




The Diffusion of Artificial Intelligence in Slovak Enterprises

Forecasting Adoption and Its Trajectory

*Vladimír Hojdík
Andriana Baziuk
Simona Kosztanko*

EKONOMICKÁ UNIVERZITA V BRATISLAVE
FAKULTA PODNIKOVÉHO MANAŽMENTU

The Diffusion of Artificial Intelligence in Slovak Enterprises
Forecasting Adoption and Its Trajectory

Vladimír Hojdík
Andriana Baziuk
Simona Kosztanko

Vydavateľstvo EKONÓM
Bratislava 2026

Title: The Diffusion of Artificial Intelligence in Slovak Enterprises: Forecasting Adoption and Its Trajectory

Authors: Vladimír Hojdík, Andriana Baziuk, Simona Kosztanko

Edition

1st edition, 2026

Published by

Vydavateľstvo EKONÓM, Bratislava University of Economics and Business, Dolnozemská cesta 1, 852 35 Bratislava, Slovak Republic

Publication type

Scientific monograph

Reviewed by

doc. Ing. Mojmír Kokles, PhD.

Ing. Róbert Šlosár, PhD.

Copyright: © 2026 Vladimír Hojdík, Andriana Baziuk, Simona Kosztanko.

License: This work is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>.

Funding: This work was supported by the European Union – NextGenerationEU through the Recovery and Resilience Plan for Slovakia under project No. 09I05-03-V02-00003/2025/VA (AI-ImpactSK); 100% share.

ISBN: 978-80-225-5357-5

Foreword

This monograph turns from the description and comparison of artificial intelligence (AI) adoption toward its prediction. Building on prior evidence positioning Slovakia within the international benchmark landscape and assessing AI adoption among Slovak enterprises through dedicated enterprise survey waves, it models the diffusion of AI across Slovak enterprises and derives forward-looking scenarios of its future trajectory.

The monograph deploys two complementary modelling approaches. The first is the Bass diffusion model, applied both to the sample as a whole and to eight functional areas of AI use, which characterises the shape and timing of the adoption curve and the mechanism driving it. The second is a discrete-time Markov chain over three adoption states — non-adopter, planner and adopter, which makes explicit the transitions between these states and, in particular, the role of the planner pool that the binary logic of the Bass model necessarily abstracts away. The modelling is preceded by a descriptive analysis of the sample and followed by a dedicated robustness assessment, so that the forecasts are presented together with a transparent account of their sensitivity to the underlying assumptions.

The monograph complements empirical and comparative evidence on AI adoption with a predictive perspective. Its findings are addressed to researchers, policymakers and enterprise decision-makers seeking to understand not only where AI adoption stands today, but how quickly, in which functional domains, and through which behavioural transitions it is likely to spread in the coming years.

Consistent with the methodological caution that characterises this study as a whole, the forecasts presented here are explicitly conditional on a defined set of assumptions, which are stated transparently and revisited in the discussion of limitations and in the robustness analysis. They should be read as structured, model-based scenarios rather than as unconditional predictions.

Executive Summary

This monograph forecasts the diffusion of artificial intelligence across Slovak enterprises using two complementary models applied to a sample of 1,128 firms, of which 741 already use AI, 216 plan to adopt it, and 171 neither use nor plan to use it.

The Bass diffusion model, estimated for the sample as a whole and for eight functional areas, shows that AI adoption is overwhelmingly driven by imitation rather than by autonomous innovation: across every analysis the imitation coefficient exceeds the innovation coefficient by one to two orders of magnitude. Aggregate adoption reached 65.7% of firms by 2025, with new adoptions peaking around 2024 and 95% saturation of the sample projected around 2029. Adoption is highly uneven across functions: marketing and sales leads at 28.8% penetration in 2025 and approaches saturation in the early 2030s, while finance and ICT security trail at around 10% and saturate only in the late 2030s. The more penetrated a function is today, the earlier its peak and saturation.

The Markov chain model, which tracks transitions between the non-adopter, planner and adopter states, complements this by foregrounding the planner pool as the principal near-term driver of future adoption. Under an absorbing adopter state the sector converges toward near-universal use during the 2030s (87–99% by 2030 across scenarios); under a churn variant with a 5% annual abandonment rate, adoption instead plateaus at 71–87%, leaving a structural residual of non-adopters. A robustness analysis confirms that the imitation-driven character of diffusion and the timing of saturation are insensitive to the take-off year and to the choice of market potential, while the long-run plateau is sensitive to the abandonment rate.

For policy, the dominance of imitation implies that interventions which make successful adoption visible – case studies, benchmarks and demonstrator projects are likely to be more effective than campaigns aimed at stimulating independent first-mover adoption. The findings provide a predictive layer that links comparative evidence on Slovakia's AI position to the analysis of firm performance.

Table of Contents

Foreword.....	4
Executive Summary.....	5
List of Figures.....	9
List of Tables.....	10
List of Abbreviations	11
1 Introduction: From Comparative Diagnosis to Forward-Looking Analysis	12
1.1 Artificial intelligence, diffusion, and the case of Slovakia	12
1.2 The comparative foundation: Slovakia's position in the European AI landscape.....	13
1.3 The survey-based foundation: structure and dynamics of adoption in Slovak enterprises ...	14
1.4 From comparative diagnosis to forecasting.....	15
1.5 Objectives, contribution and structure	15
2 Theoretical Foundations and a Review of the Evidence on AI Adoption and Diffusion	17
2.1 Conceptual foundations and the theoretical framework	17
2.1.1 Defining artificial intelligence	17
2.1.2 Defining AI adoption.....	18
2.1.3 Determinants of firm-level AI adoption.....	18
2.1.4 The scale and recency of enterprise AI adoption	19
2.2 Innovation diffusion theory.....	19
2.3 Formal models of diffusion.....	20
2.3.1 The Bass diffusion model.....	20
2.3.2 State-based and Markov-chain models of adoption	21
2.3.3 Why two models? The logic of triangulation.....	22
2.4 From IT adoption to AI adoption: what is different	22
2.5 Empirical evidence on firm-level AI adoption.....	23
2.6 Determinants, capabilities and the maturity perspective	24
2.7 Forecasting AI diffusion: modelling approaches and open questions	25
2.8 Research gap and the positioning of this monograph	25
3 Methodology, Data and Sample: Extended Specification	27
3.1 Research design and the logic of a two-model strategy	27
3.2 Data provenance and the survey instrument	28
3.3 Variable construction and the three-state classification	28
3.4 Bass model: specification, estimation and uncertainty	29
3.5 Markov chain: state space, transition structure and scenarios.....	30

3.6 Reproducibility, software and triangulation	31
3.7 Methodological limitations and their bearing on interpretation	31
4 Descriptive Analysis of the Sample	33
4.1 AI-Use Status and Adoption Timing	33
4.2 Breadth and Co-adoption of AI	35
5 Aggregate AI Diffusion	37
6 AI Diffusion Across Functional Areas	39
6.1 Marketing & Sales	39
6.2 Production & Operations	40
6.3 R&D and Innovation	41
6.4 Logistics	42
6.5 Corporate Management	43
6.6 Human Resources	44
6.7 Finance	45
6.8 ICT Security	46
6.9 Cross-Functional Comparison	46
7 The Functional Architecture of AI Diffusion: A Deeper Cross-Functional Analysis	48
7.1 A unified comparative framework	48
7.2 The penetration–timing regularity and the ordered diffusion	49
7.3 The economics of heterogeneity: three tiers of diffusion	49
7.4 Convergence speed and the information in the imitation coefficient	50
7.5 Intra-firm diffusion pathways: from functions to bundles	51
7.6 The shifting centre of gravity of new adoption	51
7.7 Quantifying the functional gap and its closing	52
7.8 Cross-validation with the survey evidence and the international record	52
7.9 Uncertainty and what can be safely concluded	53
7.10 Implications of the functional structure for sequencing	53
8 State-Transition Dynamics: A Markov Chain Perspective	54
8.1 Transition Matrices and Scenarios	54
8.2 Adoption Trajectories under the Two Variants	54
8.3 The Planner Pipeline	55
9 Robustness and Sensitivity Analysis	57

9.1 Alternative Market Potential.....	57
9.2 Sensitivity to the Take-off Year	57
9.3 A Direct Adoption Transition in the Markov Model	58
9.4 Sensitivity to the Abandonment Rate	58
9.5 Out-of-sample Validation	58
9.6 Forecast Uncertainty.....	59
10 The Policy Context of AI Diffusion in the European Union and Slovakia.....	61
10.1 Why the policy environment matters for diffusion	61
10.2 The European regulatory framework: the AI Act	61
10.3 The European digital-transformation agenda	62
10.4 The national context and the position of Slovakia	62
11 Discussion	64
12 Extended Interpretation of the Forecasting Results and Their Implications.....	66
12.1 The diffusion regime: interpreting an imitation-dominated process	66
12.2 Two models, one trajectory — and the open question of the ceiling.....	67
12.3 Functional sequencing and the architecture of intra-firm diffusion	68
12.4 The planner pipeline and the calculus of conversion	69
12.5 Implications for management.....	69
12.6 Implications for policy in a small open economy	70
12.7 Implications for the AI–performance link and for future research.....	71
13 Conclusion.....	72
13.1 Key Findings.....	72
13.2 Limitations.....	72
13.3 Recommendations and Implications for Future Work.....	72
References.....	74
Appendix A Annual Adoption Counts by Functional Area	76
Appendix B Numerical Markov Transition Matrices	77
Appendix C Full Annual Bass Forecast Tables.....	78
Appendix D Technical and Reproducibility Notes	79

List of Figures

Figure 1. Distribution of firms across the three AI-use states (N = 1,128).....	33
Figure 2. Year of first AI adoption among adopters (early years before 2020 pooled; n = 741).	34
Figure 3. Number of adopters using AI in each functional area (n = 741 adopters).....	34
Figure 4. Breadth of AI adoption: number of functional areas in which adopters use AI.	35
Figure 5. Co-adoption correlation across functions.	36
Figure 6. Aggregate AI diffusion: observed cumulative adoption and fitted Bass curve with forecast to 2040 (left); observed and modelled annual new adoptions with estimated peak (right).....	37
Figure 7. Marketing & Sales: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).....	39
Figure 8. Production & Operations: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).....	40
Figure 9. R&D and Innovation: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).....	41
Figure 10. Logistics: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).....	42
Figure 11. Corporate Management: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).....	43
Figure 12. Human Resources: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).....	44
Figure 13. Finance: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).....	45
Figure 14. ICT Security: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).....	46
Figure 15. AI adoption by functional area: observed 2025 penetration and forecast penetration for 2030 and 2040 (share of all 1,128 firms).....	47
Figure 16. Markov chain adopter share by scenario and model variant, 2025–2040. Solid lines: absorbing variant; dashed lines: churn variant.	55
Figure 17. State composition over time under the conservative scenario: absorbing variant (left) and churn variant (right).....	56
Figure 18. Aggregate diffusion under alternative market potential (m = 1,128 vs. m = 957). 57	
Figure 19. Aggregate AI diffusion forecast with uncertainty bands derived by parametric bootstrap (4,000 draws). Point forecast with 80% and 95% intervals; the observed 2025 value is the anchor.	60
Figure 20. Triangulating Bass and Markov Adoption Curves (Absorbing vs. Churn, Conservative Scenario), 2025–2040.	64

List of Tables

Table 1. Forecast trajectory of aggregate AI adoption, 2025–2040 ($m = 1,128$; estimated 95% saturation around 2029).....	38
Table 2. Forecast trajectory - Marketing & Sales, 2025–2040 (estimated 95% saturation around 2033).....	40
Table 3. Forecast trajectory — Production & Operations, 2025–2040 (estimated 95% saturation around 2034).	40
Table 4. Forecast trajectory - R&D and Innovation, 2025–2040 (estimated 95% saturation around 2035).	41
Table 5. Forecast trajectory — Logistics, 2025–2040 (estimated 95% saturation around 2036).	42
Table 6. Forecast trajectory - Corporate Management, 2025–2040 (estimated 95% saturation around 2036).	43
Table 7. Forecast trajectory - Human Resources, 2025–2040 (estimated 95% saturation around 2035).....	44
Table 8. Forecast trajectory - Finance, 2025–2040 (estimated 95% saturation around 2037).	45
Table 9. Forecast trajectory - ICT Security, 2025–2040 (estimated 95% saturation around 2038).....	46
Table 10. Estimated Bass parameters by functional area ($m = 1,128$; 95% CIs from the non-linear least squares fit; saturation = first year reaching 95% of m).....	47
Table 11. Bass Diffusion Estimates by Functional Area of AI Use, 2025–2030.....	48
Table 12. Transition matrix - absorbing variant (adopter state permanent; chain converges to full saturation).	54
Table 13. Transition matrix – churn variant (5% annual outflow from adoption; chain converges to a plateau below full saturation).	54
Table 14. Scenario transition probabilities and implied expected time to adoption.	54
Table 15. Markov chain adopter share (% of 1,128) by scenario and variant.	55
Table 16. Sensitivity of the aggregate Bass estimates to the take-off year ($m = 1,128$).....	58
Table 17. Sensitivity of the conservative absorbing scenario to a direct $N \rightarrow A$ transition.	58
Table 18. Long-run adopter plateau (% of 1,128) under alternative abandonment rates, by scenario.....	58
Table 19. Out-of-sample backtest: models estimated on 2020–2024 data, predicting cumulative adoption in 2025 (counts). Mean absolute percentage error across the nine series: 15.1%.	59
Table 20. New adopters per period, overall and by functional area (counts).....	76
Table 21. Transition matrices — Pessimistic scenario (absorbing and churn variants).	77
Table 22. Transition matrices — Conservative scenario (absorbing and churn variants).	77
Table 23. Transition matrices — Optimistic scenario (absorbing and churn variants).....	77
Table 24. Full annual Bass forecast of cumulative penetration (% of 1,128), 2025–2040.	78

List of Abbreviations

Term	Definition
AI	Artificial intelligence
Bass p	Coefficient of innovation (external influence)
Bass q	Coefficient of imitation (internal influence)
m	Market potential (ultimate number of adopters)
N, P, A	Markov states: non-adopter, planner, adopter
a, b	Markov transition probabilities $N \rightarrow P$ and $P \rightarrow A$
ℓ	Annual abandonment rate ($A \rightarrow N$) in the churn variant
NLS	Non-linear least squares
CI	Confidence interval
R²	Coefficient of determination
EU	European Union

1 Introduction: From Comparative Diagnosis to Forward-Looking Analysis

This opening chapter sets out the motivation, scope and structure of the monograph. It introduces artificial intelligence as a driver of enterprise transformation and the diffusion problem at the heart of the study, situates the analysis within the comparative evidence on Slovakia's AI position and the enterprise survey evidence on which the forecasting builds, and states the objectives, contribution and organisation of the work.

1.1 Artificial intelligence, diffusion, and the case of Slovakia

Artificial intelligence (AI) has become one of the central drivers of contemporary economic and technological transformation. Its relevance extends well beyond the technical domain: it increasingly conditions productivity, business-model innovation, labour-market dynamics, public-sector capacity, and the international competitiveness of firms and nations. The assessment of AI in an economy can therefore no longer be treated as a narrow question of digitalisation. It is more accurately understood as part of a broader process of structural adaptation, in which technological readiness, organisational capability, innovation capacity, regulatory support and institutional quality jointly determine whether an economy can absorb and scale advanced technologies.

For a small open economy such as Slovakia, deeply integrated into European production and innovation networks, the speed and breadth with which enterprises adopt AI carry direct implications for the country's structural position within the European economy. AI may strengthen productivity, improve decision-making and raise the competitiveness of firms across sectors; conversely, if adoption remains shallow or concentrated in a narrow group of advanced firms, the country risks a widening gap relative to more dynamic European and global benchmark economies. The stakes of the adoption process are thus not only microeconomic but structural.

Crucially, understanding AI adoption requires more than a static measurement of how many firms currently use the technology. Adoption is a *dynamic* process that unfolds over time and follows characteristic temporal patterns, as the diffusion theory reviewed in the preceding chapter makes clear. Capturing this process calls for diffusion-based modelling, which makes it possible to estimate the mechanisms underlying adoption and to project its future trajectory. This forward-looking, model-based task is the distinctive contribution of the present monograph — but it does not begin from a blank slate. It rests on two earlier analytical foundations: a comparative diagnosis of where Slovakia stands internationally, and a structured measurement of how Slovak enterprises actually adopt AI. The remainder of this introduction draws these foundations together and shows how they frame the forecasting that follows.

1.2 The comparative foundation: Slovakia's position in the European AI landscape

A comparative analysis of Slovakia's AI position examined it through a multi-layered comparative framework that combined harmonised enterprise-level indicators, benchmark-oriented country clustering, and a broad set of international AI- and digital-related indices, including the Global Index on Responsible AI (GIRAI), the Government AI Readiness Index (GAIRI), the Network Readiness Index (NRI), the OECD AI indicators and the Digital Economy and Society Index (DESI). Read together, these layers yield a more comprehensive picture than any single source could provide on its own, and three of their conclusions are directly relevant as a baseline for forecasting.

The first is that Slovakia occupies an intermediate, transitional position rather than that of a clear laggard. The comparative evidence describes a country that has *already entered* a phase in which AI is present in the business and policy environment, yet where this presence remains uneven, only partially embedded, and structurally weaker than in advanced benchmark economies. Slovakia is best characterised not as an economy without AI potential but as one undergoing a *real but incomplete AI transition*, in which local progress coexists with wider systemic constraints. In both the 2020 and the more recent benchmark configurations, the country-level cluster analysis placed Slovakia outside the advanced frontier, in the lower-intensity adoption grouping.

The second relevant conclusion concerns the *internal stratification* of the Slovak AI environment. The comparative analysis documented a consistent gap between a dedicated enterprise survey, which captures comparatively stronger adoption, and official Eurostat-based statistics, which provide a more conservative macro-level estimate. This divergence was visible not only in headline adoption but across firm-size categories and sectors, and in the use of specific technologies such as natural language processing, computer vision and predictive analytics. The two lenses are not contradictory: official statistics describe the enterprise population as a whole, whereas the dedicated survey captures a more engaged, adoption-prone subset. The implication for forecasting is that the Slovak adoption process is not homogeneous; any credible projection must accommodate a stratified population in which a more advanced core coexists with a broader, more cautious majority.

The third conclusion concerns the persistence of structural barriers. The comparative evidence showed that stronger adoption does not dissolve the obstacles to AI; on the contrary, as firms move from awareness toward practical use, barriers related to cost, skills, internal data readiness, legal clarity, infrastructure and trust become *more* salient. The Slovak AI challenge is therefore not merely one of raising adoption rates but of strengthening the absorptive environment in which adoption is sustained and scaled. This finding cautions against forecasts that assume a frictionless approach to full saturation, and it provides a substantive rationale for modelling a long-run ceiling below universal adoption.

1.3 The survey-based foundation: structure and dynamics of adoption in Slovak enterprises

Alongside this international diagnosis, a structured analysis turned inward, characterising the structure and recent dynamics of AI adoption among Slovak enterprises on the basis of dedicated enterprise survey waves conducted in 2025 and 2026 (primarily through computer-assisted web interviewing) and benchmarking selected results against Eurostat data. Its findings supply the empirical micro-structure that the forecasting models operationalise.

A central message of this survey evidence is that AI adoption in Slovak enterprises is best understood as a *layered and uneven process* rather than a binary division between users and non-users. It encompasses actual use, testing and piloting, planning, organisational readiness, business sentiment and perceived barriers. Most consequentially for forecasting, the comparison across survey waves revealed a *redistribution from current use toward planned or preparatory adoption*: the 2026 CAWI data, while showing a somewhat lower share of firms actively using or testing AI than the 2025 wave, contained a particularly large *planning* segment. The correct reading is not that AI has lost strategic relevance, but that a substantial share of enterprises has moved into a preparatory stage — interested in AI, and intending to adopt, but not yet operational. This intermediate “planner” population is precisely the analytical object that the Markov-chain model of the present monograph is designed to represent and whose conversion over time it seeks to project.

The survey evidence also established the *functional heterogeneity* of adoption that motivates the disaggregated modelling strategy adopted here. AI use and testing are not evenly distributed across business functions: in the 2026 CAWI data, AI appeared most prominently in marketing and sales, followed by company management, research and development, finance, human resources, ICT security and production-related activities, with the strongest increases relative to 2025 in marketing and sales and in company management. Logistics, by contrast, declined and remained among the least prominent functions. This pattern — stronger adoption in data-rich, communication- and analysis-intensive functions, weaker adoption in operationally and physically constrained domains — is the empirical justification for estimating diffusion separately for each functional area rather than relying on a single aggregate curve. It is also broadly consistent with the international and grey-literature evidence reviewed in the preceding chapter, in which marketing, sales and product development repeatedly emerge as the leading functions.

Finally, these results sit within the wider policy context of the EU digital-transformation agenda and the international evidence on firm-level adoption growth, emphasising the transition *from experimentation to scaled implementation* and identifying skills, capabilities and organisational readiness as the principal bottlenecks. The benchmarking against Eurostat confirmed that, even where the survey records relatively vigorous adoption, the planning-stage and capability indicators temper any expectation of rapid, frictionless diffusion. These contextual findings reinforce the case, carried forward into the forecasting, for treating the long-run trajectory as conditional on the resolution of capability constraints rather than as a foregone approach to universal use.

1.4 From comparative diagnosis to forecasting

These two strands of evidence together establish *where* Slovakia stands and *how* its enterprises adopt AI; they do not, however, address *where the process is heading*. This is the gap the present monograph fills, and the connection is deliberate at the level of model design rather than mere thematic continuity.

Three bridges link the diagnosis to the forecast. First, the *stratified, below-frontier position* identified in the comparative analysis informs the treatment of market potential and of the long-run ceiling: rather than assuming convergence to full adoption, the analysis examines alternative saturation assumptions and interprets the resulting scenarios against Slovakia's documented structural constraints. Second, the *planner pool* foregrounded in the survey evidence maps directly onto the intermediate state of the three-state Markov model (non-adopter → planner → adopter); the size and conversion rate of this pool, observed in the survey data, become the behavioural parameters that drive the state-transition forecasts. Third, the *functional heterogeneity* documented across both strands of evidence justifies estimating the Bass diffusion model not only for the enterprise sample as a whole but separately for eight functional areas of AI use, producing a differentiated picture of which functions lead the diffusion process and which lag.

In this way the monograph converts a static comparative diagnosis into a dynamic, forward-looking interpretation. It uses this comparative and survey evidence as the empirical and contextual baseline, and applies two complementary diffusion frameworks — the Bass model and the Markov chain — to project the trajectory of AI adoption among Slovak enterprises over a horizon long enough to expose the approach to saturation.

1.5 Objectives, contribution and structure

Against this background, the monograph pursues three interrelated objectives. First, it estimates the diffusion dynamics of AI adoption among Slovak enterprises using the Bass model, both for the sample as a whole and for eight functional areas of AI use. Second, it models the transitions between the non-adopter, planner and adopter states by means of a Markov chain under alternative scenarios and assumptions about the long-run ceiling. Third, it derives forward-looking scenarios of AI penetration over a horizon long enough to expose the approach to saturation, and interprets them in relation to the wider objectives of this study.

The contribution lies in the combination of an aggregate and a functional-area perspective and of two distinct modelling philosophies, anchored in a comparative diagnosis specific to a small, catching-up European economy. Rather than treating AI adoption as a single undifferentiated phenomenon, the analysis identifies which functions lead the diffusion process, which lag behind, how the timing of the adoption wave differs across them, and how the stock of prospective adopters is likely to convert over time. In doing so it connects comparative and survey evidence on Slovakia's AI position with a forward-looking, model-based interpretation of the country's AI trajectory.

The remainder of the monograph is organised as follows. The next chapter sets out the theoretical foundations and reviews the literature on innovation diffusion and firm-level AI

adoption. The methodology chapter then specifies the data, the Bass and Markov models, and the estimation and triangulation strategy. The empirical chapters present, in turn, a descriptive analysis of the sample, the aggregate diffusion results, the functional-area diffusion results, and the state-transition dynamics, followed by a dedicated robustness and sensitivity analysis. A discussion situates the findings within the comparative and international evidence, and the conclusion summarises the key results, acknowledges the limitations, and sets out implications for policy, management and future research.

2 Theoretical Foundations and a Review of the Evidence on AI Adoption and Diffusion

This chapter establishes the conceptual and theoretical foundations on which the analysis rests and reviews the relevant literature. It defines artificial intelligence and AI adoption, sets out the theory of innovation diffusion and the two formal modelling frameworks employed — the Bass model and the Markov chain — and surveys the empirical evidence on firm-level AI adoption, closing by identifying the research gap the monograph addresses.

2.1 Conceptual foundations and the theoretical framework

The diffusion of artificial intelligence (AI) through the enterprise sector is, at its core, a problem of *technology diffusion* — a field with a mature theoretical apparatus in economics, management and the information-systems sciences. Before the empirical models of this monograph are specified, it is therefore necessary to establish the conceptual ground on which they rest: what diffusion means, why it tends to follow regular temporal patterns, which mechanisms generate those patterns, and how the specific characteristics of AI as a general-purpose technology complicate the classical picture. This chapter pursues that task in two movements. The first (Sections 2–3) consolidates and extends the formal diffusion theory that underpins the Bass and Markov models employed later. The second (Sections 4–8) reviews the rapidly expanding empirical and conceptual literature on firm-level AI adoption, organised thematically rather than chronologically, and closes by identifying the research gap that the present work addresses.

A terminological clarification is warranted at the outset, because the word *diffusion* carries two distinct meanings in the contemporary AI literature. In the sense used throughout this monograph, diffusion denotes the **socio-economic spread of a technology across a population of potential adopters** — the tradition of Rogers (2003) and Bass (1969). A second, unrelated meaning has become prominent since 2020, in which *diffusion models* refer to a **class of generative machine-learning architectures** (denoising diffusion probabilistic models) used to synthesise images, text and molecular structures (e.g. Chen, Mei, Fan & Wang, 2024; Wang, He & Peng, 2024). These two literatures share a word but not a subject. The present review concerns exclusively the first; the generative-modelling sense is noted here only to forestall confusion, and the corresponding technical papers are deliberately excluded from the evidentiary base.

2.1.1 Defining artificial intelligence

Any rigorous treatment of AI diffusion must begin by fixing the meaning of its central term, because “artificial intelligence” is used in public, commercial and scholarly discourse with widely varying scope. The reference point adopted here is the internationally harmonised definition of the Organisation for Economic Co-operation and Development, revised in 2024, according to which “an AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions,

content, recommendations, or decisions that can influence physical or virtual environments,” with systems varying in their levels of autonomy and adaptiveness after deployment (OECD, 2024). The 2024 revision is consequential for the present study: the explicit addition of the word “content” to the earlier 2019 formulation was made precisely to bring generative AI — the technology whose diffusion this monograph models — unambiguously within the definition’s scope. The same definition has been given legal force in the European Union, whose Artificial Intelligence Act adopts an aligned wording (European Union, 2024), so that the conceptual and the regulatory boundaries of AI now substantially coincide.

This definition is deliberately broad, encompassing a family of techniques — machine learning, natural-language processing, computer vision and, most recently, generative models — rather than a single technology. For the purposes of this monograph, what matters is not the internal architecture of any particular technique but its character as a general-purpose technology whose economic value is realised only when it is embedded in concrete business functions and complemented by data, skills and organisational change (Brynjolfsson, Rock, & Syverson, 2021). AI is therefore operationalised here at the level of the firm, as the use of AI-based systems within one or more functional areas of the enterprise, and adoption is measured against that operational criterion rather than against any particular technical definition of intelligence.

2.1.2 Defining AI adoption

If AI is the object, *adoption* is the process, and it too requires conceptual precision. Three distinctions structure the analysis. The first is between adoption and *use*: a firm may be classified as an adopter the moment it begins using an AI system in any function, but this binary marker conceals wide variation in how intensively and how widely AI is actually deployed. The second is between the *extensive margin* of adoption — whether the firm uses AI at all — and the *intensive margin* — how much, how well and across how many functions it does so. Recent firm-level evidence stresses that the extensive margin, on which most statistics are based, can substantially misrepresent the economic reality captured by the intensive margin (Cirera, Comin, & Cruz, 2022), a distinction that recurs throughout this monograph and that the maturity literature operationalises as a staged capability process (Hansen, Lillesund, Mikalef, & Altwaijry, 2024). The third distinction is temporal: following the diffusion-of-innovations tradition (Rogers, 2003), adoption is not an instantaneous switch but a staged decision process in which a firm passes from awareness through intention to implementation. This staged conception is why the present study treats *intention to adopt* as a distinct analytical state rather than collapsing it into non-adoption, and it is the conceptual basis for the three-state classification — non-adopter, planner, adopter — used in the modelling.

2.1.3 Determinants of firm-level AI adoption

A substantial and rapidly growing literature seeks to explain *why* some firms adopt AI while others do not, and it provides the conceptual scaffolding for interpreting the diffusion parameters estimated later. The dominant organising framework at the level of the firm is the technology–organisation–environment (TOE) framework (Tornatzky & Fleischer, 1990), which groups the determinants of technology adoption into technological factors (such as

relative advantage and compatibility), organisational factors (such as data resources, skills, absorptive capacity and management readiness) and environmental factors (such as competitive pressure, market demand and government support). Recent systematic reviews of AI adoption increasingly integrate the TOE framework with the diffusion-of-innovations theory and confirm that AI adoption is driven jointly by the perceived attributes of the technology and by the organisational and environmental context, with digital maturity, data readiness, skills and absorptive capacity emerging repeatedly as the binding constraints, and cost, skills shortages, data-governance and trust as the principal barriers (Madan & Ashok, 2023; Hansen et al., 2024). The substantive implication for this monograph is important: the parameters that govern diffusion — the speed at which intentions convert to adoption and the ceiling adoption ultimately reaches — are not technological constants but the aggregate expression of these capability and contextual determinants, and they should be read accordingly.

2.1.4 The scale and recency of enterprise AI adoption

Finally, the conceptual discussion must be anchored in the empirical scale of the phenomenon. Despite intense public attention, enterprise AI adoption remains, at the level of the firm population, comparatively low: official European statistics indicate that only about 13.5% of enterprises in the EU27 — and a similar share across the OECD area — had adopted AI by 2024 (Eurostat, 2024), a figure that rises sharply with firm size but leaves the median firm a non-adopter. At the same time, the period since late 2022 has seen an unprecedented acceleration, driven by the public availability of generative AI through accessible, low-barrier interfaces. This combination of a low base and a steep recent acceleration is the defining empirical feature of contemporary AI diffusion and the principal motivation for modelling it dynamically rather than describing it statically. It also raises the prospect of widening divides — between large and small firms, and between leading and lagging economies — as adoption accelerates unevenly (OECD, 2025), a concern of particular salience for a small, catching-up economy such as Slovakia and one to which the analysis returns. Against this backdrop, the remainder of this chapter develops the theoretical apparatus required to model the diffusion process and reviews the evidence on its firm-level drivers.

2.2 Innovation diffusion theory

The study of how innovations spread through a population of potential adopters has a long tradition in the social and economic sciences. Its foundational framework is Rogers' (2003) theory of the diffusion of innovations, which describes adoption as a process driven by communication among the members of a social system and characterised by a typical S-shaped cumulative adoption curve. Adopters are classified into successive categories — innovators, early adopters, the early and late majority, and laggards — that enter the process at different stages and for different reasons. Each category is distinguished not merely by timing but by a distinct risk posture, resource endowment and exposure to communication channels, so that the shape of the aggregate curve is the macroscopic imprint of a structured, heterogeneous micro-process.

A central insight of diffusion theory is that adoption is shaped by two broad classes of influence. *External influence* operates through mass communication or autonomous decision-making independent of the behaviour of other adopters, while *internal influence* operates through interpersonal communication, social learning and imitation among members of the system. The relative weight of these two mechanisms governs the shape and speed of the diffusion curve and has direct economic counterparts in the analysis of technological adoption (Geroski, 2000; Mahajan, Muller, & Bass, 1990). Geroski (2000) in particular shows that several seemingly competing accounts of diffusion — epidemic models, probit (heterogeneity-threshold) models, and information-cascade models — can each rationalise the S-curve while implying very different policy levers, a reminder that curve-fitting alone does not identify the underlying mechanism.

Rogers' framework also conceives of adoption as a *staged decision process* — knowledge, persuasion, decision, implementation and confirmation — in which a prospective adopter passes through qualitatively distinct cognitive and organisational states before use becomes routine. This staged conception is of more than descriptive interest for the present study: it directly motivates the explicit modelling of an intermediate “planner” state in the Markov-chain approach adopted later in the monograph, in which firms that intend to adopt but have not yet done so are treated as an analytically separate population rather than being collapsed into the non-adopter pool.

2.3 Formal models of diffusion

This section presents the two formal modelling frameworks used throughout the monograph. It first develops the Bass diffusion model as a continuous, aggregate representation of adoption, then the discrete-time Markov chain as a state-transition representation, and finally explains why the two are deployed together as complementary, mutually validating approaches.

2.3.1 The Bass diffusion model

The Bass model (Bass, 1969) provides a parsimonious mathematical formalisation of these two mechanisms. It describes the instantaneous rate of adoption as a function of the remaining market potential and of two behavioural parameters, the coefficient of innovation p (external influence) and the coefficient of imitation q (internal, contagion-type influence). Denoting by $N(t)$ the cumulative number of adopters at time t and by m the market potential, the model is defined by the differential equation:

$$\frac{dN(t)}{dt} = \left(p + \frac{q}{m} N(t) \right) [m - N(t)] \quad (1)$$

The first term, $p[m - N(t)]$, represents adoption arising from external influence and is independent of how many firms have already adopted; the second term, $(q/m)N(t)[m - N(t)]$, represents adoption arising from imitation and is proportional to the interaction between existing adopters and the remaining potential. Integrating the equation with the initial condition $N(0) = 0$ yields the closed-form cumulative adoption curve:

$$F(t) = \frac{1 - e^{-(p+q)t}}{1 + \frac{q}{p} e^{-(p+q)t}}, \quad N(t) = m F(t) \quad (2)$$

The interaction of p and q determines the timing and height of the peak in new adoptions. Differentiating the adoption rate and setting it to zero gives the time of the peak of new adoptions:

$$t^* = \frac{\ln(q/p)}{p+q} \quad (3)$$

The interaction of p and q determines both the timing and the height of that peak. When q exceeds p , the peak occurs at a strictly positive time and the adoption rate traces the familiar bell shape while the cumulative curve follows the S-shape; when q greatly exceeds p , the process starts slowly and then accelerates sharply through imitation, producing a late but steep wave. The ratio q/p is therefore of primary interpretive interest, since it summarises the balance between autonomous and contagion-driven adoption. The Bass model and its extensions have been applied extensively to durable goods, telecommunications and information technologies and remain a standard tool for technological forecasting (Mahajan et al., 1990; Meade & Islam, 2006). In their twenty-five-year review, Meade and Islam (2006) document both the durability of the basic specification and the proliferation of extensions — generalised Bass models incorporating marketing-mix variables, multi-generation models, and spatial formulations — while cautioning that added flexibility frequently fails to improve out-of-sample forecast accuracy. The literature also emphasises that parameter estimation is most reliable once the adoption series has passed its peak; for series still rising, the simultaneous estimation of m , p and q is unstable, a consideration that directly informs the methodological choices adopted in this monograph.

2.3.2 State-based and Markov-chain models of adoption

Whereas the Bass model represents diffusion as an aggregate curve, an alternative tradition models adoption as a sequence of transitions between discrete behavioural states. A discrete-time Markov chain describes a population distributed across a finite set of states, with movement governed by a stochastic matrix $\mathbf{P}$ of transition probabilities and with the future evolution of the system depending only on its current distribution — the Markov property (Norris, 1997). If π_t denotes the row vector of state shares at time t , the system evolves as:

$$\pi_{t+1} = \pi_t \mathbf{P}, \quad \pi_t = \pi_0 \mathbf{P}^t \quad (4)$$

where the t -step transition probabilities follow from the Chapman–Kolmogorov relation as powers of $\mathbf{P}$. For an irreducible, aperiodic chain the distribution converges to a unique stationary vector π satisfying:

$$\pi = \pi \mathbf{P}, \quad \sum_i \pi_i = 1 \quad (5)$$

When one state is absorbing — once entered, never left — the chain instead converges to that state, and the dynamics of interest concern the time to absorption. Partitioning the matrix into its transient block Q and collecting the expected number of visits to each transient state in the fundamental matrix gives the expected time to absorption; for the staged structure used here, this reduces to a simple sum of the reciprocals of the forward transition probabilities:

$$\mathbf{N} = (\mathbf{I} - \mathbf{Q})^{-1}, \quad E[T_{N \rightarrow A}] = \frac{1}{a} + \frac{1}{b} \quad (6)$$

The appeal of the Markov approach in the present context is twofold. First, it represents explicitly the intermediate state of intention — firms that plan to adopt AI but have not yet done so — which the binary logic of the Bass model collapses; because the size of this planner pool directly determines how many firms ultimately adopt and over what period, modelling the transitions into and out of it provides information the aggregate curve cannot. Second, it allows the long-run ceiling to be treated as a modelling choice: an absorbing adopter state implies eventual full saturation, whereas allowing a small outflow ℓ from adoption yields an irreducible chain whose stationary adopter share lies below unity. For the three-state staged structure with abandonment, the stationary adopter share has the closed form:

$$\pi_A = \frac{1}{1 + \frac{\ell}{a} + \frac{\ell}{b}} \quad (7)$$

where a and b are the forward transition probabilities ($N \rightarrow P$ and $P \rightarrow A$) and ℓ is the annual outflow from adoption. This expression makes transparent how the long-run plateau depends jointly on the speed of forward progression and on the abandonment rate.

2.3.3 Why two models? The logic of triangulation

The decision to deploy both an aggregate continuous model and a discrete state-transition model is methodological rather than merely additive. The two frameworks rest on different primitives — a continuous contagion process versus a memoryless transition matrix — and therefore fail in different ways. Where their forecasts agree, confidence in the substantive conclusion is strengthened; where they diverge, the divergence localises the assumption responsible and converts a hidden modelling risk into an explicit, interpretable result. This triangulation logic, increasingly emphasised in the forecasting literature (Meade & Islam, 2006), is a central organising principle of the empirical chapters that follow.

2.4 From IT adoption to AI adoption: what is different

A recurring theme in the most recent literature is that AI adoption, while continuous with decades of research on information-technology (IT) adoption, also strains the inherited conceptual apparatus. A synthesis of information-systems scholarship under the heading *Beyond IT Adoption* argues that the emergence of AI, digital platforms and algorithmic infrastructures has transformed the scale, speed and complexity of innovation processes, raising fundamental questions about what counts as “adoption” in the first place, how value is

distributed among actors, and whose outcomes are measured (panel synthesis, 2025; the authors build on Fichman, 2004, and related IS work). Three differences are repeatedly identified.

First, AI is best understood as a *general-purpose technology* (GPT) whose productive value is realised only in combination with complementary investments in data, skills and organisational redesign. Brynjolfsson, Rock and Syverson (2021) formalise this as a “productivity J-curve”: measured productivity initially *falls* as firms divert resources into the accumulation of intangible complements, before rising once those complements mature. The implication for diffusion analysis is significant — observed adoption may run ahead of measured benefit for an extended period, and the relevant ceiling of adoption is endogenous to the stock of complementary capabilities rather than fixed.

Second, the cost structure of AI adoption is unusually heterogeneous across functions and firm sizes. Cirera, Comin and Cruz (2022), drawing on firm-level evidence from developing economies, show that the binary question “does the firm use the technology?” obscures wide variation in the *intensity* and *sophistication* of use, and that general-purpose technologies diffuse through firms unevenly along a capability gradient. This motivates a treatment of AI adoption that distinguishes functional areas rather than relying on a single aggregate indicator.

Third, the arrival of generative AI from late 2022 introduced a discontinuity in accessibility. Tools delivered through natural-language interfaces lowered the entry barrier dramatically for content-related and customer-facing tasks, compressing into months a diffusion process that earlier enterprise technologies underwent over years. This acceleration, and its uneven incidence across functions, is documented both in scholarly work (Daniotti et al., 2026) and in industry surveys (McKinsey, 2024), discussed below.

2.5 Empirical evidence on firm-level AI adoption

The empirical literature on AI adoption has grown rapidly and now spans several continents, methodologies and units of analysis. Read together, these studies converge on three stylised facts: adoption is *low in the median firm but concentrated in large firms*; it is *highly heterogeneous across sectors and functions*; and it has *accelerated sharply since the diffusion of generative AI*.

The most authoritative large-sample evidence comes from the United States. McElheran et al. (2024), analysing the 2018 Annual Business Survey of roughly 850,000 firms, find that fewer than 6% of firms used any of five measured AI-related technologies, yet employment-weighted adoption exceeded 18%, reflecting the concentration of AI use among very large establishments. AI use clustered with other frontier technologies such as cloud computing and robotics and was disproportionately present in dynamic young firms led by more educated and more experienced owners. The gap between the unweighted and employment-weighted figures is itself a substantive finding: it implies that aggregate “share of firms” statistics can substantially understate the economic exposure of the workforce to AI.

European firm-level studies reinforce the heterogeneity theme while adding a theoretical lens. Garlatti Costa, Venier and Pugliese (2025) examine 237 Italian enterprises through the explicit framework of diffusion-of-innovations theory, mapping firms onto Rogers’ innovator and

early-adopter profiles and confirming the role of AI in enhancing operational efficiency and fostering innovation. Berger and von Garrel (2024) survey 162 decision-makers in German manufacturing and, using the Business Model Canvas as an analytical scaffold, identify eleven distinct “service values” through which AI is integrated into business models, underscoring that adoption is not a single act but a portfolio of functionally specific decisions. At the global scale, Daniotti et al. (2026), publishing in *Science*, document marked heterogeneity in generative-AI adoption across demographics, seniority, sectors and geography, and caution — drawing the historical analogy to steam, electricity and computing — that the macro-level productivity effects of general-purpose technologies have historically taken considerable time to materialise.

Industry surveys, while methodologically distinct from scholarly studies and to be read with appropriate caution, supply timely corroboration. McKinsey’s (2024) global survey reports that regular use of generative AI roughly doubled within ten months to 65% of responding organisations, that overall AI adoption rose to 72%, and that the functions in which generative AI is most often deployed are marketing and sales and product and service development. The functional pattern in these grey-literature sources is broadly consistent with the scholarly evidence and with the functional disaggregation pursued in this monograph, even where the headline adoption levels — reflecting self-selected respondent panels — are higher than those obtained from representative firm censuses.

2.6 Determinants, capabilities and the maturity perspective

A second strand of literature shifts attention from *whether* firms adopt to *why* and *how well*. Here the dominant finding is that AI adoption is less a discrete event than a process of organisational maturation. Hansen, Lillesund, Mikalef and Altwaijry (2024), through a two-phase qualitative case study, develop and validate an *AI Capability Maturity Model* in which diffusion proceeds through identifiable stages of organisational capability rather than as a one-off switch; the model foregrounds data governance, talent and managerial readiness as gating constraints. This maturity framing complements the staged Markov structure used in the present study, since both treat the intermediate, partially-committed firm as a first-class analytical object.

The determinants literature also stresses institutional and cognitive antecedents. Reis & Pinheiro Junior (2025), writing in the *Brazilian Administration Review*, integrate institutional theory with diffusion-of-innovations theory to show that coercive, normative and mimetic *institutional pressures* shape AI adoption alongside the technology’s intrinsic attributes. Madan and Ashok (2023), in a systematic literature review of AI adoption and diffusion in public administration, identify *absorptive capacity* and *data governance* as pivotal enablers and surface a set of recurring ethical “AI tensions” — around fairness, transparency, privacy and human rights — that condition adoption in the public sector and, increasingly, in regulated private-sector contexts. Goel and Nelson (2025), in the *Journal of Technology Transfer*, take a complementary upstream view, constructing search-based indices of public *awareness* of AI and of ChatGPT across U.S. states and showing that awareness — a precondition for adoption — is itself patterned by economic prosperity, with implications for entrepreneurship and regulation.

Taken together, this literature implies that the parameters governing diffusion (the speed of progression through the planner state, the achievable ceiling) are not technological constants but outcomes of capability and institutional context — a point that frames the interpretation of the estimated parameters in later chapters.

2.7 Forecasting AI diffusion: modelling approaches and open questions

The forward-looking literature, to which this monograph contributes directly, remains comparatively thin and methodologically diverse. Three lines of work are salient.

The first applies classical and network-based diffusion dynamics to AI specifically. Roman and Bertolotti (2025), in *Frontiers in Artificial Intelligence*, model the diffusion of technological innovation on networks and show that, under plausible conditions, the process can generate *chaotic boom-bust cycles* rather than smooth monotone saturation — a formal articulation of the recurring “AI winter” phenomenon and a caution against extrapolating early exponential growth indefinitely. The second develops unifying theoretical frameworks: Olfati-Saber (2025), in a preprint that has not undergone peer review and whose quantitative claims should therefore be treated as provisional, proposes a “Universal Dynamics Theory” of technology adoption and reports a critical-mass threshold in the region of 5–15% adoption beyond which transition to rapid growth accelerates, together with productivity gains for AI tools in a wide 15–70% range. The third is cross-country and structural: Andersson Lipcsey (2025) analyses AI diffusion to low- and middle-income countries through a tripartite framework of global value chains, research collaboration and inter-firm knowledge transfers, documenting a substantial but gradually narrowing gap between developed and developing economies.

Practitioner-oriented analyses extend the same questions to the industry level. The industry report *Decoding AI Diffusion* (Hodak, Lakhani, & Cheng, 2025), prepared for a policy-research audience and likewise to be read as grey literature, maps diffusion *enablers* and *bottlenecks* across sectors as varied as high-frequency trading and retail banking, and argues that sector-specific bottlenecks — regulatory, infrastructural and organisational — dominate the realised pace of diffusion. Its central message, that the binding constraint on AI diffusion is increasingly organisational and institutional rather than purely technological, aligns with the maturity and determinants literature reviewed above.

Two open questions emerge from this body of work and orient the present monograph. The first is the *ceiling problem*: what level of long-run saturation is plausible, and should it be assumed or inferred? The second is the *heterogeneity problem*: to what extent does an aggregate diffusion curve mask divergent functional trajectories that carry distinct managerial and policy implications? The literature gestures at both but rarely addresses them jointly with firm-level data for a single national economy.

2.8 Research gap and the positioning of this monograph

The literature reviewed above establishes a mature theory of diffusion, a fast-growing but still largely *descriptive* empirical record of AI adoption, and a nascent and fragmented forecasting literature. Three gaps are evident. First, forward-looking, model-based analyses of *enterprise* AI diffusion remain scarce, and scarcer still for small open economies of the kind that dominate

the European periphery. Second, the pervasive finding of functional heterogeneity is seldom carried through into *functionally disaggregated forecasts*; most projections treat AI adoption as a single undifferentiated series. Third, the two dominant modelling traditions — aggregate Bass-type curves and discrete state-transition models — are rarely deployed *together* on the same data, despite their complementary failure modes.

This monograph addresses all three gaps in the setting of Slovakia, a small, highly open economy deeply integrated into European value chains, for which the speed and breadth of enterprise AI adoption carry direct implications for structural competitiveness. It applies the Bass model both to the full sample of Slovak enterprises and separately to eight functional areas of AI use; it models transitions between the non-adopter, planner and adopter states by means of a Markov chain under alternative assumptions about the long-run ceiling; and it derives forward-looking scenarios over a horizon long enough to expose the approach to saturation. In so doing it connects comparative cross-country evidence on Slovakia's AI position — where the country lies below the European average — with a forward-looking, dual-model interpretation of the country's AI trajectory. The result is intended as a contribution both to the substantive understanding of AI diffusion in a catching-up European economy and to the methodology of forecasting technology diffusion under genuine uncertainty about its limits.

3 Methodology, Data and Sample: Extended Specification

This chapter specifies the data and the analytical methods in the detail required for reproducibility. It describes the research design and the two-model strategy, the provenance and construction of the dataset and its variables, the specification and estimation of the Bass and Markov models together with the treatment of uncertainty, and the principal methodological limitations that condition the interpretation of the results.

3.1 Research design and the logic of a two-model strategy

The empirical strategy of this monograph rests on the joint application of two structurally distinct diffusion models — the Bass model and a discrete-time Markov chain — to a single, consistently constructed enterprise-level dataset. The design is deliberately not a sequence of independent exercises but a *triangulation*: the two models share the same data basis, the same time frame and the same observed anchor (the 2025 adopter share), so that their outputs can be read against one another as mutually validating or, where they diverge, as diagnostic of the assumptions responsible.

The rationale for combining the two frameworks is methodological rather than additive, and it is worth stating precisely because it governs how the results should be interpreted. The Bass model represents diffusion as a continuous aggregate curve and is well suited to characterising the *mechanism* (the balance of innovation and imitation) and the *timing* (the peak of new adoptions and the approach to saturation) of the process. It is, however, silent on the behavioural states through which firms pass and treats adoption as a binary outcome. The Markov chain, by contrast, represents diffusion as a set of transitions between discrete states and is well suited to making the *intermediate state of intention* explicit and to treating the long-run ceiling as a modelling choice rather than an assumption fixed in advance; it is, however, agnostic about the continuous shape of the adoption curve and, in the present application, relies on assumed rather than estimated transition probabilities. The two models therefore fail in different ways and along different dimensions. Where their forecasts coincide, the shared conclusion is unlikely to be an artefact of either functional form; where they diverge, the divergence localises the responsible assumption and converts a hidden modelling risk into an explicit, interpretable finding. This complementarity is the organising principle of the analytical design and the reason the two models are reported in parallel throughout.

The analytical sequence proceeds in four stages. A descriptive analysis first establishes the distribution of firms across adoption states, the timing of adoption and the functional composition of use. The Bass model is then estimated, first for the full sample and subsequently for each of eight functional areas, yielding an aggregate and a disaggregated characterisation of the diffusion mechanism. The Markov chain is then applied to the three-state classification under three scenarios and two structural variants. A dedicated robustness and sensitivity analysis finally probes the dependence of the results on the principal modelling choices. Each stage is designed to be reproducible from the same prepared dataset, in keeping with the open-science commitments of this study.

3.2 Data provenance and the survey instrument

The analysis draws on a dedicated enterprise-level dataset of 1,128 Slovak firms. These data were collected over an eight-month window between July 2025 and February 2026 through two consecutive computer-assisted web interviewing (CAWI) survey waves, the samples of which were pooled into a single analytical cross-section in which each enterprise contributes one record. Because both waves employed the same web-based (CAWI) mode and the same core instrument, their combination into a single forecasting dataset is methodologically consistent and avoids the mode effects that a mixed-mode pooling could introduce. A separate computer-assisted telephone interviewing (CATI) wave, carried out within the wider survey programme, was used for comparative benchmarking against external statistics but does not form part of the forecasting dataset analysed here. Each record captures the firm's AI-use status at the time of the survey together with the retrospective calendar year of first AI adoption, which provides the temporal basis for the diffusion series; the survey design and sampling have been documented in earlier comparative analyses, and this monograph takes the harmonised, pooled CAWI dataset as its point of departure and concentrates on its forward-looking modelling rather than re-documenting the instrument.

Because the forecasting models are estimated on this consolidated cross-section, two features of the data provenance bear directly on interpretation and are stated here explicitly. First, the dataset is a *single cross-section* enriched with a retrospective adoption date for each adopter; it is not a panel in which the same firms are observed transitioning between states over time. This has a specific methodological consequence for the Markov model, addressed in Section 5: the transition probabilities cannot be estimated from observed state-to-state movements and must instead be specified by scenario. Second, the unit of observation is the enterprise, and the variables describe the firm's AI-use status at the time of the survey together with the year in which adoption first occurred, which together provide the temporal basis for constructing the diffusion time series.

3.3 Variable construction and the three-state classification

The central variable is the firm's AI-use status, recoded into three mutually exclusive and ordered categories that constitute the common backbone of both models. *Adopters* are firms already using AI in at least one functional area (741 firms, 65.7% of the sample). *Planners* are firms not yet using AI but intending to adopt it (216 firms, 19.1%). *Non-adopters* are firms neither using nor planning to use AI (171 firms, 15.2%). This three-way classification is more informative than the binary user/non-user distinction conventional in adoption studies precisely because it isolates the planner population — the reservoir of latent adoption whose conversion, as the modelling shows, is the principal near-term driver of the trajectory. The ordering of the categories (non-adopter → planner → adopter) reflects an assumed progression of commitment and supplies the state space of the Markov chain.

For each adopter the dataset records the calendar year of first AI adoption, which provides the temporal raw material for the diffusion series. Inspection of this variable confirms the structural features that motivate the modelling choices. Adoption is negligible before 2020, with only sparse earlier records spanning back to 2000; it then accelerates sharply, the annual count of

new adopters rising through the early 2020s to an observed maximum in 2025. This concentration — with the overwhelming majority of adoption dates falling in 2023, 2024 and 2025 — is the empirical signature of the generative-AI wave and the direct justification for specifying the diffusion process from a 2020 take-off rather than fitting it to the full historical record (Section 4).

The functional composition of use is captured by eight binary indicators, one for each functional area in which an adopter may apply AI: marketing and sales; production and operations; corporate management; human resources; logistics; ICT security; finance; and research, development and innovation. These eight indicators serve two purposes. They define the eight series on which the disaggregated Bass models are estimated, and they support the descriptive analysis of the *breadth* of adoption (the number of functions in which a firm uses AI) and of *co-adoption* (the pairwise correlation of use across functions). It should be noted that the adoption year is recorded at the level of the firm rather than separately for each function; consequently, when the diffusion model is estimated by functional area, the firm-level adoption year is attributed to every function in which the firm reports using AI. The functional-area curves therefore differ principally in their level and composition, while the within-function timing of adoption is approximated rather than directly observed — a limitation made explicit in Section 7.

3.4 Bass model: specification, estimation and uncertainty

The Bass model was estimated in its cumulative form by non-linear least squares, fitting the closed-form cumulative adoption curve to the observed cumulative adoption series. Three modelling decisions, applied uniformly across the aggregate analysis and all eight functional-area analyses, require explicit statement and justification.

The first is the treatment of *market potential*. The market potential m was fixed exogenously at the sample size, $m = 1,128$, rather than estimated jointly with the behavioural parameters. Two considerations motivate this choice. Substantively, in the absence of a defensible external estimate of the eventual adopter population, the sample provides the only consistent and reproducible reference frame, and applying the same m to every series ensures that the aggregate and functional analyses are mutually comparable. Statistically, the simultaneous estimation of m together with the innovation and imitation coefficients is notoriously unstable for series that have not yet passed their peak, a difficulty long recognised in the diffusion-modelling literature (Meade & Islam, 2006); fixing m removes this instability at the cost of making the forecasts explicitly conditional on the chosen ceiling. The sensitivity of the results to this assumption is examined directly in the robustness analysis, where an alternative market potential is considered.

The second is the *take-off specification*. Because adoption in the sample is dominated by the recent wave, the diffusion process was modelled from a take-off point in 2020, with the sparse earlier adoptions (2000–2019) entering as an initial installed base rather than as part of the estimation window. This decision is not a matter of convenience but of estimability: a naive specification fitted over the full historical series was found to be degenerate, collapsing the innovation coefficient to its lower bound and understating the recent acceleration that is the

central feature of the data. The take-off specification thus recovers a process consistent with the observed structural break.

The third concerns the *quantities derived* from each fit. For every series the estimation yields point estimates of the innovation coefficient p and the imitation coefficient q , from which the coefficient of determination (R^2), the timing of the peak of new adoptions, the year of 95% saturation and the full cumulative forecast trajectory are computed. Approximate 95% confidence intervals for p and q were obtained from the covariance matrix of the fit using a t -distribution with the appropriate degrees of freedom. Given the brevity of the series these intervals are necessarily wide, and they are reported not as a formality but as a transparent statement of the genuine estimation uncertainty.

Beyond the parameter intervals, the uncertainty in the *aggregate forecast* was characterised by parametric bootstrap. Four thousand parameter pairs were drawn from the sampling distribution of the fitted coefficients, using the estimated covariance matrix, and a forecast trajectory was computed for each draw; the resulting distribution of trajectories yields the reported 80% and 95% forecast bands and, by extension, an interval for the year of 95% saturation. This procedure propagates parameter uncertainty into the forecast in a principled way, while remaining — as Section 7 makes clear — silent on the deeper structural uncertainty about the market potential itself, which the scenario analysis addresses separately.

3.5 Markov chain: state space, transition structure and scenarios

The Markov chain represents the enterprise sample as a population distributed across the three ordered states defined in Section 3 — non-adopter (N), planner (P) and adopter (A) — with the observed counts at the start of the horizon (171, 216 and 741 respectively) serving as the initial distribution. Movement between states over discrete annual steps is governed by a 3×3 row-stochastic transition matrix, and the future distribution evolves by repeated multiplication of the state vector by this matrix.

The transition structure embodies two behavioural assumptions that reflect the substantive logic of adoption. Progression is *upward and single-step*: a non-adopter must first become a planner before becoming an adopter, with no direct $N \rightarrow A$ transition, on the reasoning that operational adoption is preceded by an interval of intention. Forward progression is *not reversed* at the planning stage: planners do not revert to non-adoption, since the formation of an intention is treated as a durable change of state. These assumptions reduce the free parameters of the matrix to two forward probabilities — a ($N \rightarrow P$) and b ($P \rightarrow A$) — together with the treatment of the adopter state, which distinguishes the two structural variants below.

Two *structural variants* of the adopter state are examined in parallel as bounds on the long-run ceiling. In the *absorbing* variant the adopter state is permanent ($A \rightarrow A = 100\%$), so the chain converges to full saturation, consistent with the fixed market potential of the Bass model. In the *churn* variant a small annual outflow from adoption is permitted ($A \rightarrow N = 5\%$, $A \rightarrow A = 95\%$) to represent the modest abandonment that may occur in practice; this renders the chain irreducible rather than absorbing, so that adoption converges to a stationary plateau strictly below full saturation. Reporting the two variants side by side makes the dependence of the long-

run outcome on the reversibility of adoption explicit rather than hidden, and the sensitivity to the precise abandonment rate is examined in the robustness analysis.

Within each variant the two free probabilities are set by *scenario* rather than estimated, for the reason given in Section 2: the single available cross-section does not record observed state-to-state movements, so the transitions cannot be estimated directly. Three scenarios are specified — conservative (a central, plausible pace), optimistic (a faster pace) and pessimistic (a slower pace) — and in every case the constraint $a < b$ is imposed, on the reasoning that the residual pool of non-adopters is a more resistant segment that has not yet formed an intention, whereas planners have already signalled intent and convert more readily. For each scenario the implied expected time for a non-adopter to progress through to adoption is summarised as the sum of the reciprocals of the two forward probabilities ($1/a + 1/b$), providing an intuitive behavioural reading of each parameter set. Because the probabilities are assumed, the Markov results are properly understood as structured *what-if* analyses that bracket the plausible range, not as point predictions; their value lies in the transparency with which they map assumptions onto outcomes.

3.6 Reproducibility, software and triangulation

The estimation was carried out in a scripted, fully reproducible analytical environment, in keeping with the open-science and reproducibility commitments of this study: the same prepared adoption series, the same estimation routines and the same random seed for the bootstrap reproduce the reported results exactly. To guard against dependence on a single implementation, the prepared series can additionally be processed in an independent software implementation, whose Bass forecasting module provides a separate estimation route; where the two implementations yield comparable parameters, the convergence is treated as a validation through triangulation at the level of computation. This computational triangulation complements the methodological triangulation between the Bass and Markov models described in Section 1, so that the analysis is corroborated both across methods and across implementations.

3.7 Methodological limitations and their bearing on interpretation

The credibility of a forecasting exercise depends on the candid statement of its limitations, and several follow from the data and the modelling choices. They are grouped here by the dimension of the results they affect.

The most consequential limitation concerns the *representativeness of the sample*, which bears on the absolute levels of the forecasts. Within the sample, 65.7% of firms are already adopters — a share high relative to economy-wide benchmarks of enterprise AI adoption, which for the corresponding period place the proportion of AI-using enterprises substantially lower in official Eurostat and OECD statistics. This divergence is consistent with a sample skewed toward digitally engaged firms, which are more likely both to participate in a survey of this kind and to have already adopted AI. Two consequences follow. The fixed market potential $m = 1,128$ and the resulting saturation forecasts should be read as the saturation dynamics *of this sample of predominantly digitally active firms*, not as a projection across the full population of Slovak

enterprises; absolute penetration levels therefore likely overstate population-wide adoption. By contrast, the *relative and temporal* patterns — the dominance of imitation over innovation, the ordering of functional areas, the timing of the diffusion wave and the role of the planner pool — describe the *shape* of the process rather than its absolute reach and are considerably more robust to this skew. The comparative positioning of Slovakia against international benchmarks, which addresses the level question directly, is examined separately in comparative analyses, to which this monograph is a forward-looking complement.

A second group of limitations concerns the *Bass specification*. The assumption $m = 1,128$ implies that the curves tend toward full saturation of the sample, so the forecasts are conditional on this explicit ceiling rather than unconditional. The attribution of a firm-level adoption year to each functional area means that within-function timing is approximate, so the ordering of functions is more robust than the precise timing within any one of them. The annual series is short and, for most functions, has only recently passed its peak, which widens the confidence intervals. Finally, the high R^2 values reflect the fitting of monotonic cumulative curves and should be read as a goodness-of-fit check rather than as proof of dynamic forecasting accuracy; the out-of-sample exercise in the robustness analysis provides the more demanding test.

A third group concerns the *Markov specification*. The transition probabilities are assumed rather than estimated, so the scenarios are structured what-if analyses bounded by the plausible parameter range rather than empirical point forecasts. The single-step, upward-only structure is a deliberate simplification, although the robustness analysis shows the central results to be insensitive to relaxing it. The absorbing-versus-churn dichotomy brackets the long-run ceiling but does not, on the present cross-section, allow the abandonment rate to be calibrated from data — a gap that successive survey waves are expected to close.

Taken together, these limitations define the proper register in which the forecasts should be read: as *conditional, shape-robust projections* whose qualitative conclusions are well supported and whose absolute levels are deliberately tied to an explicit, stated ceiling. This is the appropriate standard for medium-term technology forecasting on a short series, and it is the standard against which the robustness analysis evaluates the results.

4 Descriptive Analysis of the Sample

Before turning to the diffusion models, this chapter characterises the sample descriptively: the distribution of firms across the three AI-use states, the timing of adoption, the relative frequency of AI use across functional areas, and the breadth and co-occurrence of AI use within firms. These features motivate and contextualise the modelling that follows.

4.1 AI-Use Status and Adoption Timing

Of the 1,128 enterprises in the sample, 741 (65.7%) are adopters, 216 (19.1%) are planners and 171 (15.2%) are non-adopters. Nearly two thirds of firms therefore already use AI in at least one area, while a further fifth intend to adopt – a stock of stated intentions that, as the Markov analysis shows, is the principal near-term driver of future adoption. The distribution across states is shown in Figure 1.




Figure 1. Distribution of firms across the three AI-use states ($N = 1,128$).

The timing of adoption among the 741 adopters is strongly concentrated in the most recent years. As Figure 2 shows, adoption was negligible before 2020 and then accelerated sharply, with the annual number of new adopters rising in every year from 2021 and reaching its observed maximum in 2025. This concentration is the empirical signature of the generative-AI wave and is precisely the pattern that motivates the take-off specification of the Bass model.

Two features of this distribution carry forward into the modelling. The first is the size of the planner pool: with 216 firms (19.1%) already intending to adopt, the sample contains a substantial reservoir of latent adoption that has not yet materialised, and it is the conversion of this pool – rather than the recruitment of committed non-adopters – that drives the near-term forecasts of both models. The second is the shape of the adoption timeline. The near-absence of adoption before 2020 followed by a steep, sustained rise does not describe a gradual diffusion that began years earlier; it describes a structural break concentrated in the post-2022 period, coinciding with the public emergence of generative AI. This is why the diffusion process is

specified from a 2020 take-off rather than fitted to the full historical record, and why the estimated process is dominated by imitation: the wave reflects firms responding to a rapidly spreading, highly visible technology rather than to independent, long-gestating adoption decisions.




Figure 2. Year of first AI adoption among adopters (early years before 2020 pooled; n = 741).

The eight functional areas differ markedly in how often firms apply AI within them. Marketing and sales is by far the most common (325 adopters), followed by production and operations (204) and R&D and innovation (181), while finance (115) and ICT security (114) are the least common. The full ordering is shown in Figure 3. Because the same fixed market potential is applied to every area in the diffusion model, these differing levels translate directly into the differing current penetration rates analysed in Chapter 6.




Figure 3. Number of adopters using AI in each functional area (n = 741 adopters).

This ordering is not arbitrary. The functions at the top of the distribution are those in which AI can be deployed with the lowest organisational and technical commitment and with the most immediately visible return – generative tools for content, customer communication and sales analytics require little integration with core systems and produce quick, observable results, which both lowers the adoption threshold and amplifies the peer imitation that the diffusion

model identifies as the dominant mechanism. The functions at the bottom of the distribution, finance and ICT security, are precisely those where adoption is gated by specialised capability, data sensitivity and regulatory caution, so that diffusion proceeds more slowly and depends less on simple imitation. The descriptive ranking therefore already prefigures the staggered diffusion timing estimated in Chapter 6.

4.2 Breadth and Co-adoption of AI

Firms differ not only in whether they use AI but in how broadly. Among adopters, the mean number of functional areas in which AI is used is 1.89, with the modal firm using AI in a single area; a smaller group of firms deploys AI across many functions simultaneously, up to all eight. The distribution of this breadth is shown in Figure 4, and indicates that AI adoption in most firms is still concentrated in one or two functions rather than spread across the enterprise.




Figure 4. Breadth of AI adoption: number of functional areas in which adopters use AI.

The functional areas also tend to be adopted in characteristic combinations. Figure 5 reports the pairwise co-adoption correlations among the eight indicators. The associations are uniformly positive – firms that use AI in one area are more likely to use it in others – but moderate in magnitude, the strongest being between R&D and ICT security and between corporate management and finance. These clusters suggest two recognisable patterns: a technologically advanced group adopting AI across specialised technical functions, and an administrative group digitalising control and finance functions together. This co-adoption structure complements the diffusion analysis by indicating that the spread of AI across functions is correlated rather than independent.




Figure 5. Co-adoption correlation across functions.

Read together, the breadth and co-adoption patterns suggest that AI has so far spread broadly but not deeply. Most adopters use AI in only one or two functions, so that the high headline adoption rate reflects many firms taking a first, narrow step rather than a smaller number embedding AI across the enterprise. Where AI does extend across functions, it does so along two recognisable pathways: a technologically advanced archetype that combines AI in R&D and ICT security functions requiring specialised technical capability and an administrative archetype that combines AI in corporate management and finance – the control and reporting core of the firm. These archetypes indicate that within-firm diffusion is structured rather than random, following the lines of a firm's existing technical or administrative orientation. Taken as a whole, the descriptive picture a large pool of stated intentions, an adoption wave concentrated in the most recent years, an ordered hierarchy of functions and a still-shallow average breadth sets up the two questions the remainder of this monograph addresses through formal modelling: how fast the process will continue, and how far it will ultimately reach.

5 Aggregate AI Diffusion

This chapter presents the Bass diffusion model estimated for the sample as a whole, before the analysis is disaggregated by functional area in Chapter 6. The forecast horizon is extended to 2040 so that the levelling-off of the S-curve and the approach to saturation are visible.




Figure 6. Aggregate AI diffusion: observed cumulative adoption and fitted Bass curve with forecast to 2040 (left); observed and modelled annual new adoptions with estimated peak (right).

The aggregate model fits the data closely ($R^2 = 0.966$) and yields an innovation coefficient of $p = 0.0153$ and an imitation coefficient of $q = 0.721$. The imitation coefficient exceeds the innovation coefficient by a factor of roughly fifty, indicating that the diffusion of AI across the sample is overwhelmingly driven by imitation and contagion among firms rather than by autonomous, externally motivated adoption. This is the characteristic signature of a technology that spreads rapidly through peer influence once an initial base of adopters has formed, and it is fully consistent with the acceleration of AI uptake following the emergence of generative AI from 2022 onward.

The economic significance of this imitation-dominated pattern goes beyond the shape of the curve. A diffusion process in which contagion vastly outweighs autonomous adoption is characteristic of a general-purpose technology whose value to each firm rises as adoption spreads and as complementary skills and tools become widely available (Brynjolfsson, Rock, & Syverson, 2021). Two implications follow. On the one hand, such a process supports rapid catch-up: once a critical mass of visible adopters has formed, the remaining firms convert quickly, which is consistent with the steep projected rise over the second half of the 2020s. On the other hand, it implies a dependence on that critical mass actually forming and remaining visible; because the process is driven far more by peer example than by independent initiative, it is sensitive to whether early adoption is observable and seen to succeed. This dual character — fast diffusion conditional on visible momentum — is the empirical basis for the policy reading developed in Chapter 9, which favours measures that make successful adoption visible over measures aimed at stimulating isolated first-mover adoption.

By 2025, 65.7% of all firms had adopted AI in at least one area. The model places the peak of new aggregate adoptions around 2024 and the 95% saturation point around 2029. Under the

fixed market-potential assumption, projected penetration rises to 98.6% by 2030 and effectively reaches the ceiling thereafter. The forecast trajectory is summarised in Table 1.

Table 1. Forecast trajectory of aggregate AI adoption, 2025–2040 ($m = 1,128$; estimated 95% saturation around 2029).

Year	Cumulative adopters	Penetration (% of 1128)
2025 (observed)	741	65.7%
2027	995	88.2%
2030	1112	98.6%
2035	1128	100.0%
2040	1128	100.0%

6 AI Diffusion Across Functional Areas

This chapter applies the Bass model separately to each of the eight functional areas of AI use, presented in descending order of their 2025 adoption level. Each section reports the estimated diffusion profile, interprets the underlying mechanism and applications, and presents the forecast trajectory to 2040. A cross-functional synthesis follows in Section 6.9.

6.1 Marketing & Sales




Figure 7. Marketing & Sales: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).

Marketing and sales is the single largest field of AI application in the sample and, at the same time, the functional area in which diffusion is most advanced. By 2025, 28.8% of all firms reported using AI in this area, well ahead of any other function. The estimated imitation coefficient ($q = 0.50$) is roughly fifty times larger than the innovation coefficient ($p = 0.010$), indicating a diffusion process driven almost entirely by imitation, competitive pressure and word-of-mouth among firms rather than by autonomous, externally motivated adoption.

The model places the peak of new adoptions around 2027 and the 95% saturation point around 2033. Although marketing leads in absolute penetration, its own diffusion wave has not yet crested, so the function is expected to remain the most dynamic area of AI uptake in the near term, with projected penetration approaching 84% of all firms by 2030 and effectively reaching the ceiling by the mid-2030s.

The applications driving this uptake are predominantly customer-facing: generative content production, conversational customer service and chatbots, automated lead qualification and segmentation, and campaign analytics, all of which have comparatively low entry barriers. From a fit-diagnostic perspective the model tracks the observed series closely across all post-take-off years; because the function remains ahead of its peak, however, the forecast is an extrapolation beyond the most informative part of the curve, and the relatively wide confidence interval on q should be borne in mind.

Table 2. Forecast trajectory - Marketing & Sales, 2025–2040 (estimated 95% saturation around 2033).

Year	Cumulative adopters	Penetration (% of 1128)
2025 (observed)	325	28.8%
2027	598	53.0%
2030	949	84.1%
2035	1112	98.6%
2040	1127	99.9%

6.2 Production & Operations




Figure 8. Production & Operations: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).

Production and operations is the second most penetrated functional area (18.1% by 2025). The diffusion profile is again strongly imitation-dominated ($q = 0.49$, $p = 0.006$), but the lower innovation coefficient and the later estimated peak (around 2028) reflect the more capital-intensive and process-embedded nature of AI in this domain, where adoption typically requires integration with existing operational technology, data infrastructure and quality-control systems.

Projected penetration reaches roughly 73% of firms by 2030 and approaches saturation around 2034. The trajectory suggests that operational AI is following the customer-facing wave with a lag of approximately one to two years, consistent with the longer implementation cycles typically observed for production-side technologies.

Representative applications include predictive maintenance, automated quality inspection, process optimisation and demand-driven scheduling. The model fit is very strong ($R^2 = 0.98$), and the observed series already shows the early acceleration of the imitation phase, lending additional confidence to the near-term part of the forecast.

Table 3. Forecast trajectory — Production & Operations, 2025–2040 (estimated 95% saturation around 2034).

Year	Cumulative adopters	Penetration (% of 1128)
2025 (observed)	204	18.1%
2027	424	37.5%
2030	821	72.7%
2035	1093	96.9%
2040	1125	99.7%

6.3 R&D and Innovation




Figure 9. R&D and Innovation: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).

Research, development and innovation displays a comparatively high adoption level (16.0% by 2025) and the tightest model fit in the sample ($R^2 = 0.99$). The estimated imitation coefficient ($q = 0.48$) confirms the same contagion-driven dynamic observed elsewhere, with the peak of new adoptions expected around 2028 and saturation around 2035.

The relatively early position of R&D in the adoption ranking is consistent with the knowledge-intensive character of innovation activities, where experimentation with AI tools tends to emerge early among more technologically capable firms. Projected penetration approaches 70% by 2030.

Typical applications span AI-assisted design and simulation, literature and patent analysis, and data-driven discovery. As shown in the co-adoption analysis (Chapter 4), R&D use is most strongly associated with ICT security use, suggesting a cluster of technologically advanced firms that adopt AI across several specialised functions simultaneously.

Table 4. Forecast trajectory - R&D and Innovation, 2025–2040 (estimated 95% saturation around 2035).

Year	Cumulative adopters	Penetration (% of 1128)
2025 (observed)	181	16.0%
2027	394	34.9%
2030	793	70.3%
2035	1088	96.5%
2040	1125	99.7%

6.4 Logistics




Figure 10. Logistics: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).

Logistics occupies a middle position (14.4% by 2025), with a diffusion profile ($q = 0.45$, $p = 0.005$) and an estimated peak around 2028–2029 that place it slightly behind the operational core. Projected penetration reaches approximately 64% by 2030 and saturation around 2036.

Table 5. Forecast trajectory — Logistics, 2025–2040 (estimated 95% saturation around 2036).

Year	Cumulative adopters	Penetration (% of 1128)
2025 (observed)	162	14.4%
2027	346	30.7%
2030	721	63.9%
2035	1066	94.5%
2040	1121	99.4%

The pattern is consistent with the use of AI for routing, demand forecasting, warehouse optimisation and inventory management. As with production, adoption depends on integration with established logistics and enterprise-resource systems, which moderates the pace of uptake relative to customer-facing applications.

The estimated coefficients are typical of the middle tier of functions: a small innovation coefficient and a moderate imitation coefficient, producing a steady, peer-driven diffusion that has only recently entered its acceleration phase.

6.5 Corporate Management




Figure 11. Corporate Management: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).

AI in corporate management – decision support, analytics-based planning and reporting – reaches 13.9% of firms by 2025. The estimated coefficients ($q = 0.44$, $p = 0.005$) and the later peak (around 2029) indicate a diffusion process that is still in an earlier phase than the customer- and operations-facing functions, with saturation projected around 2036.

Projected penetration approaches 60% by 2030, the model fit being somewhat weaker here ($R^2 = 0.95$) than in other areas. This is plausibly related to the more diffuse and heterogeneous ways in which AI enters general management, which makes the timing of adoption harder to pin down from a single firm-level adoption year.

Table 6. Forecast trajectory - Corporate Management, 2025–2040 (estimated 95% saturation around 2036).

Year	Cumulative adopters	Penetration (% of 1128)
2025 (observed)	157	13.9%
2027	318	28.2%
2030	682	60.4%
2035	1054	93.4%
2040	1120	99.3%

The co-adoption analysis shows management AI use to be most strongly associated with finance, pointing to a cluster of administrative and control functions that tend to be digitalised together.

6.6 Human Resources




Figure 12. Human Resources: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).

Human resources reaches 12.9% of firms by 2025, yet it exhibits one of the highest imitation coefficients in the sample ($q = 0.51$) alongside a very low innovation coefficient ($p = 0.004$). This combination implies a function in which autonomous early adoption is rare, but contagion is strong once peer firms begin to adopt – a classic late-but-steep diffusion signature.

The estimated peak lies around 2029 and saturation around 2035, with projected penetration reaching roughly 66% by 2030.

Table 7. Forecast trajectory - Human Resources, 2025–2040 (estimated 95% saturation around 2035).

Year	Cumulative adopters	Penetration (% of 1128)
2025 (observed)	146	12.9%
2027	333	29.5%
2030	747	66.2%
2035	1085	96.2%
2040	1125	99.7%

The profile is consistent with the spread of AI in recruitment screening, candidate communication and administrative HR tasks, where adoption tends to follow visible practice in comparable organisations rather than independent initiative. The high imitation coefficient implies that, once a critical mass of peer adoption is reached, HR could converge relatively quickly despite its current low base.

6.7 Finance




Figure 13. Finance: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).

AI in finance reaches 10.2% of firms by 2025 and is among the latest-developing functions, with an estimated peak around 2030, saturation around 2037 and projected penetration of approximately 54% by 2030. The low innovation coefficient ($p = 0.004$) and moderate imitation coefficient ($q = 0.45$) describe a cautious, follower-type adoption pattern.

This is consistent with the comparatively conservative and regulated character of financial functions, where applications such as fraud detection, forecasting and the automation of financial workflows are typically introduced once reliability and compliance considerations have been addressed.

The slower diffusion places finance among the laggard functions whose adoption gap relative to the leaders will persist for much of the coming decade before closing in the late 2030s.

Table 8. Forecast trajectory - Finance, 2025–2040 (estimated 95% saturation around 2037).

Year	Cumulative adopters	Penetration (% of 1128)
2025 (observed)	115	10.2%
2027	252	22.4%
2030	603	53.5%
2035	1036	91.8%
2040	1118	99.1%

6.8 ICT Security




Figure 14. ICT Security: observed and fitted cumulative adoption with forecast to 2040 (left); observed and modelled annual new adoptions (right).

ICT security is, together with finance, the least penetrated functional area (10.1% by 2025) and exhibits the lowest imitation coefficient in the sample ($q = 0.41$). Its estimated peak lies around 2030 and saturation only around 2038 – the latest of all functions – with projected penetration reaching approximately 49% by 2030.

The slower contagion dynamic is consistent with the specialised and capability-dependent nature of AI in cybersecurity, where adoption depends on specific technical expertise that diffuses less readily through simple peer imitation than customer-facing applications.

Notwithstanding its low base, ICT security shows the strongest co-adoption association with R&D, again indicating that its adopters are concentrated among technologically sophisticated firms rather than spread evenly across the population.

Table 9. Forecast trajectory - ICT Security, 2025–2040 (estimated 95% saturation around 2038).

Year	Cumulative adopters	Penetration (% of 1128)
2025 (observed)	114	10.1%
2027	240	21.3%
2030	556	49.3%
2035	1001	88.8%
2040	1111	98.5%

6.9 Cross-Functional Comparison

Taken together, the eight functional analyses reveal a consistent and ordered pattern. Across every function the imitation coefficient dominates the innovation coefficient, confirming at the disaggregated level the contagion-driven dynamic identified for the sample as a whole. The functions differ, however, in their current penetration and in the timing of their adoption wave, and a clear regularity emerges: the more penetrated a function is today, the earlier its estimated peak and the earlier its saturation. Marketing and sales, the front-runner, approaches saturation in the early 2030s, whereas the least penetrated functions – finance and ICT security – are projected to saturate only in the late 2030s.

The confidence intervals are wide, reflecting the short annual series on which the estimates rest. While the point estimates and the ordering of the functions are robust and mutually consistent, the precise values of the coefficients – and therefore the exact timing of the projected peaks and saturation years – carry a non-trivial margin of uncertainty, examined further in Chapter 8.




Figure 15. AI adoption by functional area: observed 2025 penetration and forecast penetration for 2030 and 2040 (share of all 1,128 firms).

What the Bass analysis cannot show, however, is how this aggregate trajectory arises from the movement of individual firms between distinct states of engagement – above all, how the large pool of firms that today only intend to adopt will convert over time. The next chapter takes up exactly this question.

Table 10. Estimated Bass parameters by functional area ($m = 1,128$; 95% CIs from the non-linear least squares fit; saturation = first year reaching 95% of m).

Functional area	p (95% CI)	q (95% CI)	R ²	Peak	Saturation
Overall sample	0.0153 [0.000, 0.034]	0.721 [0.376, 1.065]	0.966	2024.2	2029
Marketing & Sales	0.0098 [0.003, 0.017]	0.501 [0.308, 0.694]	0.981	2026.7	2033
Production & Operations	0.0059 [0.002, 0.010]	0.486 [0.304, 0.668]	0.982	2028.0	2034
R&D and Innovation	0.0053 [0.002, 0.008]	0.484 [0.340, 0.628]	0.989	2028.2	2035
Logistics	0.0054 [0.002, 0.009]	0.450 [0.283, 0.616]	0.985	2028.7	2036
Corporate Management	0.0051 [0.000, 0.010]	0.440 [0.168, 0.713]	0.953	2029.0	2036
Human Resources	0.0037 [0.001, 0.006]	0.506 [0.349, 0.662]	0.987	2028.7	2035
Finance	0.0035 [0.001, 0.006]	0.451 [0.282, 0.620]	0.983	2029.7	2037
ICT Security	0.0042 [0.001, 0.007]	0.413 [0.236, 0.590]	0.980	2030.0	2038

7 The Functional Architecture of AI Diffusion: A Deeper Cross-Functional Analysis

This chapter deepens the functional-area forecasting by reading the eight estimated diffusion curves as a single, structured system rather than as isolated results. It develops a unified comparative framework, draws out the economic logic behind the ordering of functions, quantifies the gap between leading and lagging functions and its closing, and derives the implications of this functional structure for sequencing adoption.

7.1 A unified comparative framework

The disaggregated Bass analysis estimates the same model independently for each of eight functional areas of AI use. Read one by one, these estimates describe eight diffusion curves; read together, they reveal a *system* with a definite internal architecture. The purpose of this chapter is to make that architecture explicit, because the functional structure of diffusion — which functions lead, which lag, how far apart they are, and how the gap evolves — carries managerial and policy information that the aggregate curve, by construction, suppresses.

The estimated profiles are summarised in the Table 11, which collects the 2025 penetration level, the imitation coefficient, the timing of the peak of new adoptions, the year of 95% saturation, and the projected 2030 penetration for each function. The functions are ordered, as throughout, by descending 2025 penetration.

Table 11. Bass Diffusion Estimates by Functional Area of AI Use, 2025–2030.

Functional area	2025 penetration	Imitation coeff. q	Peak of new adoptions	95% saturation	Projected 2030 penetration
Marketing & sales	28.8%	0.50	~2027	~2033	~84%
Production & operations	18.1%	0.49	~2028	~2034	~73%
R&D and innovation	16.0%	0.48	~2028	~2035	~70%
Logistics	14.4%	0.45	~2028–29	~2036	~64%
Corporate management	13.9%	0.44	~2029	~2036	~60%
Human resources	12.9%	0.51	~2029	~2035	~66%
Finance	10.2%	0.45	~2030	~2037	~54%
ICT security	10.1%	0.41	~2030	~2038	~49%

Two regularities are immediately visible and organise the discussion that follows. First, the imitation coefficient dominates the innovation coefficient in *every* function, by one to two orders of magnitude — the contagion-driven dynamic identified for the sample as a whole is reproduced at the disaggregated level without exception. Second, with the instructive exception of human resources (discussed in Section 7.4), the functions line up in a near-monotone relationship between current penetration and the timing of the wave: the more penetrated a function is today, the earlier its estimated peak and the earlier its saturation. This penetration–timing regularity is the single most important structural feature of the functional system, and it is the subject of the next section.

7.2 The penetration–timing regularity and the ordered diffusion

That current penetration and the timing of the diffusion wave move together is not a tautology. In principle a function could be highly penetrated today yet still early in its wave (a fast starter with a distant peak), or lightly penetrated yet close to saturating a small ceiling. The estimates do not show this. Instead, marketing and sales — the front-runner at 28.8% — also peaks earliest (around 2027) and saturates earliest (around 2033), while finance and ICT security — the laggards at roughly 10% — peak latest (around 2030) and saturate latest (2037 and 2038 respectively). The functions thus occupy ordered positions along a *single underlying diffusion timeline*, each having entered the same contagion-driven process but at a different date.

This ordering licenses a specific and useful reframing: the eight functions can be read not as eight separate phenomena but as eight snapshots of one process observed at different stages of maturity. The laggard functions today look much as the leading functions did a few years ago, and the leading functions today prefigure where the laggards will be a few years hence. This is a strong, falsifiable claim about the structure of enterprise AI adoption, and it has an immediate analytical payoff: it means the experience of the leading functions is informative about the future of the lagging ones, which is precisely the kind of inference a forecasting exercise exists to support. It also implies that the *sequence* in which functions adopt is more robustly identified than the *exact date* at which any one of them reaches a given milestone — a distinction that matters for how the forecasts should be used, and one to which Section 9 returns.

7.3 The economics of heterogeneity: three tiers of diffusion

Behind the ordered timeline lies an economic logic. The functions sort into three recognisable tiers, distinguished by the height of the adoption threshold and the speed of the contagion once it begins.

The *leading tier* — marketing and sales, and to a lesser extent production and operations — comprises functions in which AI can be deployed with low organisational and technical commitment and with rapid, observable returns. Generative tools for content, conversational customer service, lead qualification and campaign analytics require little integration with core systems and produce quick, visible results. Low entry barriers lower the adoption threshold, and visible results amplify the peer imitation that the high imitation coefficient ($q \approx 0.50$) captures. These functions led the early wave precisely because they combine the two ingredients a contagion process needs: easy entry and conspicuous success.

The *middle tier* — R&D and innovation, logistics, and corporate management — comprises functions in which AI is more deeply embedded in existing processes and data infrastructures. Predictive maintenance and quality inspection in operations, routing and inventory optimisation in logistics, and analytics-based planning in management all require integration with established operational or enterprise-resource systems, which lengthens the implementation cycle and pushes the peak one to two years beyond that of the customer-facing leaders. The imitation coefficients remain substantial ($q \approx 0.44$ – 0.49) but the lower innovation coefficients and later peaks reflect the heavier complementary investment that adoption in these functions

demands — the intangible complements whose accumulation the productivity-J-curve literature identifies as the rate-limiting step (Brynjolfsson, Rock, & Syverson, 2021).

The *lagging tier* — finance and ICT security — comprises functions gated by specialised capability, data sensitivity and regulatory caution. AI in finance (fraud detection, forecasting, workflow automation) is typically introduced only once reliability and compliance considerations are satisfied; AI in cybersecurity depends on scarce technical expertise that diffuses less readily through simple peer imitation. ICT security exhibits the lowest imitation coefficient in the sample ($q = 0.41$), the clearest signature of a function in which contagion is weak because the relevant capability is not easily observed and copied. These functions will be the last to saturate, in the late 2030s, and the gap between them and the leaders is the structural feature that policy must reckon with (Section 10).

The three-tier structure is not merely descriptive. It maps the functional system onto the determinants identified in the literature — entry cost, complementary-investment intensity, and regulatory and capability gating — and thereby explains *why* the penetration–timing ordering takes the form it does, rather than simply recording that it does.

7.4 Convergence speed and the information in the imitation coefficient

While the penetration levels order the functions along the timeline, the imitation coefficients carry distinct information about the *speed* at which each function will traverse the remainder of its wave once it accelerates. The instructive case is human resources, which breaks the otherwise near-monotone penetration–timing relationship. HR is only the sixth most penetrated function (12.9%) yet exhibits one of the highest imitation coefficients in the sample ($q = 0.51$), alongside a very low innovation coefficient ($p = 0.004$). This combination is the classic *late-but-steep* signature: autonomous early adoption is rare — few firms introduce AI into recruitment screening or candidate communication on their own initiative — but once peer adoption becomes visible, contagion is strong and conversion is rapid. The implication is that HR, despite its current low base, could converge relatively quickly, compressing into a few years a transition that its present penetration would not lead one to expect. Its projected 2030 penetration of roughly 66% — higher than that of the more penetrated logistics and management functions — is the quantitative expression of this catch-up.

The contrast with ICT security sharpens the point. ICT security shares a similar low base (10.1%) but has the *lowest* imitation coefficient ($q = 0.41$), so its diffusion is both late and comparatively slow; its projected 2030 penetration of roughly 49% is the lowest of all functions. Reading the base and the imitation coefficient together therefore yields a richer prediction than either alone: a function’s *position* on the timeline is set by where it starts, but its *velocity* through the wave is set by the strength of the contagion mechanism, and the two need not coincide. This two-dimensional reading — level and speed — is the analytical advantage of estimating the model function by function rather than relying on a single aggregate curve, and it is the basis for the differentiated managerial sequencing discussed in Section 10.

7.5 Intra-firm diffusion pathways: from functions to bundles

The functional analysis so far treats each function as an independent diffusion process. The descriptive co-adoption evidence qualifies this in an important way: functions are adopted in correlated clusters rather than independently. The pairwise associations are uniformly positive — a firm using AI in one function is more likely to use it in others — and two clusters stand out, the strongest associations being between R&D and ICT security, and between corporate management and finance. These define two recognisable intra-firm pathways: a *technical* pathway, in which technologically sophisticated firms adopt AI across R&D and ICT security together, and an *administrative* pathway, in which control and reporting functions (management and finance) are digitalised in tandem.

This clustering changes the interpretation of the functional forecasts in two ways. First, it implies that the contagion captured by the high imitation coefficients operates along *two* dimensions simultaneously: between firms (the inter-firm peer effect that drives each function's curve) and between functions within a firm (the tendency for adoption in one function of a cluster to pull adoption in the others). The unit of diffusion is therefore better conceived as a *bundle* of complementary functional applications than as a single discrete act of adoption. Second, it partially reconciles the apparent paradox of the lagging technical functions: ICT security and finance lag on the aggregate timeline not because their adopters are uncommitted, but because they are concentrated among a specific subset of firms — the technically or administratively advanced — and diffuse to the broad population only as those firms complete their bundles. The co-adoption structure thus supplies the microfoundation for the macro-level functional ordering: within-firm pathways and between-firm contagion together generate the observed sequence.

7.6 The shifting centre of gravity of new adoption

A static reading of the functional table understates what the forecasts imply for the *composition* of new adoption over the horizon. Because the leading functions are already closer to their peaks, the marginal new adopter increasingly comes from the middle and lagging tiers as the decade progresses. In the near term, new adoption is dominated by the still-cresting customer-facing functions; by the early 2030s, as marketing and sales and production approach saturation, the centre of gravity of new adoption shifts toward the operational, administrative and support functions whose waves crest later. The functional composition of the AI-adoption “frontier” is therefore not fixed but rotates over time, from customer-facing applications toward operational, managerial and specialised ones.

This dynamic has a concrete consequence for any analysis that uses the diffusion forecasts as an input — for instance, the assessment of AI's effect on firm performance pursued in subsequent research. The functions in which adoption is growing fastest at a given date determine where the marginal performance effects of AI are most likely to be realised, and that location moves. An analysis conducted in 2026 would find the action in marketing and operations; the same analysis in 2032 would find it in management, HR, finance and ICT security. The rotation of the adoption frontier is thus not a curiosity of the curves but a

substantive prediction about where, within the enterprise, the economic impact of AI will successively concentrate.

7.7 Quantifying the functional gap and its closing

The functional system can be summarised by a single quantity of policy interest: the *spread* between the leading and lagging functions, and the date at which it closes. On the saturation metric, the leaders reach 95% penetration around 2033 (marketing and sales) and the laggards around 2037–2038 (finance and ICT security), a spread of roughly five years. On the 2030 penetration metric, the gap is starker — approximately 84% for marketing and sales against approximately 49% for ICT security, a difference of about 35 percentage points at a single point in time. The functional gap visible today is therefore not a transient artefact of measurement but a structural feature projected to persist for the better part of a decade before the laggards close it in the late 2030s.

The persistence of this gap is consequential. It means that, for much of the forecast horizon, the enterprise sector will be characterised by a pronounced internal unevenness in AI adoption across functions, even as the aggregate penetration rate climbs toward saturation. A headline statistic such as “two-thirds of firms have adopted AI” conceals this unevenness entirely; the functional decomposition restores it. For the interpretation of Slovakia’s AI transition — characterised in comparative analyses as real but incomplete — this is the precise sense in which incompleteness will manifest at the functional level: not as a uniform lag, but as a maturing customer-facing core coexisting with a still-developing operational, administrative and specialised periphery.

7.8 Cross-validation with the survey evidence and the international record

The functional ordering estimated by the diffusion model is not an isolated artefact of the modelling; it is corroborated from two independent directions, which materially strengthens confidence in it. The enterprise survey observed directly — in its own waves, and without any diffusion modelling — that AI use was most prominent in marketing and sales, followed by company management, R&D, finance, HR, ICT security and production-related activities, with the strongest recent increases in marketing and sales and in company management and a notable lag in logistics. The leading position of marketing and sales and the broad ordering coincide with what the Bass analysis recovers from the adoption-timing data, so two methodologically distinct exercises on the same population agree on the functional structure.

Beyond the present study, the international and grey-literature evidence reviewed earlier points the same way. Industry surveys repeatedly find marketing and sales and product development to be the functions in which AI — and generative AI in particular — is most often deployed (McKinsey, 2024), and the global evidence documents the same customer-facing leadership and functional heterogeneity (Daniotti et al., 2026). The convergence of an internal diffusion model, an internal survey and external benchmarks on the same functional ordering is strong evidence that the sequence reflects a structural property of how AI diffuses through the enterprise rather than a feature peculiar to the Slovak sample or to the modelling choices.

7.9 Uncertainty and what can be safely concluded

The deeper functional reading must be tempered by an honest account of its uncertainty. The estimates rest on short annual series, and the confidence intervals on the coefficients are correspondingly wide; the out-of-sample exercise and the bootstrap analysis in the robustness chapter quantify this directly. The appropriate conclusion is not that the functional forecasts are unreliable but that their *components* differ in robustness, and they should be used accordingly.

What is robust is the *ordering* of the functions and the *qualitative* structure of the system: the dominance of imitation in every function, the penetration–timing regularity, the three-tier economic structure, the co-adoption pathways, and the persistence of a multi-year functional gap. These conclusions recur across specifications and are corroborated by the independent survey and international evidence. What is less robust is the *precise timing* — the exact peak and saturation years for any single function — which carries a non-trivial margin given the brevity of the series and, for the functional analyses specifically, the attribution of a firm-level adoption year to each function. The safe register in which to use these forecasts is therefore comparative and ordinal: confident statements about which functions lead and lag, and about the approximate magnitude and persistence of the gap, are warranted; point claims about the calendar year in which a given function will reach a given level are not, and should be read as central estimates within a band rather than as precise predictions.

7.10 Implications of the functional structure for sequencing

The functional architecture translates into differentiated guidance that neither the aggregate forecast nor a single per-function reading would yield. For management, the three-tier structure and the co-adoption pathways jointly imply a *sequencing* strategy: beginning AI adoption in the firm’s most data-rich, low-integration functions — typically marketing and sales — where the entry threshold is lowest and returns are most quickly visible, and then extending adoption along the relevant cluster (technical or administrative) to exploit the within-firm contagion that the co-adoption evidence reveals. Such sequencing economises on the complementary investments in data, skills and process redesign that the maturity literature identifies as the binding constraint (Hansen, Lillesund, Mikalef, & Altwaijry, 2024), by reusing capabilities built in earlier functions for later ones in the same bundle.

For policy, the functional structure identifies where intervention is both most needed and most difficult. The lagging tier — finance and ICT security — will be the last to saturate and is gated not by simple awareness but by specialised capability and regulatory caution, exactly the barriers that the weak imitation coefficient of ICT security reflects. Because contagion is weakest precisely where the capability is hardest to observe and copy, the visibility-based instruments that work well for the customer-facing leaders are likely to be *less* effective here, and they need to be complemented by capability-building and regulatory-clarity measures targeted at the specialised functions. The differential incidence of the AI Act’s compliance obligations, which fall most heavily on the more regulated functions, reinforces this conclusion: the functional tier that diffusion theory predicts will lag is also the tier on which the regulatory burden is heaviest, so the two forces compound and warrant deliberate policy attention if the functional gap is not to harden into a permanent feature of the Slovak enterprise landscape.

8 State-Transition Dynamics: A Markov Chain Perspective

The Bass model treats adoption as a binary outcome and abstracts from the intermediate state of intention. This chapter complements that perspective with the discrete-time Markov chain developed in Section 2.3, which makes the transitions between the three states explicit and brings the planner pool, a direct determinant of future adoption, into the analysis.

8.1 Transition Matrices and Scenarios

Firms move between states over annual steps according to a transition matrix whose rows sum to one. Under the assumptions of single-step, upward-only progression and no reversal of forward progress, the matrix takes the general form set out in Section 2.3, where a denotes the probability that a non-adopter forms an intention ($N \rightarrow P$) and b the probability that a planner converts to adoption ($P \rightarrow A$). The two structural variants are shown in Tables 12 and 13.

Table 12. Transition matrix - absorbing variant (adopter state permanent; chain converges to full saturation).

Absorbing: From \ To	Non-adopter	Planner	Adopter
Non-adopter	$1 - a$	a	0
Planner	0	$1 - b$	b
Adopter	0	0	1

Table 13. Transition matrix – churn variant (5% annual outflow from adoption; chain converges to a plateau below full saturation).

Churn: From \ To	Non-adopter	Planner	Adopter
Non-adopter	$1 - a$	a	0
Planner	0	$1 - b$	b
Adopter	0.05	0	0.95

The two free probabilities are set by scenario: the conservative scenario represents a central, plausible pace of transition, the optimistic scenario a faster pace, and the pessimistic scenario a slower one. The corresponding values, together with the implied expected time for a non-adopter to progress through to adoption (computed as $1/a + 1/b$), are given in Table 14.

Table 14. Scenario transition probabilities and implied expected time to adoption.

Scenario	a ($N \rightarrow P$)	b ($P \rightarrow A$)	Expected time $N \rightarrow A$ (years)
Pessimistic	0.20	0.30	8.3
Conservative	0.35	0.50	4.9
Optimistic	0.60	0.75	3.0

8.2 Adoption Trajectories under the Two Variants

Applying the transition matrices to the initial state distribution and iterating over annual steps yields the adopter-share trajectories shown in Figure 16. The absorbing variant (solid lines) converges toward full saturation in every scenario, differing only in speed; the churn variant

(dashed lines) instead levels off at a stationary plateau that leaves a residual share of firms permanently outside adoption.




Figure 16. Markov chain adopter share by scenario and model variant, 2025–2040. Solid lines: absorbing variant; dashed lines: churn variant.

The headline results are summarised in Table 15. In the absorbing variant, the conservative scenario reaches 94.6% adoption by 2030 and approaches full saturation during the 2030s, bracketed by the optimistic (99.3% by 2030) and pessimistic (87.0% by 2030) scenarios. In the churn variant, the same scenarios level off at long-run plateaus of 80.5%, 87.0% and 70.6% respectively, implying that between roughly 13% and 29% of firms would remain non-adopters in the long run under a sustained 5% annual abandonment rate.

Table 15. Markov chain adopter share (% of 1,128) by scenario and variant.

Scenario	Absorbing 2030	Absorbing 2040	Churn 2030	Plateau
Pessimistic	87.0%	98.5%	71.3%	70.6%
Conservative	94.6%	99.9%	79.9%	80.5%
Optimistic	99.3%	100.0%	87.0%	87.0%

8.3 The Planner Pipeline

The distinctive contribution of the Markov model is its representation of the planner pool as a transient reservoir that feeds future adoption rather than as a static category. Figure 17 shows the evolution of the three states over time for the conservative scenario under both variants. In the absorbing variant the non-adopter and planner stocks are progressively drawn down as their mass flows into the permanent adopter state; in the churn variant the same downward flow is offset by the small return flow from adoption, so that the non-adopter and planner stocks stabilise at non-zero levels and the adopter share settles below full saturation.




Figure 17. State composition over time under the conservative scenario: absorbing variant (left) and churn variant (right).

This pipeline perspective complements the Bass results in a concrete way. The 216 firms currently classified as planners constitute the principal near-term source of new adoption, and the speed at which they convert – captured by the parameter b – is the main driver of how quickly the adopter share rises over the next several years. The framework thus links the present stock of stated intentions directly to the future trajectory of adoption, information that the aggregate Bass curve, by construction, cannot isolate.

9 Robustness and Sensitivity Analysis

This chapter probes the sensitivity of the principal results to the main modelling choices: the fixed market potential and take-off year of the Bass model, and the single-step structure and abandonment rate of the Markov model. The aim is to establish which conclusions are robust and which are conditional on specific assumptions.

9.1 Alternative Market Potential

The baseline analysis fixes the market potential at the full sample ($m = 1,128$), implying that even current non-adopters eventually adopt. An alternative, behaviourally grounded ceiling is the sum of adopters and planners ($m = 957$), which excludes the firms that today neither use nor plan to use AI. Re-estimating the aggregate model under this alternative leaves the qualitative picture unchanged: the imitation coefficient remains dominant (q rises from 0.721 to 0.826 as the lower ceiling is approached sooner), while the estimated 95% saturation year moves only marginally, from 2029 to 2028. Figure 18 compares the two trajectories.




Figure 18. Aggregate diffusion under alternative market potential ($m = 1,128$ vs. $m = 957$).

The two specifications differ mainly in the level at which the curve plateaus, not in the shape or timing of the diffusion. The choice of ceiling therefore affects the interpretation of the long-run saturation level but not the central finding of rapid, imitation-driven diffusion, and the $m = 1,128$ baseline can be read as a conservative (later-saturating) bound.

9.2 Sensitivity to the Take-off Year

The baseline models diffusion from a take-off in 2020. Re-estimating the aggregate model from alternative take-off years tests whether this choice drives the results. As Table 16 shows, the innovation coefficient p is sensitive to the take-off year – it rises as earlier, sparse years are excluded – but the imitation coefficient q is highly stable (around 0.70–0.72) and the estimated

95% saturation year is essentially invariant. Since the substantive conclusions rest on the dominance of q and on the timing of saturation, they are robust to the take-off choice.

Table 16. Sensitivity of the aggregate Bass estimates to the take-off year ($m = 1,128$).

Take-off year	p	q	Saturation (95%)
2019	0.0074	0.724	2029
2020 (baseline)	0.0153	0.721	2029
2021	0.0330	0.704	2029

9.3 A Direct Adoption Transition in the Markov Model

The Markov model assumes single-step progression, so that a non-adopter cannot adopt without first becoming a planner. Because low-friction generative-AI tools can in practice be adopted without a distinct planning phase, this assumption is relaxed by introducing a small direct $N \rightarrow A$ transition into the conservative absorbing scenario. As Table 17 shows, adding a direct transition of 5% or even 10% raises the 2030 adopter share only modestly (from 94.6% to 95.8% and 96.6% respectively) and leaves the long-run outcome unchanged. The staged-progression assumption is therefore not a material driver of the results.

Table 17. Sensitivity of the conservative absorbing scenario to a direct $N \rightarrow A$ transition.

Direct $N \rightarrow A$ probability	Adopter share 2030	Adopter share 2040
0 (baseline)	94.6%	99.9%
0.05	95.8%	100.0%
0.1	96.6%	100.0%

9.4 Sensitivity to the Abandonment Rate

The churn variant assumes a 5% annual outflow from adoption. Because this rate is itself an assumption, Table 18 reports the long-run adopter plateau under abandonment rates of 2%, 5% and 10% for each scenario. The plateau is materially sensitive to this rate: a 2% rate leaves a plateau of 86-94%, close to full saturation, whereas a 10% rate depresses it to 55-77%. This sensitivity is itself an important finding – the long-run ceiling of AI adoption depends not only on how fast firms adopt but, critically, on how durably they retain the technology.

Table 18. Long-run adopter plateau (% of 1,128) under alternative abandonment rates, by scenario.

Abandonment rate	Pessimistic	Conservative	Optimistic
2%	85.7%	91.1%	94.3%
5% (baseline)	70.6%	80.5%	87.0%
10%	54.5%	67.3%	76.9%

9.5 Out-of-sample Validation

To assess the predictive validity of the model rather than its in-sample fit, an out-of-sample backtest was performed. For each series the Bass model was re-estimated using only the adoption data up to 2024, and the resulting parameters were used to predict the cumulative number of adopters in 2025, which was then compared with the observed value. Because the

most recent year is withheld from estimation, this provides a genuine one-step-ahead test of the model's forecasting accuracy. The results are reported in Table 19.

Table 19. Out-of-sample backtest: models estimated on 2020–2024 data, predicting cumulative adoption in 2025 (counts). Mean absolute percentage error across the nine series: 15.1%.

Series	Observed 2025	Predicted 2025	Error	Error (%)
Overall Sample	741	621	-120	-16.1%
Marketing and Sales	325	277	-48	-14.6%
Production and Operations	204	181	-23	-11.2%
RD and Innovation	181	208	+27	+15.0%
Logistics	162	189	+27	+16.4%
Corporate Management	157	108	-49	-31.0%
Human Resources	146	129	-17	-11.6%
Finance	115	97	-18	-15.6%
ICT Security	114	109	-5	-4.0%

Across the nine series the mean absolute percentage error is 15.1%; for the aggregate sample the model predicts a 2025 penetration of 55.1% against an observed 65.7%, an under-prediction of 16.1%. The errors are predominantly negative: estimated on data only up to 2024, the model tends to under-predict the 2025 outcome because it cannot fully anticipate the continued acceleration of adoption that actually occurred. Two conclusions follow. First, one-year-ahead forecasts are accurate to within roughly fifteen percent on average, which is reasonable given the short series and supports the use of the model for medium-term projection. Second, the systematic under-prediction implies that, if anything, the main forecasts – estimated on the full series through 2025 – are more likely to be conservative than optimistic about the near-term pace of diffusion. The single weak case, corporate management (–31%), is the function with the weakest in-sample fit and the most diffuse adoption pattern, consistent with its lower estimability.

9.6 Forecast Uncertainty

The forecasts presented so far are point estimates. To convey the uncertainty around them, the parameter estimates for the aggregate model were propagated into the forecast by parametric bootstrap: 4,000 parameter pairs were drawn from the sampling distribution of the fitted coefficients, using the covariance matrix of the fit, and a forecast trajectory was computed for each draw. The resulting distribution of trajectories is summarised in Figure 19.




Figure 19. Aggregate AI diffusion forecast with uncertainty bands derived by parametric bootstrap (4,000 draws). Point forecast with 80% and 95% intervals; the observed 2025 value is the anchor.

The bands are widest in the near term – where the steepness of the curve makes the trajectory most sensitive to the estimated coefficients – and narrow as the process approaches its ceiling. By 2030 the point forecast of 98.6% penetration carries an 80% interval of roughly 96–99% and a 95% interval of roughly 93–99%, so that even at the lower bound the sample is projected to be close to saturation by the end of the decade. Translating the same parameter uncertainty into the timing of saturation, the year in which the sample reaches 95% penetration has a median of 2029 and a 95% interval of 2028–2031. The narrowness of this interval indicates that, while the precise level of penetration in any given year is uncertain, the broad conclusion – near-saturation of the sample around the end of the 2020s – is robust to the estimation uncertainty in the parameters.

10 The Policy Context of AI Diffusion in the European Union and Slovakia

This chapter situates the forecasting analysis within the regulatory and policy environment that conditions AI diffusion. It explains why the policy setting matters for the diffusion mechanism identified in the empirical chapters, and reviews in turn the European regulatory framework (the AI Act), the European digital-transformation agenda, and the national context and starting position of Slovakia.

10.1 Why the policy environment matters for diffusion

The diffusion models developed in this monograph describe a behavioural process — the spread of AI adoption through peer observation, competitive pressure and the conversion of stated intentions. That process does not unfold in an institutional vacuum. The pace at which firms adopt, and the ceiling toward which adoption tends, are shaped by a regulatory and policy environment that can lower or raise the cost of adoption, alter the visibility and perceived risk of AI use, and address the very capability and trust barriers that comparative analysis of Slovakia's AI position identified as binding. This chapter therefore complements the model-based analysis with a brief account of the principal policy layers relevant to Slovakia: the European regulatory framework, the European digital-transformation agenda, and the national context. The treatment is intentionally selective; a comprehensive policy analysis lies beyond the scope of a forecasting monograph.

10.2 The European regulatory framework: the AI Act

The most consequential regulatory development for AI diffusion in the European Union is Regulation (EU) 2024/1689, the Artificial Intelligence Act, the first comprehensive horizontal regulation of AI by a major jurisdiction. The Act adopts a risk-based architecture, distinguishing prohibited practices, high-risk systems subject to substantial conformity obligations, limited-risk systems subject to transparency requirements, and minimal-risk systems left largely unregulated. Its provisions apply on a staggered timetable: the general provisions and the prohibitions on unacceptable-risk practices took effect on 2 February 2025; the obligations for providers of general-purpose AI models and the governance architecture from 2 August 2025; the bulk of the rules, including those for high-risk systems and the transparency obligations, from 2 August 2026; and the framework becomes fully applicable, including for general-purpose models already on the market, by 2 August 2027.

For the diffusion process modelled here the Act cuts two ways, and the ambivalence is itself analytically relevant. On the one hand, by establishing a predictable legal framework and addressing the trust and legal-clarity concerns found to intensify as firms move from awareness to practical use, the Act may *reduce* a barrier to adoption and support the conversion of the planner pool. On the other hand, the compliance obligations it imposes — particularly on high-risk applications — raise the fixed cost of certain categories of adoption and may slow diffusion in the more regulated functions. This is consistent with the functional ordering estimated in the

empirical chapters, in which the lightly regulated, customer-facing functions of marketing and sales lead the diffusion process while the more regulated or data-sensitive functions, such as finance and ICT security, trail. The regulatory burden is thus not uniform across the functional areas but falls differentially, reinforcing the staggered functional sequence that the diffusion model identifies on independent grounds.

10.3 The European digital-transformation agenda

Beyond the regulatory framework, the European Union pursues an affirmative digital-transformation agenda that sets quantitative ambitions for technology uptake. Under the Digital Decade Policy Programme 2030, the Union has committed to a target whereby at least 75% of enterprises take up cloud computing, artificial intelligence or big-data analytics by 2030, alongside the goal that the large majority of small and medium-sized enterprises reach at least a basic level of digital intensity. These targets establish a policy benchmark against which national progress is monitored, most visibly through the Digital Economy and Society Index (DESI), whose assessment situates Slovakia below the European average.

The agenda is operationalised through instruments directly relevant to the adoption mechanism this monograph identifies. European Digital Innovation Hubs, the Digital Europe Programme and the financing available under the Recovery and Resilience Facility provide advisory, demonstration and co-financing support that lowers the marginal cost of moving from intention to operation. Read against the imitation-dominated diffusion regime estimated in the empirical chapters, instruments of this kind — which fund demonstrators, diffuse case evidence and build shared capability — work with the grain of the diffusion mechanism, since they increase the visibility and lower the perceived risk of adoption rather than relying on the empirically minor channel of autonomous, externally motivated decisions. The policy architecture and the estimated diffusion dynamics are therefore mutually legible: the instruments most emphasised by the European agenda are also those the diffusion analysis identifies as most likely to be effective.

10.4 The national context and the position of Slovakia

For Slovakia the European framework is mediated by the national policy context and by the structural starting position documented in comparative analyses, which placed Slovakia below the European frontier across a range of AI- and digital-readiness indices — including the DESI, the Government AI Readiness Index and the Network Readiness Index — and characterised the country as undergoing a real but incomplete AI transition marked by persistent barriers of cost, skills, data-readiness and trust. The national transposition and implementation of the European agenda — the designation of competent authorities under the AI Act, the national digital strategy, and the deployment of European funding instruments — therefore operate on an enterprise base that is internally stratified, with an advanced, adoption-prone core coexisting with a broader and more cautious majority.

The implication for the reading of this monograph's forecasts is direct. The principal long-run risk the modelling identifies is not a failure of aggregate adoption — which the forecasts suggest will be substantial in any plausible scenario — but the entrenchment of a divide between an

advanced core and a lagging periphery of firms, the churn-driven residual that the Markov analysis brings into focus. National policy that channels the European instruments toward the absorptive capacity of the broad enterprise base, and particularly toward the smaller and less resourced firms most exposed to permanent non-adoption, is therefore the policy posture most consonant with the structure of the diffusion process. In this sense the policy context does not merely surround the forecasting analysis; it identifies the margin on which policy can most plausibly shift the trajectory the models project, and it links the predictive layer of this monograph to the cohesion objectives that motivate this research agenda.

11 Discussion

The two modelling approaches converge on a coherent and mutually reinforcing interpretation of how artificial intelligence is diffusing through the Slovak enterprise sector. The Bass model identifies the mechanism and timing of the aggregate and functional diffusion curves, while the Markov model traces the behavioural transitions that underlie them; read together, and corroborated by the robustness analysis, they triangulate the same underlying process from two directions.




Figure 20. Triangulating Bass and Markov Adoption Curves (Absorbing vs. Churn, Conservative Scenario), 2025–2040.

A direct comparison of the two models reinforces this reading. Figure 20 overlays the aggregate Bass penetration curve with the Markov adopter-share trajectories under the conservative scenario. Although the two models are estimated independently and rest on different logics – a continuous diffusion curve fitted to adoption timing, versus a discrete state-transition process – they share the observed 2025 adopter share as a common anchor and then trace closely comparable paths. Under the absorbing assumption, which matches the full-saturation ceiling of the Bass specification, the Markov trajectory tracks the Bass curve to within a few percentage points throughout the horizon, both converging toward near-universal adoption during the 2030s. The agreement of two methodologically distinct approaches on the near-term pace and the long-run direction of adoption is the substance of the triangulation: it indicates that the central conclusion does not depend on the particular functional form of either model. The churn trajectory, by contrast, illustrates how far the long-run outcome is governed by the abandonment assumption rather than by the diffusion mechanism – opening a gap below the jointly implied path that widens only after the near-term trajectories have already coincided.

The most consistent finding across both methods and all analyses is the dominance of imitation over innovation. In every Bass estimate the imitation coefficient exceeds the innovation

coefficient by one to two orders of magnitude, a result that the sensitivity analysis shows to be robust to the take-off year and to the market-potential assumption, and the Markov scenarios that best match the observed pace of adoption are those in which planners convert rapidly once peer adoption is visible. Substantively, AI adoption in Slovakia is propagating through competitive pressure, peer observation and social learning rather than through autonomous, externally initiated decisions. For policy design the concern of innovation and technology policy this implies that interventions which increase the visibility of successful adoption (case studies, sectoral benchmarks, demonstrator projects and shared standards) are likely to be more effective than broad awareness campaigns aimed at stimulating independent first-mover adoption.

The functional analysis adds an ordered, time-staggered structure to this picture. Customer-facing applications lead, followed by operational and knowledge-intensive functions, with management, human resources and the more regulated or specialised functions trailing; because the more penetrated functions are also closer to their peak and saturation, the centre of gravity of new adoption will shift over the forecast horizon from customer-facing toward operational, administrative and support functions. The extended horizon makes clear that, while the leading functions approach saturation in the early 2030s, the laggard functions do not do so until the late 2030s, so that the functional gap observed today will persist for the better part of a decade before closing. The descriptive co-adoption analysis qualifies this picture by showing that functions are adopted in correlated clusters – a technical cluster around R&D and ICT security, and an administrative cluster around management and finance – rather than independently, which suggests that the diffusion of AI within a firm proceeds along recognisable pathways.

The Markov model contributes two further insights. First, it foregrounds the planner pool: the conversion of the 216 firms that currently intend to adopt is the dominant near-term driver of the adoption trajectory, which directs attention to the determinants of that conversion. Second, the contrast between the absorbing and churn variants, together with the abandonment-rate sensitivity, frames the long-run ceiling as a substantive question rather than a modelling artefact. If adoption is effectively irreversible, the sector converges to near-universal use; if a modest abandonment persists, a structural residual of firms remains outside adoption, the size of which depends sharply on the abandonment rate. This range is of direct relevance to the assessment of firm performance pursued in subsequent research, since it delineates the extent to which AI use will become a universal baseline rather than a persistent differentiator.

These interpretations must be qualified by the assumptions set out in the methodology. The fixed market potential and the assumed transition probabilities mean that the quantitative forecasts are conditional rather than unconditional, and the ordering of functions is more robust than the precise timing within them. The robustness analysis nonetheless shows that the central qualitative conclusions – imitation-driven diffusion, the ordered functional sequence, and the pivotal role of the planner pool – hold across the range of specifications examined.

12 Extended Interpretation of the Forecasting Results and Their Implications

This chapter draws the empirical results together into an extended interpretation and sets out their implications. It examines the imitation-dominated diffusion regime, the reconciliation of the two models and the open question of the long-run ceiling, the functional sequencing of adoption and the role of the planner pool, and then develops the implications for management, for policy in a small open economy, and for the analysis of the AI–performance link and future research.

12.1 The diffusion regime: interpreting an imitation-dominated process

The single most robust empirical regularity of the analysis is the dominance of imitation over innovation. In the aggregate Bass estimate the imitation coefficient ($q = 0.721$) exceeds the innovation coefficient ($p = 0.0153$) by a factor of roughly fifty, and the same ordering — by one to two orders of magnitude — recurs in every functional-area estimate and survives the robustness checks on the take-off year and the market-potential assumption. This section draws out what such a regime means economically, why it is plausible for AI specifically, and what it does and does not allow one to infer about the future.

In the language of diffusion theory (Rogers, 2003; Mahajan, Muller, & Bass, 1990), a high q/p ratio describes a process in which adoption is propagated overwhelmingly through *internal* influence — peer observation, competitive pressure and social learning — rather than through *external*, autonomous decision-making. The economic interpretation is that, for the typical Slovak enterprise, the decision to adopt AI is conditioned far more on the visible behaviour of comparable firms than on independent strategic initiative or exposure to external information. This is precisely the signature one would expect of a general-purpose technology whose value to each adopter rises as adoption spreads and as complementary skills, tools and service markets thicken around it (Brynjolfsson, Rock, & Syverson, 2021). It is also consistent with the discontinuity introduced by generative AI from late 2022, whose highly visible, low-entry-cost tools created exactly the kind of observable, imitable adoption that a contagion process feeds upon.

The magnitude of the estimated q warrants comparison with the wider empirical literature. Reviews of Bass applications across durable goods, telecommunications and information technologies typically report imitation coefficients in the range of roughly 0.3–0.5, with innovation coefficients an order of magnitude smaller (Mahajan et al., 1990; Meade & Islam, 2006). An imitation coefficient above 0.7 therefore places AI diffusion in the Slovak enterprise sample at the *fast* end of the historical spectrum — a finding that should be read in light of two features of the data. First, the adoption series is short and was estimated through a period of exceptional acceleration, which tends to load recent momentum onto the imitation term. Second, the out-of-sample exercise shows the model *under*-predicting the realised 2025 outcome (55.1% against an observed 65.7%), which implies that, if anything, the estimated pace is conservative rather than inflated. The high q is thus not an artefact of optimistic fitting;

it reflects a genuinely rapid, peer-driven diffusion, while the short series counsels against treating its precise value as a structural constant.

This reading connects directly to two strands of the recent forecasting literature. The notion of a *critical mass* beyond which adoption accelerates sharply — placed by Olfati-Saber (2025) in the region of 5–15% of the population, albeit in a non-peer-reviewed framework — is consistent with the steep mid-decade rise the model projects once an initial base of visible adopters has formed. At the same time, the work of Roman and Bertolotti (2025) on network diffusion and boom-bust dynamics is a useful caution: an imitation-dominated process is, by its nature, sensitive to the continued visibility and perceived success of early adoption, and a contagion that can accelerate quickly can in principle also stall if early adopters are seen to fail. The model’s projection of smooth approach to saturation should therefore be understood as conditional on momentum being sustained, not as an inevitability — a qualification that has direct implications for policy, developed in Section 5.

12.2 Two models, one trajectory — and the open question of the ceiling

A central methodological claim of the analysis is that two structurally distinct models — a continuous Bass curve fitted to adoption timing, and a discrete three-state Markov chain — triangulate the same underlying process. The value of that claim rests on where the models agree and, equally informatively, on where they diverge.

They agree, closely, on the near-term trajectory and the long-run *direction*. Anchored on the common observed 2025 adopter share, the Markov absorbing variant under the conservative scenario tracks the aggregate Bass curve to within a few percentage points across the whole horizon, both converging toward near-universal adoption during the 2030s (the Bass curve reaches 98.6% by 2030; the Markov conservative-absorbing path 94.6% by the same year). Because the two approaches rest on entirely different primitives, their agreement on the pace and direction of diffusion indicates that this central conclusion does not depend on the functional form of either model. This is the substance of triangulation: a result that survives a change of method is more credible than one that does not.

They diverge — by construction — on the long-run *ceiling*, and this divergence is the analytically richest feature of the forecast. The absorbing variant, in which adoption is irreversible, converges to near-universal use; the churn variant, with a sustained 5% annual outflow from adoption, instead levels off at plateaus of 80.5%, 87.0% and 70.6% in the conservative, optimistic and pessimistic scenarios respectively, leaving between roughly 13% and 29% of firms permanently outside adoption. The gap between these two families of trajectories opens only *after* the near-term paths have already coincided, which localises the disagreement precisely: the data and the diffusion mechanism govern the next several years, whereas the long-run outcome is governed by an assumption — the reversibility of adoption — that the present data cannot adjudicate.

This is not a weakness of the analysis but a clarification of what is and is not knowable from current evidence, and it is here that the comparative diagnosis of Slovakia’s AI position becomes analytically load-bearing. That diagnosis characterised Slovakia as undergoing a *real but incomplete* AI transition, positioned below the European frontier and marked by persistent

barriers that intensify as firms move from awareness to practical use. Read against the two Markov variants, this diagnosis tilts the interpretation away from the frictionless absorbing path: an environment in which cost, skills, data-readiness and trust barriers remain salient is one in which some abandonment, and hence a ceiling below universal adoption, is the more plausible long-run expectation. The churn variant should therefore not be regarded as a pessimistic curiosity but as the scenario most consistent with the structural conditions independently documented for Slovakia. The policy question that follows — whether AI use becomes a universal baseline or settles as a persistent differentiator across firms — is taken up in Section 5 and bears directly on the subsequent analysis of firm performance.

12.3 Functional sequencing and the architecture of intra-firm diffusion

Disaggregating the diffusion process by functional area converts a single S-curve into an ordered, time-staggered family of curves, and this structure carries more information than the aggregate alone. Marketing and sales leads the sequence, at 28.8% penetration in 2025, followed by operational and knowledge-intensive functions, while the more regulated or specialised functions — and corporate management — trail at around 10%. Because the leading functions are already closer to their peak and to saturation, the *centre of gravity* of new adoption shifts over the horizon: the marginal new adopter increasingly comes from operational, administrative and support functions rather than from the customer-facing applications that led the early wave. The leading functions approach saturation in the early 2030s, the laggards not until the late 2030s, so that the functional gap visible today persists for the better part of a decade before closing.

This ordering is not idiosyncratic to Slovakia; it mirrors the international and grey-literature evidence reviewed earlier, in which marketing, sales and product development repeatedly emerge as the leading functions (McKinsey, 2024; Daniotti et al., 2026), and it coincides with the functional pattern observed directly in the enterprise survey waves, where marketing and sales and company management showed the strongest recent increases and logistics lagged. The convergence of an independent forecasting model, the enterprise survey and external benchmarks on the same functional ordering strengthens confidence that the sequence reflects a structural feature of how AI diffuses through the enterprise rather than a sampling artefact.

The descriptive co-adoption evidence adds a further dimension by showing that functions are adopted in correlated clusters — a technical cluster around research and development and ICT security, and an administrative cluster around corporate management and finance — rather than independently. Substantively, this implies that AI diffuses *within* a firm along recognisable pathways: an enterprise that adopts in one function of a cluster is disproportionately likely to adopt in the others, so that adoption propagates not only across firms (the inter-firm contagion captured by the high q) but also across functions inside the firm. This dual contagion — between firms and between functions — is the microstructure beneath the aggregate curve, and it suggests that the unit of diffusion is better conceived as a *bundle* of complementary functional applications than as a single, discrete act of adoption. The managerial corollary, developed in Section 4, is that capability investments which unlock one function of a cluster are likely to yield adoption spillovers within it.

12.4 The planner pipeline and the calculus of conversion

The distinctive contribution of the Markov framework is to make the *planner pool* an explicit, dynamic object rather than a static survey category. The 216 firms currently classified as planners constitute the principal near-term reservoir of new adoption, and the rate at which they convert — the transition probability from planning to adoption — is the dominant driver of how quickly the adopter share rises over the next several years. This reframing has an important practical consequence: it identifies a concrete, observable population whose behaviour can be influenced, and it locates the most actionable leverage over the adoption trajectory not in the conversion of indifferent non-adopters but in the acceleration of firms that have already declared an intention.

The significance of this pipeline is reinforced by the dynamics that the enterprise survey documented independently. The survey found a redistribution across its waves *from* current use *toward* planning, with the 2026 CAWI data showing a particularly large planning segment. Interpreted through the Markov model, this is not evidence of waning relevance but of a swelling reservoir feeding future adoption: a larger planner stock today implies a larger flow into adoption tomorrow, provided the conversion rate holds. The two strands thus describe the same phenomenon from two angles — the survey measures the growing stock of intentions, while the present model captures the rate at which that stock converts — and together they make the conversion of planners the empirical fulcrum of the medium-term forecast. The uncertainty surrounding the forecast should be stated with equal precision. The parametric bootstrap places the 2030 aggregate penetration point estimate of 98.6% within an 80% interval of roughly 96–99% and a 95% interval of roughly 93–99%, and the year in which the sample reaches 95% penetration has a median of 2029 with a 95% interval of 2028–2031. The bands are widest in the near term, where the steepness of the curve makes the path most sensitive to the estimated coefficients, and narrow as the process approaches its ceiling. Two points deserve emphasis. First, although the *level* in any given year is uncertain, the *qualitative* conclusion — near-saturation of the sample around the end of the 2020s — is robust to estimation uncertainty. Second, this bootstrap interval reflects only parameter uncertainty within the fixed-market-potential specification; it does not encompass the deeper, structural uncertainty about the ceiling itself, which the absorbing-versus-churn contrast addresses separately. A complete account of forecast uncertainty therefore reads the bootstrap bands and the scenario range together: the former bounds the trajectory toward a given ceiling, the latter bounds the ceiling.

12.5 Implications for management

The findings translate into a set of implications for enterprise decision-makers that follow directly from the structure of the diffusion process rather than from general exhortation. Three are central.

First, because diffusion is overwhelmingly imitation-driven, the *visibility* of successful adoption is a strategic asset, both for the individual firm and for the sectors in which it operates. For a firm weighing adoption, the implication is that the experience of comparable peers is the most informative available signal; for a firm that has already adopted, there is value in the credible demonstration of results, since an environment rich in visible success accelerates the

contagion from which early adopters themselves benefit through thicker skills and service markets. The corresponding risk is symmetric: in a contagion-driven regime, visible failures propagate as readily as visible successes, which places a premium on adopting where the probability of demonstrable benefit is highest rather than on first-mover adoption for its own sake.

Second, the functional and co-adoption structure implies that adoption is best sequenced rather than attempted wholesale. The evidence that AI diffuses along functional clusters — technical (R&D, ICT security) and administrative (management, finance) — suggests that capability investments unlocking one function of a cluster are likely to generate adoption spillovers within it, so that a deliberate sequencing strategy that begins in the firm's most data-rich, communication- and analysis-intensive functions and then extends along the cluster can economise on the complementary investments in data, skills and process redesign that the maturity literature identifies as the binding constraint (Hansen, Lillesund, Mikalef, & Altwaijry, 2024). The leading position of marketing and sales is consistent with this logic, since these functions combine high data availability with low entry barriers and rapid, observable returns.

Third, the contrast between the absorbing and churn trajectories reframes the managerial challenge as one of *sustaining* adoption, not merely initiating it. A sustained abandonment rate of only 5% per year is sufficient to leave between 13% and 29% of firms outside adoption in the long run; abandonment is therefore not a marginal phenomenon but a decisive determinant of whether AI becomes a durable capability. For management this directs attention to the post-adoption phase — the embedding of AI into workflows, the retention of skills, and the realisation of measurable value — as the locus at which the long-run payoff is won or lost. The productivity-J-curve logic (Brynjolfsson et al., 2021) reinforces the point: because measured returns may lag the accumulation of intangible complements, firms that abandon prematurely forgo precisely the delayed benefits that justify the initial investment.

12.6 Implications for policy in a small open economy

For policy, the analysis carries implications that are sharpened by Slovakia's documented position as a below-frontier economy undergoing an incomplete AI transition. Three follow most directly and align with broad concerns of technology policy.

The first concerns the *design* of intervention. In an imitation-dominated regime, instruments that increase the visibility of successful adoption — curated case studies, sectoral benchmarks, demonstrator and lighthouse projects, shared standards and interoperability frameworks — work *with* the grain of the diffusion mechanism and are likely to be more effective per unit of expenditure than broad awareness campaigns aimed at stimulating isolated, autonomous first-mover adoption. The model's structure thus does more than describe diffusion; it discriminates between classes of policy, favouring those that amplify peer signals over those that address the (empirically minor) channel of external influence. The second concerns the *target* of intervention. Because the planner pool is the dominant near-term driver of adoption, and because the survey evidence documents a growing planning segment, the highest-leverage policy focus is the conversion of firms that have already declared an intention but face the cost, skills, data-readiness and trust barriers found to intensify at exactly this transition. Measures

that lower the marginal cost of moving from intention to operation — targeted skills provision, advisory and implementation support, access to shared data and compute, and regulatory clarity — address the conversion bottleneck directly. This is a more precise prescription than a generalised push to raise awareness, since awareness is not the binding constraint for the population that matters most to the near-term trajectory.

The third concerns the *long-run distributional stakes*, which the absorbing-versus-churn contrast brings into focus. If adoption is effectively irreversible, AI use becomes a near-universal baseline and ceases to be a source of competitive differentiation among firms; if a modest abandonment persists, a structural residual of firms — disproportionately likely, in a below-frontier economy, to be smaller and less resourced enterprises — remains permanently outside adoption. The policy implication is that the central risk for Slovakia is not a failure of aggregate adoption, which the forecasts suggest will be substantial in any case, but the *entrenchment of a divide* between an advanced core and a lagging periphery of firms. This connects the forecasting directly to the cohesion objectives of digital policy: avoiding the churn-driven residual requires sustained attention to the absorptive capacity of the broad base of enterprises, not only to the dynamism of the frontier. Given the stratification identified within the Slovak enterprise population, this is the scenario against which policy should most deliberately guard.

12.7 Implications for the AI–performance link and for future research

The diffusion trajectories developed here are an input to, rather than the endpoint of, a broader analytical agenda. The projected paths frame the expected evolution of the relationship between AI adoption and firm performance in a specific way: if the absorbing path obtains and adoption approaches universality, the cross-sectional association between adoption and performance will weaken over time as the comparison group of non-adopters shrinks, and the relevant performance margin will shift from *whether* a firm adopts to *how intensively and effectively* it does so. If, instead, the churn path obtains and a structural residual persists, adoption status will remain a meaningful differentiator for longer. The choice between these readings is therefore not merely methodological; it determines the very design of any future analysis of AI’s performance effects, which should accordingly distinguish extensive-margin (adoption) from intensive-margin (intensity and maturity) effects.

Several directions for further work follow. The transition probabilities of the Markov model, presently assumed and varied by scenario, should be estimated directly from observed state-to-state movements as successive survey waves accumulate, converting the scenario range into an empirically grounded distribution. The fixed market-potential assumption and the absorbing-versus-churn dichotomy could be developed into a fuller set of saturation scenarios anchored to external estimates of the relevant enterprise population, and the abandonment rate — shown here to be decisive for the long-run ceiling — could be calibrated empirically from longitudinal data. Independent validation of the Bass estimates, including through alternative estimation routines, would further consolidate the predictive layer. Together these steps would tighten the conditional forecasts presented here into progressively less conditional ones, and would embed the predictive layer firmly within a broader research agenda on AI adoption and firm performance.

13 Conclusion

This final chapter draws together the principal findings, restates the main limitations, and sets out the recommendations for further work.

13.1 Key Findings

This monograph has applied two complementary models — the Bass diffusion model and a three-state Markov chain — to enterprise-level survey data on Slovak firms in order to forecast the diffusion of artificial intelligence among Slovak enterprises. Four principal findings emerge. First, AI diffusion in Slovakia is strongly imitation-driven: across both methods and robust to the specifications tested, contagion among firms rather than autonomous adoption is the primary mechanism. Second, diffusion is highly uneven across functions, with marketing and sales leading at 28.8% penetration in 2025 and the more regulated or specialised functions trailing at around 10%, and with saturation staggered from the early to the late 2030s. Third, the planner pool is the dominant near-term driver of future adoption, so that the pace at which current intentions convert governs the trajectory over the next several years. Fourth, the long-run ceiling depends on whether adoption is irreversible: under an absorbing adopter state the sector converges to near-universal use, whereas a sustained abandonment leaves a structural residual of firms outside adoption whose size is highly sensitive to the abandonment rate.

13.2 Limitations

The findings are conditional on the modelling assumptions discussed in Chapter 3 and probed in Chapter 8. The Bass market potential was fixed at the sample size, so its forecasts describe trajectories toward saturation of the sample; the firm-level adoption year was attributed to each functional area, making within-function timing approximate; and the high R^2 values reflect the fitting of cumulative curves. For the Markov model, the transition probabilities are assumed rather than estimated, and the single-step, upward-only structure is a deliberate simplification, although the robustness analysis shows the results to be insensitive to relaxing it. The absolute penetration levels also depend on the representativeness of the sample, which cannot be fully verified.

13.3 Recommendations and Implications for Future Work

Three directions for further work follow. First, the forecasts should be validated through independent estimation, including alternative estimation routines, and revisited as additional survey waves extend the adoption series beyond its peak and, crucially, make it possible to estimate the Markov transition probabilities directly from observed state-to-state movements rather than assuming them. Second, the fixed market-potential assumption and the absorbing-versus-churn dichotomy could be developed into a fuller set of saturation scenarios anchored to external estimates of the relevant enterprise population, and the abandonment rate could be calibrated empirically as longitudinal data accumulate. Third, the functional-area and state-transition forecasts provide a natural input to subsequent analysis of firm performance, where the projected diffusion trajectories can frame the expected evolution of the relationship between

AI adoption and firm performance. Together, these steps would consolidate the predictive layer developed here within a broader research agenda on AI adoption and firm performance.

References

- Andersson Lipcsey, R. (2025). *AI diffusion to low- and middle-income countries: A blessing or a curse?* Working paper, arXiv:2405.20399v3.
- Bass, F. M. (1969). A new product growth model for consumer durables. *Management Science*, 15(5), 215–227.
- Berger, P., & von Garrel, J. (2024). Diffusion of AI value-driven services in the German manufacturing industries — an empirical examination of value-driven service references classified by the Business Model Canvas. *Frontiers in Engineering Management*, 2, 1407367.
- Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1), 333–372.
- Cirera, X., Comin, D., & Cruz, M. (2022). *Bridging the technological divide: Technology adoption by firms in developing countries*. Washington, DC: World Bank.
- Daniotti, S., et al. (2026). Global adoption and use of generative artificial intelligence. *Science*. Advance online publication. <https://doi.org/10.1126/science.adz9311>
- European Commission. (2024). *Digital Decade Policy Programme 2030: Targets for the digital transformation of enterprises*. Brussels: European Commission.
- European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union.
- Eurostat. (2024). *Use of artificial intelligence in enterprises*. Luxembourg: Eurostat.
- Garlatti Costa, G., Venier, F., & Pugliese, R. (2025). Unveiling organizational AI adoption patterns in Italian companies through the lens of the diffusion of innovations theory. *Managing Global Transitions*, 23(1), 79–111.
- Geroski, P. A. (2000). Models of technology diffusion. *Research Policy*, 29(4–5), 603–625.
- Goel, R. K., & Nelson, M. A. (2025). Awareness of artificial intelligence: Diffusion of AI versus ChatGPT information with implications for entrepreneurship. *The Journal of Technology Transfer*, 50, 96–113.
- Hansen, H. F., Lillesund, E., Mikalef, P., & Altwaijry, N. (2024). Understanding artificial intelligence diffusion through an AI capability maturity model. *Information Systems Frontiers*, 26, 2147–2163.
- Hodak, M., Lakhani, M., & Cheng, D. (2025). *Decoding AI diffusion: Mapping the path of transformative AI across industries*. Industry report, SSRN No. 5475227. [Non-peer-reviewed.]
- Madan, R., & Ashok, M. (2023). AI adoption and diffusion in public administration: A systematic literature review and future research agenda. *Government Information Quarterly*, 40(1), 101774.
- Mahajan, V., Muller, E., & Bass, F. M. (1990). New product diffusion models in marketing: A review and directions for research. *Journal of Marketing*, 54(1), 1–26.

- McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L., & Zolas, N. (2024). AI adoption in America: Who, what, and where. *Journal of Economics & Management Strategy*, 33(2), 375–415.
- McKinsey & Company. (2024). *The state of AI in early 2024: Gen AI adoption spikes and starts to generate value*. McKinsey Global Survey. [Industry survey.]
- Meade, N., & Islam, T. (2006). Modelling and forecasting the diffusion of innovation – A 25-year review. *International Journal of Forecasting*, 22(3), 519–545.
- Norris, J. R. (1997). *Markov chains*. Cambridge: Cambridge University Press.
- OECD. (2024). *Explanatory memorandum on the updated OECD definition of an AI system*. Paris: OECD Publishing.
- OECD. (2025). *Emerging divides in the transition to artificial intelligence*. Paris: OECD Publishing.
- Olfati-Saber, R. (2025). *A unified framework for technology adoption and innovation diffusion: Theory, evidence, and AI policy implications*. SSRN No. 5860839. [Preprint; not peer-reviewed.]
- Reis, J. F., & Pinheiro Junior, L. P. (2025). Institutional pressures and the diffusion of artificial intelligence: An institutional-theory and diffusion-of-innovations perspective. *BAR — Brazilian Administration Review*, 22(4), e250060.
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). New York: Free Press.
- Tornatzky, L. G., & Fleischer, M. (1990). *The processes of technological innovation*. Lexington, MA: Lexington Books.

Appendix A Annual Adoption Counts by Functional Area

The following table reports the number of new adopters recorded in each period for the overall sample and for each functional area, with adoptions before 2020 pooled into a single take-off period and 2026 adoptions assigned to 2025. These are the empirical inputs to the diffusion models.

Table 20. New adopters per period, overall and by functional area (counts).

Period	Overall	Marketing & Sales	Production & Operations	R&D & Innovation	Logistics	Corporate Management	Human Resources	Finance	ICT Security
≤2019	70	28	15	11	12	17	11	8	12
2020	35	16	10	6	5	9	6	7	5
2021	11	2	5	4	4	1	1	1	0
2022	52	29	16	12	13	10	10	8	8
2023	103	46	22	33	24	20	22	19	18
2024	213	80	59	59	58	30	36	25	32
2025	257	124	77	56	46	70	60	47	39

Appendix B Numerical Markov Transition Matrices

The following tables give the fully specified transition matrices for each scenario under both the absorbing and the churn variant, with rows summing to one.

Table 21. Transition matrices — Pessimistic scenario (absorbing and churn variants).

Pessimistic — absorbing	N	P	A
N	0.80	0.20	0.00
P	0.00	0.70	0.30
A	0.00	0.00	1.00
Pessimistic — churn	N	P	A
N	0.80	0.20	0.00
P	0.00	0.70	0.30
A	0.05	0.00	0.95

Table 22. Transition matrices — Conservative scenario (absorbing and churn variants).

Conservative — absorbing	N	P	A
N	0.65	0.35	0.00
P	0.00	0.50	0.50
A	0.00	0.00	1.00
Conservative — churn	N	P	A
N	0.65	0.35	0.00
P	0.00	0.50	0.50
A	0.05	0.00	0.95

Table 23. Transition matrices — Optimistic scenario (absorbing and churn variants).

Optimistic — absorbing	N	P	A
N	0.40	0.60	0.00
P	0.00	0.25	0.75
A	0.00	0.00	1.00
Optimistic — churn	N	P	A
N	0.40	0.60	0.00
P	0.00	0.25	0.75
A	0.05	0.00	0.95

Appendix C Full Annual Bass Forecast Tables

The following table reports the full year-by-year Bass forecast of cumulative penetration (as a percentage of $m = 1,128$) for the overall sample and each functional area, from 2025 to 2040.

Table 24. Full annual Bass forecast of cumulative penetration (% of 1,128), 2025–2040.

Year	Overall	Marketing & Sales	Production & Operations	R&D & Innovation	Logistics	Corporate Management	Human Resources	Finance	ICT Security
2025	66%	29%	18%	16%	14%	14%	13%	10%	10%
2026	78%	40%	27%	25%	22%	20%	20%	15%	15%
2027	88%	53%	38%	35%	31%	28%	30%	22%	21%
2028	94%	65%	50%	47%	41%	38%	41%	31%	29%
2029	97%	76%	62%	59%	53%	49%	54%	42%	39%
2030	99%	84%	73%	70%	64%	60%	66%	54%	49%
2031	99%	90%	81%	79%	74%	71%	77%	65%	60%
2032	100%	94%	88%	86%	82%	79%	85%	74%	69%
2033	100%	96%	92%	91%	87%	85%	90%	82%	77%
2034	100%	98%	95%	94%	92%	90%	94%	88%	84%
2035	100%	99%	97%	97%	95%	93%	96%	92%	89%
2036	100%	99%	98%	98%	97%	96%	98%	95%	92%
2037	100%	100%	99%	99%	98%	97%	99%	97%	95%
2038	100%	100%	99%	99%	99%	98%	99%	98%	97%
2039	100%	100%	100%	100%	99%	99%	100%	99%	98%
2040	100%	100%	100%	100%	99%	99%	100%	99%	99%

Appendix D Technical and Reproducibility Notes

All estimation and figure generation were carried out in a scripted Python environment using standard scientific libraries (pandas, NumPy, SciPy, Matplotlib). The Bass parameters were obtained by non-linear least squares (Levenberg–Marquardt / trust-region) fitting of the cumulative adoption curve with the market potential fixed at $m = 1,128$; confidence intervals were derived from the covariance matrix of the fit using the Student t-distribution with degrees of freedom equal to the number of fitted points less the two estimated parameters. The Markov projections were obtained by iterating the initial state vector through the transition matrix over annual steps, and the stationary plateaus were verified against the closed-form expression in Equation 7. The analysis is fully reproducible from the prepared dataset and scripts, consistent with open-science principles; the same prepared series can be re-estimated independently in the an independent Bass forecasting module for triangulation.