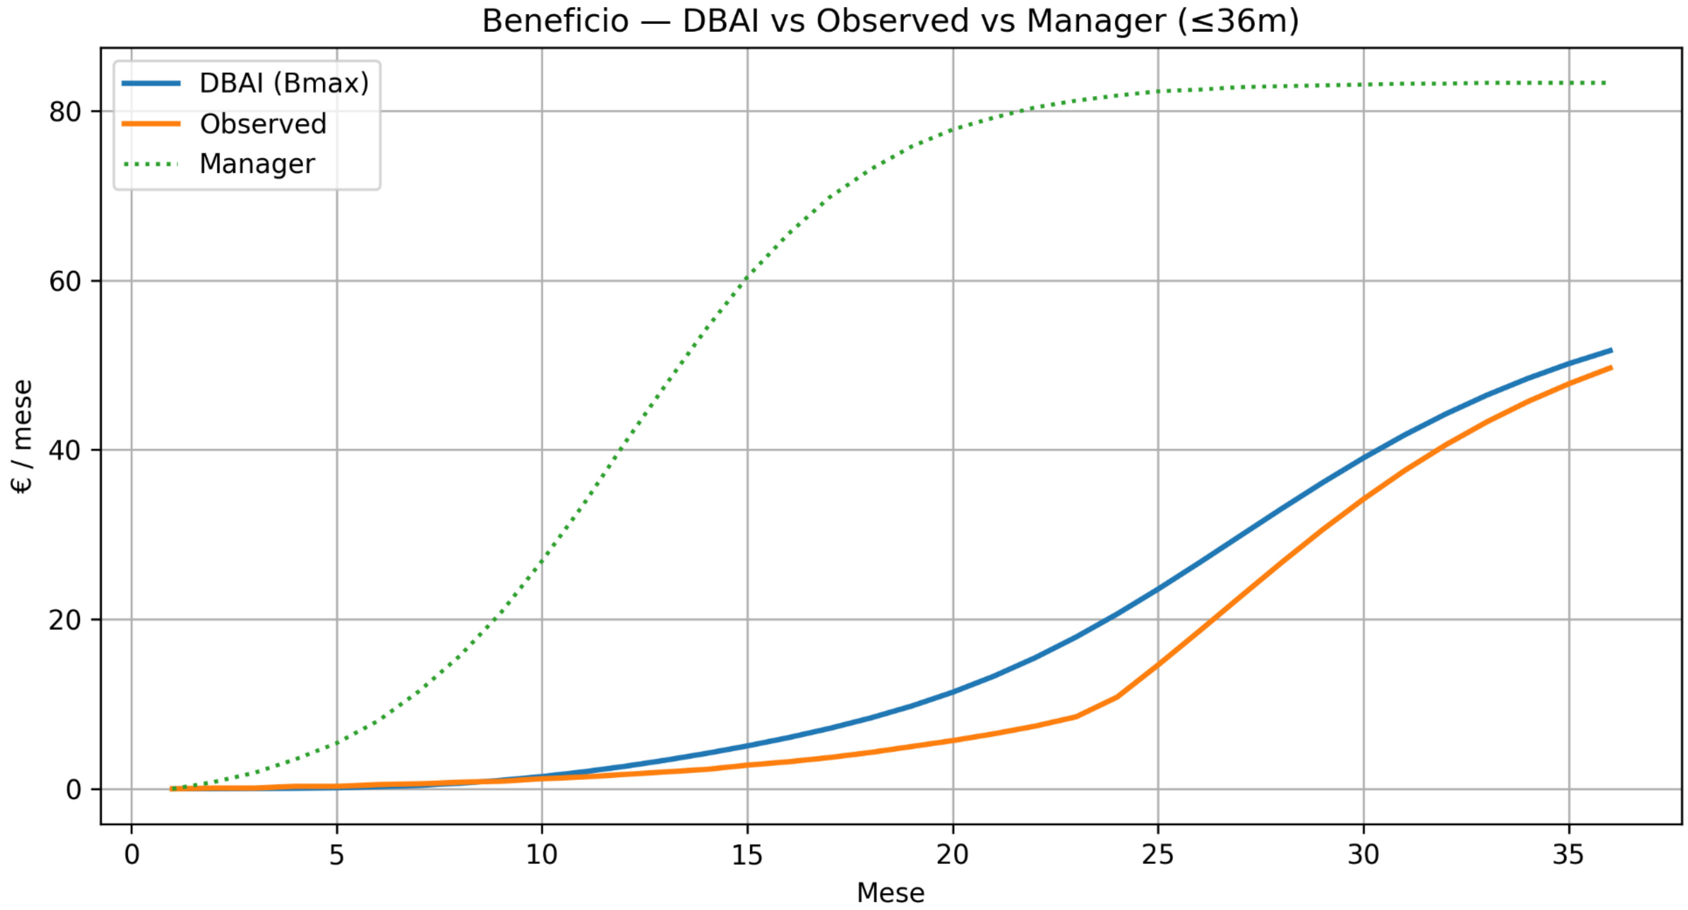

Last Remarks

By integrating a multicriteria approach into a dynamic model, our novel framework contributes to the existing state-of-the-art in the field of Artificial Intelligent benefits evaluation, providing a tool for enterprises to perform a more reliable benefit/cost analysis for their business cases and also enhancing strategic decision making to identify the most impactful improvements to gain bigger benefits, faster.

Thank you for your attention